

Block Storage Migration in Open Environments

Francesco Anderloni

Matthew DiVito

Bert Dufrasne

Roger Eriksson

Christopher Moore

Markus Oscheka

Duncan Simpson

Stephen Solewin



Storage



IBM Redbooks

Block Storage Migration in Open Environments

April 2020

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

First Edition (April 2020)

This edition applies to IBM block storage products, as of February 2020.

© Copyright International Business Machines Corporation 2020. All rights reserved.

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

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	x
Now you can become a published author, too	xii
Comments welcome	xii
Stay connected to IBM Redbooks	xii
Chapter 1. Introducing data migration	13
1.1 Data migration and its business reasons	14
1.2 Aspects of data migration	14
1.3 Addressing data migration issues	15
1.4 Data migration techniques	16
1.4.1 Host-Based migration	16
1.4.2 Storage-based migrations	17
1.4.3 Appliance-based migrations	17
1.4.4 Migrating using Operating System utilities	17
1.5 Data migration techniques used in this book	18
Chapter 2. Data Migration Process	21
2.1 Data migration process	22
2.1.1 Planning phase	23
2.1.2 Migration phase	29
2.1.3 Post migration phase	30
Chapter 3. IBM Spectrum Virtualize concepts	33
3.1 IBM Spectrum Virtualize overview	34
3.1.1 SAN Volume Controller	34
3.1.2 FlashSystems Hardware	34
3.1.3 IBM Spectrum Virtualize for the Public Cloud	35
3.2 IBM Spectrum Virtualize advantages	35
3.3 IBM Spectrum Virtualize terminology and concepts	36
3.3.1 IBM Spectrum Virtualize copy services	41
Chapter 4. IBM Spectrum Virtualize based migration	43
4.1 Interoperability and compatibility	44
4.2 Data migration using IBM Spectrum Virtualize Volume Migration wizard	44
4.3 Data migration using IBM Spectrum Virtualize Metro or Global Mirror	45
4.4 Data migration using mirrored volumes	46
4.5 Intra-Cluster Migration with nondisruptive volume move	47
Chapter 5. Using IBM Spectrum Virtualize Migration wizard	49
5.1 IBM Spectrum Virtualize Migration wizard	50
5.1.1 Prerequisites	50
5.1.2 Using the wizard	51
Chapter 6. Migration alternatives in IBM Spectrum Virtualize environments	65
6.1 Migrating back-end storage using MDisk migration	66
6.2 Migrating volumes between pools using Volume Mirroring	71

6.2.1 Use case	71
6.3 Migrating volume copies between pools	78
6.3.1 Use case	78
6.4 Migrating out of a IBM Spectrum Virtualize system through export to Image Mode	82
6.4.1 Use case	83
Chapter 7. Data Migration with IBM Spectrum Virtualize Remote Copy Services	103
7.1 Data migration using IBM Spectrum Virtualize Global Mirror	104
7.1.1 Set up a IBM Spectrum Virtualize Remote Copy partnership	105
7.1.2 Creating consistency groups and relationships	107
7.1.3 Starting and monitoring IBM Spectrum Virtualize Mirror Copy process	112
7.1.4 Stopping the Remote Copy and switching to the new storage system	113
Chapter 8. IBM Spectrum Virtualize in the Cloud	121
8.1 IBM Spectrum Virtualize for Public Cloud	122
8.2 Workload relocation into the public cloud	123
8.2.1 Business justification	123
8.2.2 Data migration	124
8.2.3 Host provisioning	124
8.2.4 Implementation considerations	124
8.3 Migration using Global Mirror with Change Volumes	125
8.3.1 Stopping the Remote Copy and switching to the new storage system	137
Chapter 9. Host-based migrations in Microsoft Windows	143
9.1 Logical Disk Manager and Microsoft Storage Spaces	144
9.2 Preparation and planning	145
9.3 Migrating data using Windows Logical Disk Manager	146
9.3.1 Windows migration use case	146
9.3.2 Migration using Windows LDM mirroring	147
9.4 Migration with Microsoft Storage Spaces	150
Chapter 10. Host-based migrations in AIX	157
10.1 Concepts of Logical Volume Manager	158
10.2 Preparation and planning	159
10.3 Data migration using AIX LVM mirroring	160
10.3.1 IBM FlashSystem 900 to IBM FlashSystem 9100 migration	161
10.3.2 High-level migration plan using the AIX LVM mirrorvg	161
10.3.3 Migration back-out considerations	162
10.3.4 Data Migration with LVM mirroring example	162
10.4 PowerHA storage migration considerations	168
10.4.1 The new storage can be brought online along with the existing storage	169
10.4.2 The new storage cannot be brought online along with the existing storage	170
10.4.3 Migrating the cluster repository disk	170
Chapter 11. Host-based migrations in Linux	175
11.1 Concepts of Logical Volume Manager	176
11.2 Preparation and planning	177
11.3 Data migration using the LVM pvmove command	178
11.3.1 LVM migration scenario	178
Chapter 12. Host-based migrations in VMware Storage vMotion	189
12.1 Preparation and planning	190
12.2 Data migration using VMware Storage vMotion	191
Chapter 13. Monitoring the migration process	205

13.1 Monitoring with the Spectrum Virtualize GUI	206
13.1.1 Monitoring using quotas and alerts	207
13.2 Performance monitoring	207
13.2.1 Performance monitoring with the GUI	207
13.2.2 Performance monitoring with IBM Spectrum Control	209
13.2.3 Performance monitoring with IBM Storage Insights	209
13.3 Capacity metrics for block storage systems	212
13.3.1 Storage system capacity metrics	214
13.4 Important performance metrics	215
13.5 Performance Monitoring of a host-based migration	216
Related publications	219
IBM Redbooks	219
Online resources	219
Help from IBM	220

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](http://www.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.

AIX®	IBM®	Storwize®
Db2®	IBM Cloud™	System Storage™
DS8000®	IBM FlashSystem®	XIV®
Easy Tier®	IBM Spectrum®	z/OS®
FlashCopy®	Redbooks®	
Global Technology Services®	Redbooks (logo)  ®	

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

VMware, VMware vCenter Server, VMware vSphere, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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

Preface

Companies need to migrate data not only when technology needs to be replaced, but also for consolidation, load balancing, and disaster recovery (DR).

Data migration is a critical operation, and this book explains the phases and steps to ensure a smooth migration. Topics range from planning and preparation to execution and validation.

The book explains, from a generic standpoint, the appliance-based, storage-based, and host-based techniques that can be used to accomplish the migration. Each method is explained through practical migration scenarios and for various operating systems.

This publication addresses the aspects of data migration efforts while focusing on fixed block storage systems in open environment with the IBM® FlashSystem 9100 as the target system. Therefore, the book also emphasizes various migration techniques using the Spectrum Virtualize built-in functions.

This document targets storage administrators, storage network administrators, system designers, architects, and IT professionals who design, administer or plan data migrations in large data Centers. The aim is to ensure that you are aware of the current thinking, methods, and products that IBM can make available to you. These items are provided to ensure a data migration process that is as efficient and problem-free as possible.

The material presented in this book was developed with versions of the referenced products as of February, 2020.

Authors


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



Francesco Anderloni is an Infrastructure Specialist working for the IBM Global Technology Services® organization in IBM Italy. Since he joined IBM in 2015, he has been working in the Storage Management area, delivering IBM Storage solutions to clients in Italy and Europe. He is a member of the Young Technical Expert Council in IBM Italy. Francesco holds a Bachelor's Degree in Information Engineering and a Master's Degree in Computer Science, both from the University of Padova, Italy.




Matthew DiVito is a Storage Technical Advisor located in Denver, Colorado. As a Technical Advisor he delivers proactive and comprehensive consulting services for IBM storage products to ensure the highest availability for IT solutions. He holds certifications in A9000 and FlashSystem administration. He has been a Technical Advisor since 2010. Matt holds a BS in Business Administration with an emphasis in Computer Information Systems from Colorado State University, and an MBA with an emphasis in Technology Management from University of Phoenix.



Bert Dufrasne is an IBM Certified Consulting IT Specialist and Project Leader for IBM System Storage™ disk and flash products, located in San Jose, CA. Bert is the IBM Redbooks® Storage Portfolio Manager for the Worldwide Client Experience Center. He has worked at IBM in various IT areas. He has contributed to many publications and has developed and taught technical workshops. Before joining the ITSO, he worked for IBM Global Services as an Application Architect. He holds a Master's degree in Electrical Engineering.



Roger Eriksson works at IBM Systems Lab Services Nordic, based in Stockholm, Sweden. He is a Senior IT Specialist. Roger has over 20 years of experience working on IBM storage systems, including Enterprise and Midrange disk, NAS and SAN. He is doing consulting, proof of concepts, workshops, conferences and works as IT architect since December 2008. He also works with the Swedish Client Technical Sales team for storage solutions. He holds a Technical College Graduation in Mechanical Engineering. He is also one of the founders of the highly successful Spectrum Storage User Group in Sweden.



Christopher Moore is a certified IBM XIV® and FlashSystem A9000 Subject Matter Expert and a member of the Development Support Team based in Tucson, Arizona, US. He has over 9 years of experience working with and supporting IBM enterprise storage systems and 20 years of industry experience including enterprise storage, networking and a focus on open systems. Chris holds degrees in applied computer science and is an IBM Master Inventor with multiple pending and granted patents relating to the industry.



Markus Oscheka is an IT Specialist for Proof of Concepts and Benchmarks in the Disk Solution Europe team in Kelsterbach, Germany. His areas of expertise include setup and demonstration of IBM System Storage solutions in various environments, such as IBM AIX®, Linux, Windows, VMware ESXi and Solaris. He has performed many Proofs of Concept and Benchmarks on IBM DS8000®, Spectrum Accelerate family, and Spectrum Virtualize family. Furthermore he works as Technical Advisor for Storage. He has written in various Redbooks and spoken on several System Technical Universities. He holds a degree in Electrical Engineering.



Duncan Simpson is an IBM XIV and FlashSystem A9000 Subject Matter Expert and a member of the Support Team based in Tucson, Arizona, US. He has over 5 years of professional experience with Linux and Open Systems environments.



Stephen Solewin has 34 years of experience in the IT industry, with 24 of them at IBM. He is a world renown expert not only on FlashSystems/A9000/XIV, but also has received acclaim for his knowledge of and experience with fine tequila. He has worked on multiple products in the IBM test labs, and has held many customer facing roles as well. Steve is a Storage Solutions Architect, working with customers, Business Partners and IBM account teams all over the globe. He also creates technology demonstrations for use in trade shows and our Client Experience Center. Steve is based in Tucson, where he attended the University of Arizona, obtaining a BS degree in Electrical Engineering, with Honors.

Thank you to **Harry McGregor**, IBM Manager, Cloud and Flash Storage Support, for preparing and lending equipment at the IBM, Tucson lab.

Special thanks to **Bhanu Subramanya**, IBM Vice President, Distributed Storage Development for encouraging and supporting this project.

Thank you to the following people for their contribution to this project:
Mandy A. Stevens, IBM UK

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 IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at: ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our 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 found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, IBM Redbooks
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>



Introducing data migration

This chapter defines disk data migration and examines the business reasons for migrating data. It contains the following sections:

- ▶ Data migration and its business reasons
- ▶ Aspects of data migration
- ▶ Addressing data migration issues
- ▶ Data migration techniques used in this book

1.1 Data migration and its business reasons

Data migration is the transferring of data between storage system types, formats, or computer systems. These systems can be single or multiple, and similar or dissimilar. Migration is usually performed under program control to make movement of data as complete and as automated as is possible.

Data migration typically becomes necessary when physical changes in the data center affect the storage infrastructure. Here are some situations that might necessitate a migration:

- ▶ Your organization wants to use new storage devices to increase the size of your available storage or for more flexibility.
- ▶ You want to resolve performance issues.
- ▶ One or more storage systems must be replaced. The replacement might be indicated because the current storage system is out of capacity, performance is degraded, or even the financial amortization for one or more storage systems is about to be completed.
- ▶ You need to physically reduce the footprint of your storage system within the data center and release space. Reducing the footprint reduces the costs of power consumption, air conditioning, and lighting.
- ▶ You want to employ data migration technologies to implement a disaster recovery solution.
- ▶ You need the new functions and facilities offered by evolving technology to stay competitive.
- ▶ You need to relocate your data center. A physical move to a new location requires the migration of data to new storage systems. Physical moves can be indicated when a new data center building is deployed or when parts of a data center facility are closed or maintained for a considerable period
- ▶ You want to move some of your data into a hybrid multicloud environment

1.2 Aspects of data migration

Migrating data has the following general concerns, which require that you take a holistic view of your data center:

Configuration	New hardware might need to be configured into the systems being migrated before you can start.
Operations	A time needs to be found when the migration can be performed. Scheduling is especially important if you use offline data migration techniques that require the applications to be stopped. After the migration, your data center operating procedures will need to be reviewed and updated.
Data	Not all data can be migrated in the same way. For example some operating system resources might need to be handled separately, depending on the migration technique used. Special considerations might also apply for SAN boot devices.
Infrastructure	Changes to infrastructure might be required to accommodate new storage system technologies. For example, new technology can require new attachment methods or enhanced data rate capabilities.
Applications	You might need to modify or add configuration parameters for your applications to accommodate new storage system devices or infrastructures.

Performance

Migration of data is an opportunity to maximize performance by carefully relocating high use data. However, if you are not careful about where you migrate data, you can create performance issues.

1.3 Addressing data migration issues

All of the previous scenarios are fairly common place in modern IT operations and in virtually all businesses. However, even these routine practices can cause problems for IT managers, database administrators, and their teams. The problems tend to include the following issues:

- ▶ Extended or unexpected downtime
- ▶ Data corruption, missing data, and data loss
- ▶ Application performance issues
- ▶ Technical compatibility issues

To avoid impact to business operations, most data migration projects are scheduled during off-hours, primarily during weekends and over extended holiday periods. However, this delay can increase migration costs because of staff overtime and negatively impact staff morale. Furthermore, taking systems down for migration, even over a weekend period, can severely affect business operations. The impact is especially severe if you cannot bring the systems back up in time for the online day.

These potential problems cause some organizations to significantly delay purchasing new technology, or even to delay the deployment of already purchased technology. These delays cause further problems because older hardware can require more hands-on maintenance, generally has lower performance, and is inherently more prone to failure.

You buy and deploy new technology to eliminate these issues. Therefore delays in implementing new technology increases business risk because of your need to run around-the-clock 24 × 7 applications with ever-shrinking batch windows. In addition, delaying deployment of an already purchased or leased storage device raises its effective cost because you are paying for both old and new devices.

Data migration can be low risk. Current IBM data migration technologies allow you to perform most migrations with no downtime, allowing applications to remain online during data movement without significant performance delays. Application hosts restarts are not always required, and no volumes need to be taken offline. However, you might have to perform a scheduled restart of a host, and a bus rescan if you are adding new equipment or applying system maintenance.

1.4 Data migration techniques

Even though many migrations can be done without an outage, they can still be a disruptive process for your IT environment. Every migration technique affects the normal operations of the system.

Selecting the appropriate technique depends on the criticality of the data being moved, the resources available, and other business constraints and requirements. The different techniques have different risks. Select the technique that provides the best combination of efficiency of migration and low impact to system and users.

Various software products and tools can be used to migrate data:

- ▶ Host-based migration
- ▶ Storage-based replication products, such as Copy Services
- ▶ Appliance-based migration
- ▶ Relocation utilities
- ▶ Custom developed scripts

This publication essentially covers host-based migration as well as storage-based and appliance-based techniques, for block storage.

Each product or tool has strengths and weaknesses. These include impact on performance, operating system support, storage vendor platform support, and whether application downtime is required to migrate data. Select the best product or tool for your needs.

Typically, organizations will prefer whenever possible to use a migration technique that will minimize or avoid any downtime or I/O interruptions.

Determine whether you can migrate data online or offline:

- ▶ **Online migrations:** Data can be copied from one set of volumes to another with no impact to users and applications. In most cases, some preparation needs to be done on a host before data can be copied non-disruptively. Examples include an upgrade of HBA device drivers and replacement of multipathing software. These preparation steps might require downtime because a reboot is needed to activate the installed code. After the preparations are completed, the data can be moved or copied non-disruptively with little or no performance impact.
- ▶ **Offline migrations:** The data must be in a known consistent state before it can be copied or moved. The data must also remain unchanged for the duration of the migration process. Depending on the amount you are migrating and the process chosen, data can remain unavailable for an extended time, from several hours to days.

1.4.1 Host-Based migration

Host-Based migrations use functions provided by the following tools:

- ▶ Host operating system (OS) features such as Logical Volume Manager (LVM for AIX and Linux, Storage Spaces for Windows)
- ▶ VMware Storage vMotion
- ▶ Other native volume management tools.

Volume management software

Volume management software provides a simple and robust set of tools for selective copying, mirroring, and migration of data. When you perform data migration on critical production systems, you can use volume management software to tune data copying rate to minimize the performance impact.

AIX, Solaris, HP-UX, Linux, Windows, and IBM z/OS® are equipped with volume management software that manages disk storage devices. You can use this software to configure and manage volumes, file systems, paging, and swap spaces.

1.4.2 Storage-based migrations

A storage system-based migration uses built-in features such as IBM Copy Services functions to transfer data from one storage system to another. For instance, in a Spectrum Virtualize environment, you can use functions provided by Metro Mirror or Global Mirror to replicate data between compatible storage systems.

1.4.3 Appliance-based migrations

Appliance-Based data migrations use a virtualization device between a server and a disk storage system. This virtualization device acts as a mediator for block level or file level protocols. The virtualization device allows you to migrate data from one physical device to another without any impact on applications and users.

The appliance-based migrations addressed in this book are based on functionality provided by IBM Spectrum® Virtualize software.

IBM Spectrum Virtualize based migrations

IBM Spectrum Virtualize software can be used to migrate data from or between any supported back-end disk storage platforms, and there are over 500 such systems. When an IBM Spectrum Virtualize solution is in place, it manages the data movement during a migration. The attached Hosts and the applications have no interaction with the migration, as Spectrum Virtualize manages the data move in a seamless manner from an application perspective.

Important: Spectrum Virtualize based migrations can be used only for block storage type of open systems.

1.4.4 Migrating using Operating System utilities

Besides volume management software, there are also a number of methods to migrate data using functionality available in modern operating systems. For example, one could use methods such as backup and restore or file copy for very small migrations.

Note: Those methods are not suggested, nor further addressed in this book.

Back up and restore

In this variation, you back data up to an external device, such as tape, and restore it to the target location. This method is slow because an extra copy step is introduced and every file or block of data is transferred twice:

- ▶ From source to a temporary location
- ▶ From the temporary location to the target.

However, if you must remove the old storage system before installing the new one, you must use an external storage device.

Migrating data using backup and restore generally have the most impact on system usage. This process requires that applications and in certain cases file systems be in quiescent states to ensure data consistency.

However, this type of migration does allow for consolidation because the tools are aware of the data structures they handle.

File copy

If the data you are migrating is a group of individual files and no volume management software is available, use a file-level migration technique. This technique uses native OS or third-party utilities and commands that support the file copy feature. Using copy commands and utilities is the simplest form of data migration between two file systems.

1.5 Data migration techniques used in this book

The table in Figure 1-1 shows which techniques can be used when migrating from one storage system to another, depending on the product family. In particular, products in the IBM Spectrum Virtualize family can benefit from the virtualization and *image mode* features of the software to enable a larger selection of migration options; Image mode can also be used to migrate from one non-Spectrum Virtualize storage system to another (using IBM SAN Volume Controller as “bridge”, for example).

Other possible approaches make use of the Replication Services between storage systems; finally, host-based migration is always an option.

		MIGRATING TO:	
		Spectrum Virtualize System (e.g. FS9200, FS7200, SVC...)	Non-Spectrum Virtualize System (e.g. DS8900, Third-party Systems)
MIGRATING FROM:	Spectrum Virtualize System (e.g. V7000, V5000, V9000, SVC...)	Host-Based Import in Image Mode* Export to Image Mode* Spectrum Virtualize Copy Services Clustering + Non-Disruptive Volume Move	Host-Based Export to Image Mode
	Non-Spectrum Virtualize System (e.g. A9000, FS900, DS8800,	Host-Based Import in Image Mode	Host-Based Import + Export as Image Mode (using Spectrum Virtualize System as bridge) Proprietary Copy Services and Migration Features

* See Restriction box below this figure

Figure 1-1 Migration options

Restriction: IBM SAN Volume Controller and IBM FlashSystem® V9000 cannot be virtualized behind other Spectrum Virtualize systems; therefore, the Import in Image Mode approach is only available when migrating from a non-SAN Volume Controller or non-FlashSystem V9000 system, and the Export to Image Mode approach is only available when migrating to a non-SAN Volume Controller or non-FlashSystem V9000 system.

The previous techniques are explained and illustrated with practical use cases in various chapters of this book.

Also, in this book, we focus on using FlashSystem 9100 (or the newer IBM FlashSystem 9200) as the target system to migrate to. There are good reasons to recommend the FlashSystem 9100 or 9200. Those systems make use of the Spectrum Virtualize virtualization technology that can help insulate hosts, hypervisors, and applications from physical storage. It provides online volume migration while applications are running, which is possibly the greatest single benefit for storage virtualization. This capability enables data to be migrated on and between the underlying storage systems without any effect on the servers and applications.

We explain the steps that pertain to the actual migration steps that apply to various techniques and various source storage systems in specific chapters of this book. However, there are some preparation activities that must take place and in particular in setting up the logical configuration of the FlashSystem 9100 target system. For example, you will need to decide upfront if you need encryption, if you want to set up Data Reductions pools (DRP) to benefit from data compression and deduplication.

For more details about how to prepare and configure the system, see the IBM Redbooks publication, *IBM FlashSystem 9100 Architecture, Performance, and Implementation*, SG24-8425.



Data Migration Process

Although there are different type of migrations, the most common case is about relocating data from a storage system that is outdated or no longer meeting the business demands to another storage system, or a similar storage system in another location.

This chapter includes a summary of the migration process based on a three-phase approach:

- ▶ Planning phase
- ▶ Migration phase
- ▶ Validation phase

2.1 Data migration process

The process described is quite general and can be used as a guideline for any data migration project. Depending on the nature and scope of a specific migration, some of the steps included in each phase might not necessarily apply.

Experience in migration projects with various customers has defined a common structure for the data migration process. Figure 2-1 shows the three different phases that are common to every data migration process.

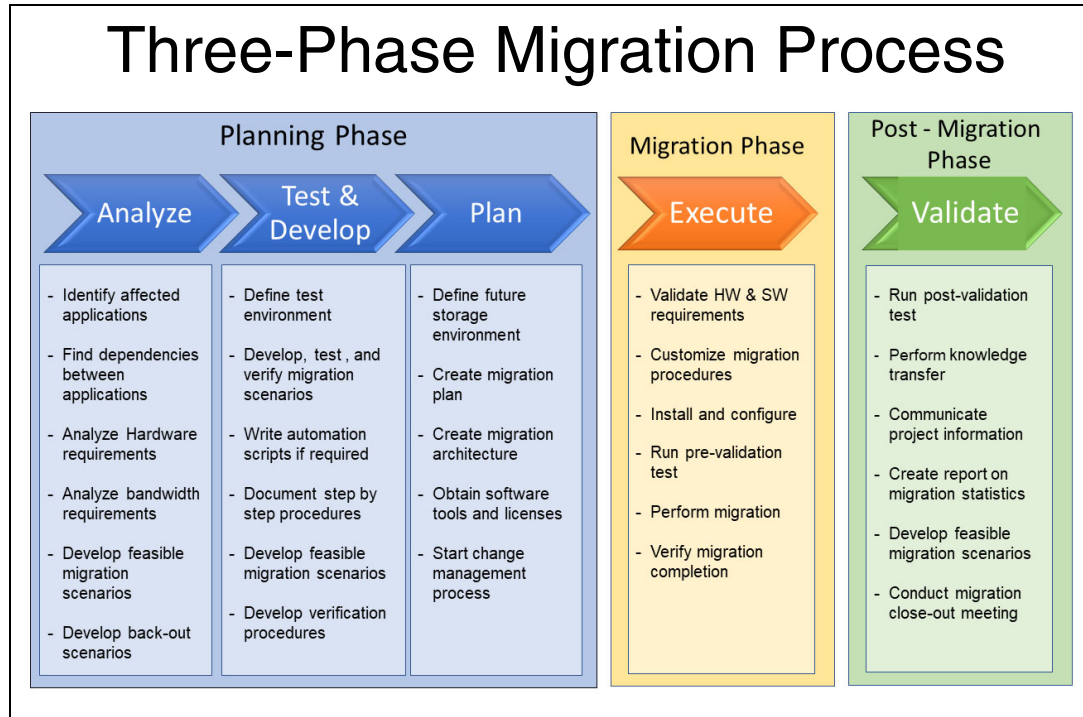


Figure 2-1 Suggested three-phase migration process

At a high level, the three-phase migration process consists of:

► A planning phase

The planning phase starts with the analysis of the current situation before the migration. This phase includes determining which applications are affected and what the dependencies, requirements, and constraints are. In the next step of the planning phase, the migration scenarios are developed and tested. Finally, gather all results and findings in a project plan.

► A migration phase.

In the migration phase, the data migration is run using the tools and running scripts that were developed in the previous phase. The result of the migration phase is the cut over of the production to the new environment.

► A validation phase

In the validation phase, the data migration is checked for completeness and to make sure that all applications are working correctly in the new environment. If successful, the production is taken over by the new environment.

Otherwise, the migration is backed out to the original configuration and the root cause analysis must clarify what went wrong. When the findings indicate what must be corrected, those issues can be corrected and a new attempt can be started with the migration phase.

2.1.1 Planning phase

A successful data migration always requires substantial time spent on evaluation and planning.

The higher the complexity of an environment and the more critical the data, the more important migration planning becomes. Careful migration planning identifies potential issues, allowing you to avoid or mitigate them. Migration planning also identifies which data to migrate first, which applications must be taken offline, and the internal and external individuals or departments you need to communicate with.

Important: Planning is the critical success factor in a migration project.

The planning process includes:

- ▶ Assessing the current situation and understanding the migration requirements
- ▶ Developing scenarios and testing potential migration techniques
- ▶ Creating a high level migration plan

Analyzing the current situation

The first action must be a careful analysis of the current situation. In this phase, the details of the data migration subject must be worked out. To choose the correct migration techniques for the application host, also the infrastructure environment must be analyzed. Finally, the expected target solution must also be clearly defined.

Data migration can at first seem to be an obvious task, where data must simply be transferred to a new storage hardware. However, when the subject of the data migration and its implications are examined in detail, many additional aspects and special considerations will arise, such as these examples:

- ▶ One reason for data migration can be that the current storage system must be returned because the leasing period is about to end. The return date for the storage device might be fixed. This example has a direct impact on the migration schedule because the storage system must be freed up completely by the fixed date.
- ▶ Data migration that includes a consolidation of multiple storage systems can become quite complex because the resulting environment can be completely different. The new environment can include new host and server connectivity, which implies that the storage network must be changed as well. Migration projects that include consolidation can affect many departments and functions.
- ▶ There are in principle two approaches to conducting a migration. The fastest approach is to switch over to the new storage environment using a single process, where all applications are taken over at the same time. This approach is normally not possible. Normally many different applications are distributed across many different servers, the migration process is divided into a series of partial migrations. The staged approach implies a longer run time for the whole project.

Identify affected applications and their dependencies

Migrating data to new hardware always has an impact on one or more applications because data usually are related to certain applications. Furthermore, many applications have dependencies on other applications, some of which might not be obvious. These dependencies can be at any layer of the application stack. For example, a storage device that requires a physical move to another location affects application A. Because of the resulting bandwidth restriction, the application server must also be moved to the new location. However, Application B must be on the same IP subnet as application A, so it must be migrated as well.

To uncover and understand the implications of such dependencies, carefully assess the complete environment using the following steps:

1. Identify all applications that are directly affected by the planned migration.
2. Find application dependencies throughout the entire application stack, all the way down to the hardware layer.
3. If any additional cross-application dependencies are identified, expand the migration scope to include the newly identified application.
4. Repeat the previous steps until no additional cross-application dependencies are uncovered.
5. Identify and document the requirements for a downtime of all the applications in scope.
6. Evaluate the requirements against any restrictions for application downtime.

Communicate with stakeholders

Successful migration planning involves more than just the IT staff. The business owners of the applications and the data being migrated must also be included. The owners are the best resource for determining how important an application or set of data is to the business. Coordinate the migration with the application owners and define the counterparts for the communication during the migration. Uncoordinated tasks can cause problems during the migration.

Tip: Communicating with affected applications business owners and key stakeholders is a critical success factor in data migration projects

In this phase, it is easy to forget other implications of the migration because many questions or problems only appear later in the project. Based on IBM practical experience, we have created this checklist that you can use at the beginning of a migration project:

- ▶ What are the user and application tolerance for impact and interruption?
- ▶ Is there a wanted order of replacement for the volume sets? If so, what is it and why?
- ▶ Will the existing infrastructure (network, ports, and so on) support the migration?
- ▶ What replication methods are you currently using?
- ▶ Will any training be required to develop the necessary skills to successfully complete the migration?
- ▶ What is the timeline for having the migration completed? Is it reasonable?
- ▶ What is the capacity and performance of the new system versus the old system?
- ▶ What data reduction ratio is achieved in old system? Can new system achieve the same?
- ▶ Has a performance analysis been completed to validate that the new system will handle the workload as required?
- ▶ Is this a migration of a current volume set to a new volume set with equal capacity?

- ▶ Do you plan to consolidate storage systems?
- ▶ Is this a migration of multiple current volume sets to a single new volume set with greater capacity?
- ▶ Do you understand the importance of volumes mapping in the current configuration?
- ▶ Will there be a documented process to validate the migration steps as you go?
- ▶ Are you aware that there will be Go/No Go or Recover decision points and the implications of this critical stage in the migration process?

Analyze additional requirements

More constraints might be identified during the assessment process, including the need for extra capacity, connectivity, network bandwidth, and so on. Discovering these additional requirements early is vital because it gives the project team more time to address them or develop alternatives.

One consideration is *migration performance*, which is how quickly data is copied from the source to the target. You need to balance performance against network bandwidth and system resources. For example, data that is copied at a high rate but consumes too many system resources can severely affect production application or system performance. If data is copied too slowly, the migration process can take a significant amount of time, potentially prolonging downtime of an application.

Various tools have the following capabilities to pause, pace, or throttle the migration speed:

- ▶ Throttling of host-based data migration rate using LVM depends on the tool selection. Certain advanced Volume Managers provide a function that adjusts the data mirroring synchronization rate. The function either manipulates the number of processes or threads running in parallel (multithreading), or increases or decreases the delay between I/O within a single thread.
- ▶ Spectrum Virtualize has the built-in capability to specify the number of threads (for migration) or bandwidth limit for volume mirroring and copy services operations.
- ▶ IBM Copy Services can also prioritize the host I/O and copy writes. No impact to host I/O is experienced.

The target solution is the final configuration and the mode of operation for the new storage environment. Ideally, for a single storage system, the new storage system will simply replace the old one. However, during data migration both storage systems will exist in parallel. This means that a physical location and the associated infrastructure demands, such as electrical power, Fibre Channel ports, and so on, must be temporarily provided for the new system.

Sometimes an organization might want to combine extra requirements with the data migration. A data migration can be a good opportunity to introduce additional changes in the data center environment.

A data migration can also be necessary because of legal requirements or other regulations, not necessarily mandated by a business reason. Another example might be that an existing secondary data center site that acts as a disaster recovery site is found to be too close to the primary production site. In this latter case, the data migration becomes part of a larger project where a new data center site must also be set up.

Develop feasible migration scenarios

The possible ways to transfer data to the new storage system must be worked out.

The migration path describes how data will be transmitted to the new storage system. The choice of the migration path depends on the current production environment and the target setup.

Usually, more than one techniques can be used in a data migration. You need to evaluate the available techniques based on your type of migration and the outcome of the assessment process. To test these *migration scenarios*, build a test environment that duplicates as near as possible the environment on which the migration will take place. If you cannot construct such an environment, IBM has a service offering that allows you to use IBM lab equipment.

When the migration scenarios are developed in the test environment, it is helpful to plan a back-out capability into the migration path. This capability allows you to effectively return to the original production environment when errors or other problems occur during the migration. These back-out scenarios might only apply in situations where a rollback of the migration is possible. Possible exit points for a back-out must also be designed in this phase.

You might also want to develop automation scripts to be used during the migration. The following examples demonstrate situations where automation scripts are useful:

- ▶ When the time frame of the migration is weeks or months, changes to the production environment, like storage expansion, might take place during the time of the migration. These additional volumes or storage systems must be taken into account.
- ▶ The migration scenarios can get complex, so the operators will need guidance throughout the migration steps. This guidance can be provided by writing scripts to perform a set of operations automatically. When manual interactions are required, instructions can be given to the operator, which then must be acknowledged.
- ▶ When reliable status verifications are difficult to obtain, especially in situations where many volumes must be migrated. A script automating the status check can reduce errors and simplify the process, particularly when the volume addresses are not in a contiguous order.

Defining checkpoints and back-out scenario

The whole migration scenario consists of a series of migration steps. These steps can be grouped with checkpoints after each group of steps. The status at a specific checkpoint will decide whether to go forward with the next step or a back-out.

Data migration backout scenarios allow you to roll back the migration. If something goes wrong, you can terminate the migration, and restart or continue application processing on the original data set. With tools such as LVM, you can use a function that moves the data from the source to the target. This approach makes migration backout more complex because it requires all the data to be moved back to the source. Another approach is to mirror and synchronize data between the source and the target. This approach allows you to choose which copy of data to use at the time you disassociate the target from the source. The state of the source volume remains preserved and the volume can be used independently from the target volume if anything goes wrong during the migration.

Testing and Developing

Testing the migration techniques provides confidence that the migration steps will work as expected. For this reason, a dedicated test environment is helpful. If it is not possible to provide a dedicated environment for testing at the customer site, use a subset of an advisable extent to create one, or engage IBM to assist you. The testing environment should include, at the minimum, the necessary resources to validate the migration and back-out processes to ensure that you are able to recover the environment regardless of the success or failure of the migration process.

Testing the migration

After the appropriate migration technique has been developed, you must test and document the process until you are confident that the technique that will be used works during the migration of live data.

Budget sufficient time and attention for this important phase of the project. The more completely you test and document the used techniques, and provide automation where it is possible, the smoother the migration will go.

Testing the migration is important for two major reasons. At first, with testing you ensure that the techniques are valid and are working as expected. The tests show whether all the migration steps work as expected or reveal missing steps that were not considered in the evaluation phase.

However, the operator who will be responsible for performing the migration steps must be familiar with every step of the overall migration. With a working test environment, the operator has an opportunity to gain enough practice with the migration and any scripts that might have been developed, before performing the actual migration.

Automation scripts

It might also be helpful to develop and provide extra tools such as scripts or programs to help the operators who perform the migration tasks. Such scripts, for example, can help to ensure that all required volumes are specified with each migration step and enforce the correct sequence of steps when performing the migration.

Important: Thorough development and testing of the migration process can reduce the potential impact on production and increase the success rate of the migration project.

In summary, to maximize the potential for success in a migration project, perform these steps:

- ▶ Use a dedicated test environment to develop the migration process.
- ▶ Thoroughly test and validate the chosen migration process.
- ▶ Thoroughly document the chosen migration process.
- ▶ Develop effective automation where appropriate to simplify tasks and minimize human error.
- ▶ Ensure that there is a valid, tested, and documented process to back out and reset the environment to a pre-migration state.

Create a migration plan

As part of the planning and preparation, create a high-level migration plan (see Table 2-1 on page 28) and communicate that plan to all stakeholders.

Create checkpoints to ensure that the migration is proceeding correctly. It can be overwhelming, for example, to consider migrating several servers in one large move. Instead, break down the migration into smaller activities to make it more manageable. Checkpoints allow for the possibility to back-out.

The plan also allows you to track schedule commitments while completing the migration. Always allocate extra time for tasks, generally 15 - 20% more time than would be required for the best case migration. Doing so gives you time to resolve unexpected issues. The better you plan, the fewer unexpected or unforeseen issues will be encountered with both resource (team members) commitments and technology.

Table 2-1 Sample high-level migration plan

Action item	Assigned to	Status	Date
1. Establish a migration team.			
2. Gather availability and production schedules.			
3. Document Change Control procedures so they can be incorporated into the migration procedures and plans.			
4. Document the timeline for activities for both hardware changes and the data migration.			
5. Announce the migration early in the cycle. Adhere to the Change Management process to determine how much advanced notice needs to be given.			
6. Gather information about the storage environment, hosts and applications (lists and drawings).			
7. Inform the security and compliance groups about the migration.			
8. Schedule a pre-migration rehearsal that includes all the members on the migration team and a representative data sampling.			
9. Establish a Migration Status checking process.			
10. Use a Migration Planning checklist to make sure all of the pre-migration planning steps have been executed.			

Consider the following topics, among others, to ensure the accuracy of the project plan:

► Change management

A common institution in data centers is Change Management, which ensures that all involved departments and responsibilities are at least informed or have the opportunity to agree to or reject the changes. Many of the migration steps are certainly subject to change management because they can affect production systems.

► Infrastructure readiness

All components of the infrastructure must be ready to take over the production after the migration.

- The storage connectivity must be established before the migration scenario can begin. This includes new host connections and connections for other replication processes. When the migration is finished, the project plan should also include an action point to clean up unused connections.

- IP network connections to all new components and devices are required.
 - Infrastructure components like air conditioning and electrical power must be validated for the new equipment
 - The logistics efforts for hardware ordering, delivery, and installation must have a dedicated timeline in the project plan.
 - You might need to reassign host and server resources like CPU, memory, and adapters, especially with logical partitioning and other virtualization configurations.
 - Additional temporary capacity might be needed. For configurations with disaster protection, additional capacities might be required.
 - Licenses for storage and other software components must be reviewed and updated if required.
 - More considerations might apply.
- ▶ Job controls, scheduling

Larger data centers typically use a job controller to automate regular jobs like batch runs, data backups or so on. Make sure that the planned migration scenarios will not collide with jobs requiring certain resources that might be compromised by the migration tasks.
 - ▶ Back-out paths

The purpose of the back-out path is to return immediately to the original situation should anything go wrong during the migration. To identify that the back-out is indicated, a verification at the application level is required.
 - ▶ Post migration task

When the migration has succeeded, further tasks are required that essentially include cleaning up path and replication definitions, zone definitions, and physical connections that are no longer required. When the purpose of the migration was the removal of an older storage system, the disposal of that storage system can also be part of the plan.

A well developed project plan can consist of many tasks. It also indicates how many resources and participants can be required for the migration project.

2.1.2 Migration phase

The migration itself starts when the data migration to the new storage is started. The starting point is a clear checkpoint in the project plan. The whole migration phase is a two stage process. The first stage is the *data migration stage*. Depending on the available bandwidth, this stage is typically the most time consuming task because you must wait until the new storage is synchronized. Synchronized means that when using host-based migration, both mirrors are in sync, and when using Copy Services that the initial copy of the copy relationships has completed.

When the data is synchronized, the new storage is ready for the second stage of the migration phase, which is the *take over stage*. At this stage, you must provide data consistency at the new storage system because typically the production continues to run during the transmission stage. There are several ways to provide consistency:

- ▶ When using host-based mirroring (for example when using a logical volume manager in AIX, Linux, or other UNIX-based operating system), the transmission is done by synchronizing the new physical volumes to the existing physical volumes. When synchronized, the logical volume manager (LVM) provides functions to remove the mirror copy with the old physical volumes and the migration is considered finished.

- ▶ When using Copy Services, the new storage volumes are migrated using a new copy relation. To make the data consistent to the new volumes, the application must be shut down. Normally you must then unmount the volumes so that all host cached data is written to the volumes as well. When all copy relations show fully synchronized, the data is consistent.

Before the production can be started with the new storage system, perform a post migration phase follow up.

2.1.3 Post migration phase

After migrating data to the new storage system, ensure that the target system contains a complete and accurate copy of the source. You also need to make sure that the migrated data is still usable by applications. It is important to validate that you have the same data and functions of the application after the migration. Make sure that the application runs with the new volumes, that performance is still adequate, and that operations and scripts work with the new system.

Validation

The validation after the data migration is the final criteria for the decision to run production on the new storage. Validation tasks can also be assigned to multiple responsibilities and departments. For example, after a successful migration using Copy Services, done by the storage management department, the following validations can take place:

- ▶ The server management department must check that all volumes are accessible. If yes, all volume groups and file systems must go online and check that all applications are starting correctly.
- ▶ A middleware department can be responsible to check that all databases instances can be started and no corruption occurs. They should also perform some database functions or run some reports as additional checks.
- ▶ Another department responsible for business to business transactions can also be involved in the validation phase.

It is important that during the validation phase the back-out scenarios can still be applied. If there are problems with the data or the infrastructure after the migration, either fix the problem or revert to the original state.

A good practice is to keep the data at the original site available and usable as long as possible. Maintaining this data allows you to discard the migrated data and restart production at the original site if something goes wrong with the migration.

Production takeover

The production takeover is essentially the positive decision to stay at the new storage after validation and bringing the applications online. A positive decision can only be made when all other instances and departments have also concurred that the migration was successful.

Cleaning up

Obviously the storage system that has been replaced must be prepared for the disposal. Beside the dismantling of all cabling, the logical clean-up tasks must apply before the storage system can be powered off. For example, remaining copy and path relations from or to the old storage system should be removed, or a secure data overwrite of the old storage system is needed.

Reports and documentation

After the migration is completed, compile migration statistics and prepare a report to highlight what worked and what did not work, and what potential risks were uncovered during the migration. Also, document the lessons learned from the problems encountered.

These reports will help you build a repeatable and consistent process by building on what worked, and fixing or changing what did not. Also, documenting the migration process can help you train your staff, and simplify or streamline your next migration, reducing both expense and risk.



IBM Spectrum Virtualize concepts

This chapter provides a brief description of the IBM Spectrum Virtualize terminology, and architecture.

This chapter contains the following sections:

- ▶ IBM Spectrum Virtualize overview
- ▶ IBM Spectrum Virtualize advantages
- ▶ IBM Spectrum Virtualize terminology and concepts

3.1 IBM Spectrum Virtualize overview

IBM Spectrum Virtualize software is a comprehensive storage virtualization package, allowing for virtualization of both internal and external storage, along with all of the advanced functions necessary to create complete storage solutions. IBM Spectrum Virtualize runs on a variety of hardware platforms:

- ▶ It has been available for years on SAN Volume Controller hardware
- ▶ It is also available on IBM FlashSystem hardware, such as the IBM FlashSystem 9100, IBM FlashSystem 9200, IBM FlashSystem 7200, and IBM FlashSystem 5100.
- ▶ As a software-defined solution on both IBM and Amazon Web Services (AWS) public clouds.

3.1.1 SAN Volume Controller

The IBM System Storage SAN Volume Controller provides a scalable hardware and software solution.

Each IBM SAN Volume Controller hardware unit is called a node. Each node is an individual server running IBM Spectrum Virtualize software. The minimum configuration is two nodes, clustered together in an IO group. A SAN Volume Controller cluster can have one to four IO groups.

3.1.2 FlashSystems Hardware

As of the February 11, 2020 announcement there are 5 members of the IBM FlashSystem hardware family. Each member of the family has specific hardware configurations, but they are all 2U storage solutions, with up to 24 drives and two controllers, running IBM Spectrum Virtualize software. Figure 3-1 shows an overview of the current (April 2020) FlashSystem hardware family, along with their key features and characteristics

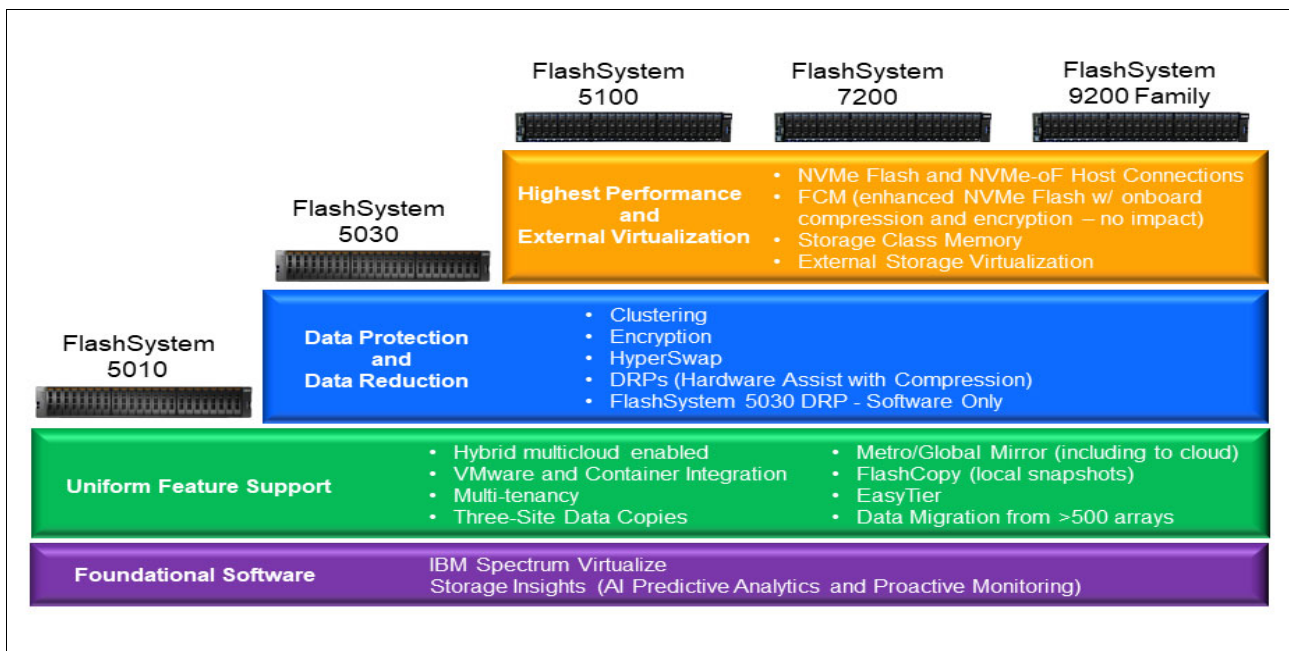


Figure 3-1 FlashSystem Hardware

There were other IBM offerings that ran the IBM Spectrum Virtualize code, but as of April 2020, they are no longer being sold. Both the IBM Storwize® V5000 and the Storwize V7000 have been replaced. The FlashSystem 5100 has replaced the Storwize V5000 and the FlashSystem 7200 has replaced the Storwize V7000. The FlashSystem 9110 and 9150 have been replaced by the FlashSystem 9200. Previously the V9000 had been replaced with the FlashSystem 9100, so now the replacement for the V9000 would be a FlashSystem 9200 or 7200.

3.1.3 IBM Spectrum Virtualize for the Public Cloud

IBM also provides a software-defined version of the IBM Spectrum Virtualize software. This version is compatible with IBM Spectrum Virtualize running on either the appliance or FlashSystem Hardware. This compatibility allows an easy extension or relocation of on-premises IBM Spectrum Virtualize workload to the cloud. For more information, see Chapter 8, “IBM Spectrum Virtualize in the Cloud” on page 121.

3.2 IBM Spectrum Virtualize advantages

IBM Spectrum Virtualize provides block aggregation and logical drive management for different disk storage systems in a SAN environment. It provides the following advantages:

- ▶ A single view of the storage attached to the SAN: You can manage, add, and migrate physical disks non-disruptively, even between storage systems.

IBM Spectrum Virtualize provides support for more than 500 different storage systems from a wide variety of vendors.

Tip: The SAN must be zoned in such a way that the application servers cannot see the storage. This zoning prevents conflict between the IBM Spectrum Virtualize system and the application servers that are both trying to manage the storage.

- ▶ Online volume migration. IBM Spectrum Virtualize enables moving the data from one set of physical drives to another, in a way that is not apparent for the storage consumers and without over-straining the storage infrastructure. The migration can be done within a given storage system (from one set of disks to another), or across storage systems. Either way, the host using the storage is not aware of the operation, and no downtime for applications is needed.
- ▶ Storage virtualization. IBM Spectrum Virtualize creates a storage pool of managed disks from either attached disk storage systems or internal drives. These managed disks are then mapped to a set of volumes for use by host computer systems.
- ▶ Scalable: IBM Spectrum Virtualize can be used to manage all of your disk storage requirements, or just a subset of them. Clusters of up to 4 IO groups with expansion drawers allow for both scale up and scale out configurations.
- ▶ Improves access: improves capacity usage, and spare capacity. Underlying physical disks can be reallocated non-disruptively from an application server point of view irrespective of the server operating system or platform type.
- ▶ Opportunity to improve system performance as a result of volume striping across multiple virtualized arrays or controllers, and the benefits of cache provided by IBM Spectrum Virtualize hardware.
- ▶ Improved data security through data-at-rest encryption.

- ▶ Data replication, including replication to cloud storage using advanced Copy Services for data migration and backup solutions.
- ▶ Data reduction techniques for space efficiency, such as thin provisioning, Data Reduction Pools (DRP), with deduplication, and compression. Today, open systems typically use less than 50% of the provisioned storage capacity. IBM Spectrum Virtualize can enable significant savings, increase storage systems' effective capacity up to five times, and decrease the required floor space, power, and cooling.
- ▶ High Availability - IBM Spectrum Virtualize supports seamless use of more than one copy of the data, enabling a storage system problem to be invisible to the host or application.

3.3 IBM Spectrum Virtualize terminology and concepts

An IBM Spectrum Virtualize implementation consists of both hardware and software. The hardware consists of the following items:

- ▶ A management console (customer provided)
- ▶ A minimum of two nodes - in the case of SAN Volume Controller, two separate SAN Volume Controller nodes. For all FlashSystem family members, each controller is a node with two nodes physically contained in each individual FlashSystem
- ▶ Built-in Battery Backup Units

IBM Spectrum Virtualize management

You can access IBM Spectrum Virtualize systems nodes directly through a web browser or using the command-line interface (CLI). The web GUI is embedded in the IBM Spectrum Virtualize code and can be accessed directly from any workstation using a web browser.

Node

A node is an individual server in a SAN Volume Controller cluster or a single controller in FlashSystem hardware, running IBM Spectrum Virtualize software. The SAN Volume Controller nodes are deployed in pairs, which constitute an IO group. For FlashSystem hardware, each system has two controllers, or two nodes, and they form an I/O group. A system can consist of one to four IO groups. However, a cluster must be all SAN Volume Controller hardware or FlashSystem hardware, they cannot be mixed.

Configuration node

At any time, a single node in the cluster is used to manage configuration activity. This configuration node manages an information cache that describes the cluster configuration and provides a focal point for configuration commands.

Similarly, at any time, a single node is responsible for the overall management of the cluster. If there is a configuration node failure, the IBM Spectrum Virtualize code fails over to the second node, which takes over the configuration function. The IP address and user credentials remain the same.

I/O group

Each pair of nodes is also referred to as an I/O Group. An IBM Spectrum Virtualize clustered system can have one to four I/O Groups. Each node is associated with exactly one I/O group. The nodes in the I/O group provide access to the volumes in the I/O group. A specific volume is always presented to a host server by a single I/O Group of the system. The I/O Group associated with a volume can be changed.

IBM Spectrum Virtualize cluster

When the first two nodes are installed and configured, they form an IBM Spectrum Virtualize cluster that can contain up to four I/O groups. Additional I/O groups can be added to an existing cluster

Internal storage and external storage

IBM Spectrum Virtualize can manage a combination of internal and external storage:

- ▶ Internal storage: These drives are used to create a Redundant Array of Independent Disks (RAID) that is presented as managed disks (MDisks) to the system. Distributed RAID does away with a dedicated spare, instead using spare capacity spread across all the drives, enhancing performance and reliability.
- ▶ External storage: IBM Spectrum Virtualize can detect logical units (LUNs) on an external storage system that is attached through Fibre Channel or iSCSI connections. These LUNs are detected as managed disks (MDisks) in the system, and must be protected from drive failures using RAID technology on the external storage system. External storage is also referred to as back-end storage.

Front-end and back-end systems

When virtualizing external storage, IBM Spectrum Virtualize acts as an initiator, or host to that storage. The virtualized storage system is referred to as the back-end system, and IBM Spectrum Virtualize is the front-end system, acting as the target, or disk system to the attached hosts.

Managed disks (MDisks)

For SAN Volume Controller-based implementations, a managed disk (MDisk) is a disk presented by a storage system and managed by IBM Spectrum Virtualize. This MDisk must be on a protected storage system, as IBM Spectrum Virtualize cannot provide data protection - like RAID - on an external system. For FlashSystem hardware implementations, an MDisk is a RAID array of some (or all) of the internal drives. In both cases, an MDisk provides usable blocks (or extents) of physical storage to the IBM Spectrum Virtualize cluster. If an MDisk is external, it must be configured so that it is not visible to host systems on the SAN. An MDisk in an IBM Spectrum Virtualize cluster can be either managed or unmanaged. A managed MDisk is an MDisk assigned to a storage pool.

Storage pool

A storage pool is a collection of MDisks that jointly contain all the data for a specified set of volumes. Each storage pool is divided into a number of *extents*. When creating a storage pool, you must choose an extent size. After it is set, the extent size stays constant for the life of that storage pool. Each storage pool can have a different extent size.

Data Reduction Pools

A Data Reduction Pools (DRP) is a type of storage pool, implementing techniques such as thin-provisioning, compression, and deduplication to reduce the amount of physical capacity required to store data.

Volumes

A volume is a logical entity that represents extents contained in one or more MDisks from a storage pool. Volumes are allocated in whole numbers of extents. Each volume is only associated with a single I/O group. The volume is then presented to the host as a LUN for use.

There are three types of volumes:

- ▶ **Striped:** A volume created in striped mode has extents allocated from each MDisk in the storage pool in a round-robin fashion.
- ▶ **Sequential:** The volumes are allocated sequentially on one MDisk to create the volume.
- ▶ **Image:** Image-mode volumes are special volumes that have a direct relationship with one MDisk. If you have a MDisk that contains data that you want to merge into the clustered system, you can create an image-mode volume.

Host

Volumes can be mapped to a *host* to allow access for a specific server to a set of volumes. A host definition in IBM Spectrum Virtualize is a collection of host bus adapter (HBA) worldwide port names (WWPNs) or iSCSI-qualified names (IQNs) that are defined on the specific server.

Host cluster

A *host cluster* is a combination of two or more hosts that are connected to an IBM Spectrum Virtualize cluster through a Fibre Channel, FCoE, NVME-oF, or an iSCSI connection. A host cluster object can see the same set of volumes. Therefore, volumes can be mapped to a *hostcluster* to allow all hosts in that cluster to have a common mapping

Easy Tier

IBM Spectrum Virtualize includes Storage IBM Easy Tier®, a function that optimizes performance at the extent level when there are multiple types of MDisks in a single storage pool. The system will promote extents that have a lot of I/O activity (a hot extent) to the MDisk with the best potential performance. It will also demote an extent that no longer has a lot of activity to an MDisk with lower potential performance. From the host perspective, this tiering is seamless, as IBM Spectrum Virtualize always knows where each extent of a volume actually resides.

Quorum disks

Quorum disks have two main functions. The first is to store information about the cluster, second is to avoid a situation where failures leave the cluster in an incorrect configuration. A quorum disk is less than 1 GiB in size. Each IBM Spectrum Virtualize cluster should have 3 quorum drives for redundancy (only one quorum disk is active at any one time).

IP quorum

The IP quorum application is a Java application that runs on a separate server or host.

With respect to the quorum disk, IP quorum is a feature that can enable the use of a low-cost IP network-attached host as a quorum disk for simplified implementation and operation.

Before IP quorum support was introduced, the third site had to be connected using Fibre Channel, and maintaining this third site and storage controller over FC makes the system costly for site recovery implementation of multi-IOgroup IBM Spectrum Virtualize systems. Note that there should still be three quorum disks in the system.

Access modes

The access mode determines how the IBM Spectrum Virtualize system uses the MDisk. The following are the types of access modes:

- ▶ **Unmanaged:** The MDisk is not used by the system. It is not assigned to a storage pool.

- ▶ **Managed:** The MDisk is assigned to a storage pool and provides extents that volumes can use.
- ▶ **Image mode:** The MDisk is assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. This mode enables you to insert the IBM Spectrum Virtualize system into the data path of an existing storage volume with minimal downtime.
- ▶ **Array:** The MDisk represents a set of drives in a RAID from internal storage.

Important: If you add an MDisk that contains existing data to a storage pool while the MDisk is in unmanaged or managed mode, the data is lost. The image mode is the only mode that preserves this data.

Mirrored volumes

Volume mirroring allows a volume to have two physical copies. The mirrored volume feature provides a simple RAID-1 function. Each volume copy can belong to a different storage pool, and can belong to separate storage systems.

The primary copy indicates the preferred volume for read requests. When a host writes to a mirrored volume, the system writes the data to both copies. You can create a volume with one or two copies, and you can convert a non-mirrored volume into a mirrored volume by adding a copy. Alternately, you can convert a mirrored volume into a non-mirrored volume by deleting one copy or by splitting one copy to create a new, non-mirrored volume.

System state

The state of the clustered system holds all of the configuration and internal data. The system state information is held in nonvolatile memory. If the power fails, the BBUs maintain internal power long enough for system state information to be stored. This information is stored on quorum disks.

Back up and restore process

Planning your backup/restore procedures is an important part of a complete backup and disaster recovery solution.

Backup is the process of extracting configuration settings from an IBM Spectrum Virtualize system and writing them to disk. The restore process uses backup configuration data files to restore system configuration.

If power fails on a system or a node in a system is replaced, the system configuration settings are automatically restored. This restoration also occurs when the repaired node is added to the system. To restore the system configuration in a disaster, plan to back up the system configuration settings to tertiary storage. Use the configuration backup functions to back up the system configuration.

Important: For complete disaster recovery, regularly back up the business data that is stored on volumes at the application server or host level.

Extents

An extent is a fixed-size unit of data that is used to manage the mapping of data between MDisks and volumes. The extent size choices are 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, and 8192 MB. The choice of extent size affects the total storage that can be managed by the SAN Volume Controller.

For the 8192 MB extent size, SAN Volume Controller supports the following capacities:

- ▶ A maximum volume capacity of 256 TB
- ▶ A maximum IBM Spectrum Virtualize system capacity of 32 PB

Figure 3-2 shows the relationship between extents and a volume.

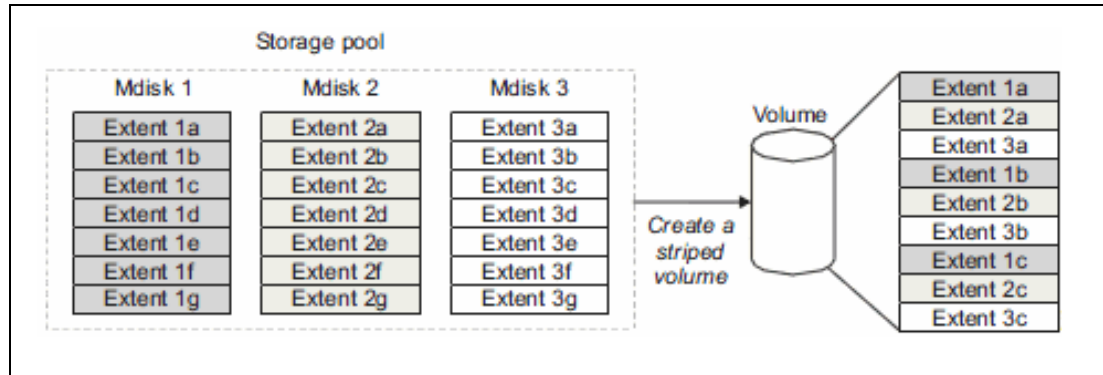


Figure 3-2 Extents used to create a volume

Figure 3-3 shows the relationship between physical and virtual disks.

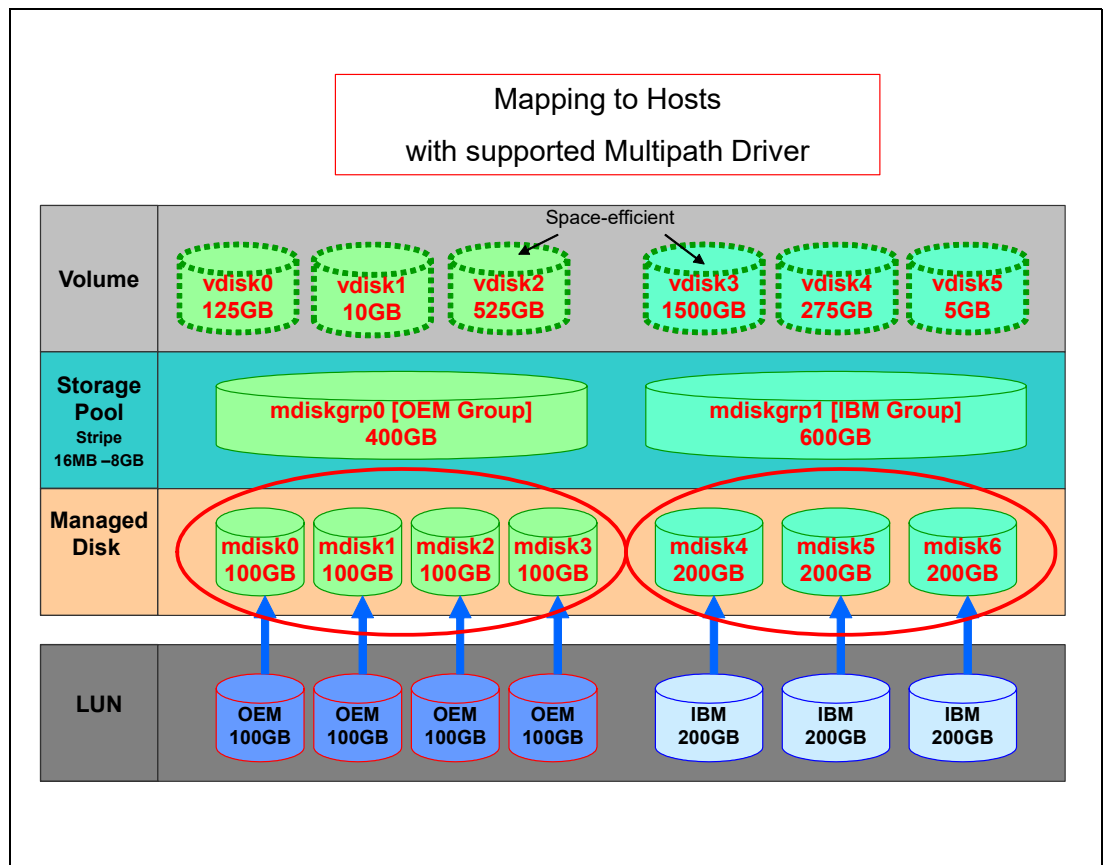


Figure 3-3 Relationships between physical and virtual disks

Consistency groups

Consistency groups preserve data consistency across multiple volumes, especially for applications that have related data that spans multiple volumes.

IBM Spectrum Virtualize data mobility

IBM Spectrum Virtualize software provides unparalleled data mobility. It manages the extents that make up a volume, allowing for data movement without interrupting host access.

This data movement capability can be used for the following tasks:

- ▶ Redistribution of volumes and their workload within an IBM Spectrum Virtualize system cluster across back-end storage
- ▶ Moving workload onto newly installed storage systems
- ▶ Moving workload off storage so that old or failing storage subsystems can be decommissioned
- ▶ Moving workload to rebalance a changed workload.

Multipath Device Driver

For resilience, availability, redundancy, and improved throughput, the use of multipath device drivers has long been considered a best practice.

Note that the IBM Multipath Subsystem Device Driver for Windows (SDDDSM), and AIX (SDDPCM) are not compatible with any IBM Spectrum Virtualize code after version 8.3.0 and those drivers will no longer be supported as of June 2020.

See the IBM Support page at:

https://www.ibm.com/support/pages/node/1106937?myns=s035&mync=E&cm_sp=s035-_-NULL-_-E

Instructions for migrating to supported multipath drivers are also included in the previous link.

3.3.1 IBM Spectrum Virtualize copy services

IBM Spectrum Virtualize supports the following copy services:

- ▶ Synchronous remote copy (Metro Mirror)
- ▶ Asynchronous remote copy (Global Mirror)
- ▶ IBM FlashCopy® (Point-in-Time copy)
- ▶ Transparent Cloud Tiering

The copy services layer sits above and operates independently of the function or IBM characteristics of the underlying disk systems used to provide storage resources to an IBM Spectrum Virtualize instance.

For details, see the IBM Redbooks publication *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize*, SG24-7933, Advanced Copy Services.

The Global Mirror function can provide additional capabilities for data migration between compatible IBM Spectrum Virtualize systems.

Global Mirror (GM) works by defining a relationship between two volumes of equal size and maintains the data consistency in an asynchronous manner.

Global Mirror has two options, Global Mirror and Global Mirror with Change Volumes. Global Mirror can provide a very low RPO, but might require high bandwidth to maintain that low RPO. Global Mirror with Change Volumes provides a low RPO, but with a lower bandwidth requirement.

A migration technique based on Global Mirror with Change Volumes is presented and illustrated in Chapter 7, “Data Migration with IBM Spectrum Virtualize Remote Copy Services” on page 103.



IBM Spectrum Virtualize based migration

This chapter addresses various methods for migrating data to IBM Spectrum Virtualize systems. To reduce complexity while maintaining data integrity, IBM Spectrum Virtualize methods are usually the best solutions for data migration.

This chapter provides an overview of the following methods:

- ▶ Data migration using IBM Spectrum Virtualize Volume Migration wizard
- ▶ Data migration using IBM Spectrum Virtualize Metro or Global Mirror
- ▶ Data migration using mirrored volumes
- ▶ Intra-Cluster Migration with nondisruptive volume move

More details and use cases illustrating each method can be found in subsequent chapters of this book.

4.1 Interoperability and compatibility

Interoperability is an important consideration when a new storage system is set up in an environment that already contains existing storage infrastructure. Before attaching any external storage systems to the IBM Spectrum Virtualize system, see the IBM System Storage Interoperation Center (SSIC):

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

Select **IBM System Storage SAN Volume Controller**, **IBM System Storage Enterprise Flash**, or **IBM System Storage Midrange Disk** in Storage Family, depending on the IBM Spectrum Virtualize system which will be used, then select the Storage Model and other characteristics of your environment to check for compatibility.

4.2 Data migration using IBM Spectrum Virtualize Volume Migration wizard

To migrate data from an existing storage system to a IBM Spectrum Virtualize system without host-based methods, using Copy Services, or clustering with an older IBM Spectrum Virtualize system, it is necessary to use the built-in external virtualization capability. This capability places externally connected Logical Unit Numbers (LUNs) under the control of the IBM Spectrum Virtualize system. After LUNs are virtualized, hosts continue to access them but do so through the IBM Spectrum Virtualize system, which acts as a proxy.

Licensing considerations: A license is required for each enclosure of any external storage systems that are being virtualized. Data can be migrated from existing storage systems to a IBM Spectrum Virtualize system by using the external virtualization function. After migration, if you decide to keep the migrated storage system under IBM Spectrum Virtualize control, a license for each enclosure in each external storage system is required.

Set the license temporarily during the migration process to prevent messages that indicate that you are in violation of the license agreement. When the migration is complete reset the license to its original limit or purchase a new license.

The following steps give an overview of the storage migration process:

- ▶ Typically, storage systems divide storage into many SCSI LUNs that are presented to hosts.
- ▶ I/O to the LUNs must be stopped and changes must be made to the mapping of the external storage system LUNs and to the fabric configuration so that the original LUNs are presented directly to the IBM Spectrum Virtualize system and not to the hosts anymore. The IBM Spectrum Virtualize system discovers the external LUNs as *unmanaged* MDisks.
- ▶ The unmanaged MDisks are *imported* to the IBM Spectrum Virtualize system as *image-mode volumes* and placed into a temporary storage pool. This storage pool is now a logical container for the LUNs.

- ▶ Each MDisk has a one-to-one mapping with an image-mode volume. From a data perspective, the image-mode volumes represent the LUNs exactly as they were before the import operation. The image-mode volumes are on the same physical drives of the external storage system and the data remains unchanged. The IBM Spectrum Virtualize system is presenting active images of the LUNs and is acting as a proxy.
- ▶ The hosts must have any existing storage system multipath device driver removed and replaced by the multipath device driver compatible with IBM Spectrum Virtualize. The compatible device driver must also be configured for IBM Spectrum Virtualize system attachment. The IBM Spectrum Virtualize system hosts are defined with worldwide port names (WWPNs) or iSCSI qualified names (IQNs), and the volumes are mapped to the hosts. After the volumes are mapped, the hosts discover the IBM Spectrum Virtualize system volumes through a host rescan or reboot operation.
- ▶ Resume I/O to the LUNs.
- ▶ IBM Spectrum Virtualize volume mirroring operations are then initiated. The image-mode volumes are mirrored to generic volumes. Volume mirroring is an online migration task, which means a host can still access and use the volumes during the mirror synchronization process.
- ▶ After the mirror operations are complete, the image-mode volumes are removed. The external storage system LUNs are now migrated and the now redundant storage can be decommissioned or reused elsewhere.

See Chapter 5, “Using IBM Spectrum Virtualize Migration wizard” on page 49 for more details about IBM Spectrum Virtualize Storage Migration and the complete procedure.

4.3 Data migration using IBM Spectrum Virtualize Metro or Global Mirror

Data migrations using Metro or Global Mirror are more complex, but are a good solution if the system are geographically distant. The method minimizes I/O downtime and provides an easy backout. Use two IBM Spectrum Virtualize clusters when migrating using Metro or Global Mirror.

By using Metro or Global Mirror, the source volumes receive the host updates at all times during the data migration. If a problem occurs between the host and its data volumes, a normal recovery is possible. In the background, the IBM Spectrum Virtualize system copies the data over to the target volumes. The target LUN must be exactly the same size as the source LUN.

The following steps give a high-level overview of the storage migration process:

- ▶ Create a zone between the two IBM Spectrum Virtualize systems.
- ▶ Create a zone between new IBM Spectrum Virtualize system and host.
- ▶ Create a Copy Services partnership between the two IBM Spectrum Virtualize systems.
- ▶ Create volumes of the same size on the new (target) system.
- ▶ Create and start Metro or Global Mirror relationships between source and target volumes.
- ▶ Wait for the Metro or Global Mirror to reach *consistent synchronized* state.
- ▶ Quiesce all I/O.
- ▶ Shutdown application and unmount file system.
- ▶ If using Global Mirror check that all writes were copied to the new system.

- ▶ Delete Mirror relationships.
- ▶ Unmap volumes on old IBM Spectrum Virtualize system.
- ▶ Map volumes on new IBM Spectrum Virtualize system.
- ▶ Remove old devices on host and rescan for the new ones.
- ▶ Mount the file system and start application.
- ▶ Remove partnership and zoning between clusters, if not needed anymore.

See Chapter 7, “Data Migration with IBM Spectrum Virtualize Remote Copy Services” on page 103 for more details about IBM Spectrum Virtualize Copy Services and the complete procedure.

4.4 Data migration using mirrored volumes

If you already have a IBM Spectrum Virtualize system and a license for External Virtualization, the easiest and an online method for storage-based data migration is volume mirroring.

IBM Spectrum Virtualize offers the capability to mirror volumes, which means that a single volume that is presented to a host can have two physical copies. Each volume copy can belong to a different pool, and each copy has the same virtual capacity. When a server writes to a mirrored volume, the storage system writes the data to both copies. When a server reads a mirrored volume, and the volume copies are synchronized, the system reads the data from the primary copy. In the management GUI, the primary volume copy is marked by an asterisk (*).

If one of the mirrored volume copies is temporarily unavailable (for example, because the storage system that provides the pool is unavailable), the volume remains accessible to servers. The system remembers which areas of the volume were modified after loss of access to a volume copy and re synchronizes only these areas when both copies are again available.

The use of mirrored volumes has the following advantages:

- ▶ Improves availability of volumes by protecting them from a single storage system failure.
- ▶ Enables concurrent maintenance of a storage system that does not natively support concurrent maintenance.
- ▶ Provides an alternative method of data migration.
- ▶ Enables conversion between fully allocated volumes and thin-provisioned volumes.
- ▶ Enables space reclamation of thin-provisioned volumes.

The two copies of a mirrored volume typically are allocated from separate storage pools that are backed by different physical hardware. Volume copies are identified in the IBM Spectrum Virtualize GUI by a copy ID, which can have value 0 or 1. Copies of the volume can be split.

Each volume copy is not a separate object and can be manipulated only in the context of the volume. A mirrored volume behaves in the same way as any other volume. In other words, all of its copies are expanded or shrunk when the volume is resized, the volume can participate in FlashCopy and remote copy relationships, is serviced by an I/O Group, and has a preferred node.

Figure 4-1 on page 47 provides an overview of volume mirroring.

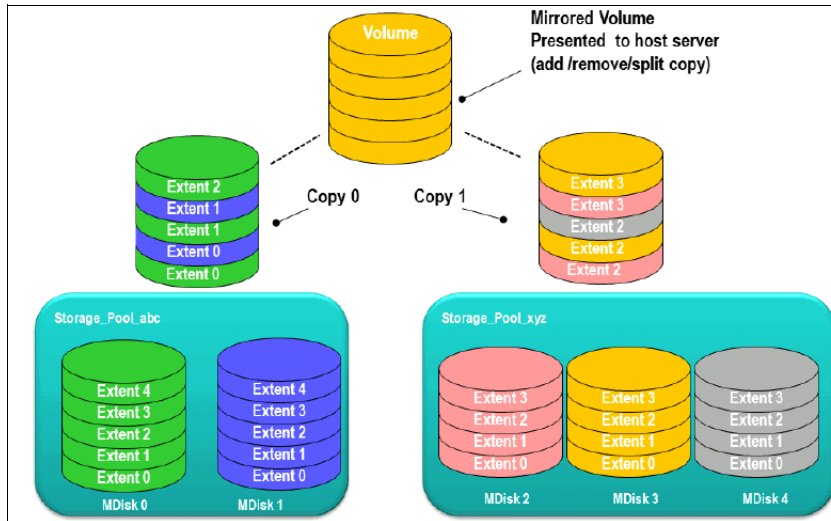


Figure 4-1 IBM Spectrum Virtualize mirrored volumes

The following steps give a high-level overview of the storage migration process:

- ▶ Create a zone between the two IBM Spectrum Virtualize systems.
- ▶ Create volumes on the virtualized IBM Spectrum Virtualize system (the system you want to migrate) and map it to the virtualization IBM Spectrum Virtualize system (the new system).
- ▶ Discover Storage on the virtualization IBM Spectrum Virtualize system.
- ▶ Create a pool on the newly discovered external MDisks.
- ▶ Add a volume copy with option **autodelete** from the new pool to the volumes which need to be migrated, **autodelete** will automatically delete the primary copy after secondary copy is synchronized.
- ▶ Remove original pool and old virtualized storage system from the virtualization system, if not needed anymore.

See 6.2, “Migrating volumes between pools using Volume Mirroring” on page 71 for more details about IBM Spectrum Virtualize Volume Mirroring and the procedure.

4.5 Intra-Cluster Migration with nondisruptive volume move

If you already have a IBM Spectrum Virtualize system which is capable of clustering with the new IBM Spectrum Virtualize system, an easy and online method for storage-based data migration is Intra-Cluster migration with nondisruptive volume move (NDVM) between I/O groups.

Some of IBM Spectrum Virtualize systems are capable of clustering with certain other IBM Spectrum Virtualize systems, for example:

- ▶ SAN Volume Controller SV2/SA2 can cluster with any SV1 and DH8 SAN Volume Controller clusters running the minimum level of software (8.3.1) up to 4 IO Groups.
- ▶ Storwize V7000 (model 524, 624 and 724) can cluster with FlashSystem 7200, 9100, and 9200 running the appropriate minimum level of software up to 4 IO Groups.
- ▶ FlashSystem 5100 can cluster with Storwize V5100 up to 2 IO groups.

- ▶ FlashSystem 5030 can cluster with Storwize V5030 up to 2 IO Groups.

The name of the cluster will be renamed to the most recent member of the cluster. For example, if you cluster a Storwize V7000 with a FlashSystem 9100 the cluster will be renamed to FlashSystem 9100. The cluster will be renamed according to the following priority:

- ▶ Storwize → FlashSystem
- ▶ Storwize V7000 → FlashSystem 7200 → FlashSystem 9100 → FlashSystem 9200

NDVM can move volumes between I/O Groups to redistribute the load between the I/O Groups or to migrate data to another I/O Group. Modifying the I/O Group that services the volume can be done concurrently with I/O operations if the host supports nondisruptive volume moves.

For hosts and operating systems supporting NDVM see the following web pages:

- ▶ <https://www.ibm.com/support/pages/node/957103>
- ▶ <https://www.ibm.com/support/pages/83x-recommended-software-levels-ibm-storwize-v7000v5000>

NDVM also requires a rescan at the host level to ensure that the multipathing driver is notified that the allocation of the preferred node changed, and the ports (by which the volume is accessed) changed.

The following steps give a high-level overview of the storage migration process, using NDVM:

- ▶ Update the existing cluster to firmware that is supported by the new system.
- ▶ Add the new nodes to the existing cluster (maximum of 8 nodes).
- ▶ Add access to the volumes for the new I/O Group.
- ▶ Move the volumes to the new I/O Group.
- ▶ Rescan on host to detect the new paths.
- ▶ Remove the access to the volumes from the old I/O group.
- ▶ Rescan on host to remove the old paths.
- ▶ Remove old I/O Group from the cluster, if not needed anymore.

To move a volume between I/O groups using the CLI, complete the following steps:

1. Issue the following command: **addvdiskaccess -iogrp iogrp id/name volume id/name**
2. Issue the following command: **movevdisk -iogrp destination iogrp -node new preferred node volume id/name**

The system disables moving a volume if the selected volume is currently performing quick initialization. After the quick initialization completes, you can move the volume to another I/O group.

3. Issue the appropriate commands on the hosts that are mapped to the volume to detect the new paths to the volume in the destination I/O group.
4. After confirming that the new paths are online, remove access from the old I/O group: **rmvdiskaccess -iogrp iogrp id/name volume id/name**
5. Issue the appropriate commands on the hosts that are mapped to the volume to remove the paths to the old I/O group.



Using IBM Spectrum Virtualize Migration wizard

This chapter describes the steps which are necessary to complete a data migration from an existing storage system to a IBM Spectrum Virtualize based storage system using the IBM Spectrum Virtualize Migration wizard.

The IBM Spectrum Virtualize Migration wizard conducts a data migration concurrently when the storage system being migrated from and the new storage system being migrated to have been properly configured in the IBM Spectrum Virtualize environment. The configuration of the storage systems may require a maintenance window to complete the necessary changes in preparation for the migration.

The IBM Spectrum Virtualize Migration wizard simplifies the task of migrating data by presenting easy to follow steps which guide the storage administrator through the process of migrating their data. The IBM Spectrum Virtualize wizard allows migration between storage systems using either Fibre Channel (FC) or Internet Small Computer Systems Interface (iSCSI) connections.

Note: This chapter does not cover migration outside of the IBM Spectrum Virtualize Migration wizard. To migrate data outside of the wizard, you must use either host based or other IBM Spectrum Virtualize based data migration methods explained in subsequent chapters of this book.

5.1 IBM Spectrum Virtualize Migration wizard

The IBM Spectrum Virtualize Migration wizard simplifies the process of successfully completing a data migration task by guiding the storage administrator through the process. Necessary steps are presented to the user as either commands or images so that the exact process can be easily followed and understood.

Important: Although, risk of failure when using the IBM Spectrum Virtualize Migration wizard is low, it is prudent to help avoid potential problems by having a current backup of all the data that is stored on the hosts and the existing storage system before migrating.

5.1.1 Prerequisites

Using the IBM Spectrum Virtualize Migration wizard requires that you either have a IBM Spectrum Virtualize environment already in place that is virtualizing MDisks from an external storage system or that you are migrating data to a new IBM Spectrum Virtualize system, behind which you can present data from your existing external storage systems for migration.

An example for completing a migration from an external storage system like the IBM FlashSystem A9000 to a IBM Spectrum Virtualize system would include attaching the FlashSystem A9000 system as a back end to the new IBM Spectrum Virtualize system. The wizard will guide you through the process of validating the connectivity between the FlashSystem A9000 and the IBM Spectrum Virtualize system, connecting and configuring the hosts and the IBM Spectrum Virtualize system and updating the host MDisk (Volume) mappings to point to the IBM Spectrum Virtualize system MDisks. After successful migration to the IBM Spectrum Virtualize system, the hosts can be unmapped from the IBM FlashSystem A9000 and use the IBM Spectrum Virtualize system migrated MDisks for their storage.

Before starting a data migration using the IBM Spectrum Virtualize Migration wizard, the external storage system and the system being migrated to must be visible to IBM Spectrum Virtualize. Both systems must be connected either using iSCSI or Fibre Channel connections, have proper multipathing already configured, and be following best practices for the existing storage system and IBM Spectrum Virtualize connectivity procedures.

Note: Validate that all systems involved in the data migration are using currently supported drivers and are running compatible system versions as indicated by the IBM System Storage Interoperation Center (SSIC) tool at:

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

Prerequisites for Fibre Channel connections

The following are prerequisites for Fibre Channel connections:

- Verify that IBM Spectrum Virtualize system is properly cabled into the same storage area network (SAN) as the external storage system that you are migrating. If using traditional Fibre Channel, connect the Fibre Channel cables to the Fibre Channel ports in *all* nodes of the IBM Spectrum Virtualize system and to the Fibre Channel network switch.

Note: A leading practice is to test any setting changes on a non-production environment before proceeding with data migration on production systems

Prerequisites for iSCSI connections

The following are prerequisites for iSCSI connections:

- ▶ Connect the IBM Spectrum Virtualize system to the external storage system with redundant connections. Migrating iSCSI external storage requires that the IBM Spectrum Virtualize system and the storage system be connected through an Ethernet switch. Symmetric ports on *all* nodes of the system must be connected to the same switch and must be configured on the same subnet.

In addition, modify the Ethernet port attributes to enable the external storage on the Ethernet port so that the IBM Spectrum Virtualize system allows proper connectivity to the external storage system. To modify the Ethernet port attribute for the external storage system, select **Network** → **Ethernet Ports** and right-click a configured port. Select **Modify Storage Ports** to enable that port for external storage connections.

Cable the Ethernet ports on the storage system to the Ethernet switch in the same way as the IBM Spectrum Virtualize system and ensure that the ports are configured to use the same subnet. Optionally, you can use a virtual local area network (VLAN) to define network traffic for the system ports and storage system ports, if all associated ports are contained within that VLAN.

For full redundancy, configure and use separate Ethernet switches in the data path to prevent network resource issues from impacting the connected systems. If the source system nodes and the external storage system both have more than two Ethernet ports, additional redundant iSCSI connections can be established for increased throughput.

5.1.2 Using the wizard

Perform the following steps to complete the migration by using the IBM Spectrum Virtualize Migration wizard.

1. Click **Pools** → **System Migration**, as shown in Figure 5-1. The System Migration pane provides access to the storage migration wizard and displays information about the migration progress.

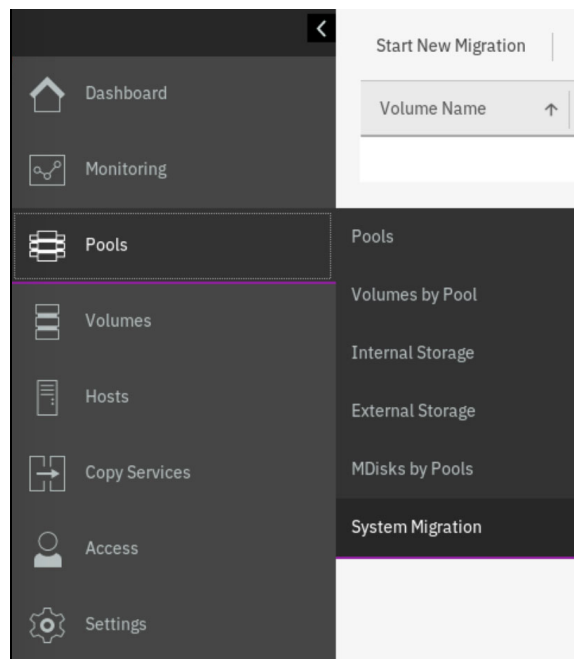


Figure 5-1 Going to System Migration Using IBM Spectrum Virtualize

2. Click the **Start New Migration** button to begin the storage migration wizard, as shown in Figure 5-2.




Figure 5-2 The Start New Migration button located above the data migration pane.

Note: Starting a new migration creates a volume to be migrated in the list of data migrations displayed in the pane. After a volume is migrated, it remains in the list until you *finalize* the migration.

3. If both Fibre Channel and iSCSI external systems are detected, a dialog is shown asking you which connection protocol should be used. Select the wanted method of attachment between the IBM Spectrum Virtualize system and the external storage system to be used for the migration and click **Next**. If only one type of attachment is detected, this dialog is not displayed.

Note: If the external storage system is not detected, the warning message shown Figure 5-3 is displayed when you attempt to start the migration wizard. Click **Close** and correct the problem before reattempting use of the migration wizard.

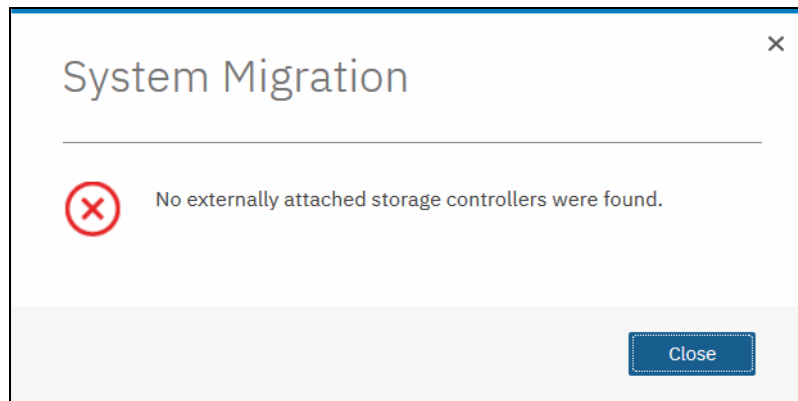


Figure 5-3 Error message if no external storage is not detected

4. When first started, the IBM Spectrum Virtualize Migration wizard prompts the user to verify that all necessary restrictions and prerequisites are accounted for. See Figure 5-4 on page 53.
 - Restrictions:
 - You are not using the IBM Spectrum Virtualize Migration wizard to migrate clustered hosts, including clusters of VMware hosts or Virtual I/O Servers (VIOS) due to additional configuration necessary for these types of hosts to be successfully migrated.
 - You are not using the IBM Spectrum Virtualize Migration wizard to migrate SAN boot images.

If either of these situations apply, the migration must be performed using an alternative migration method among the methods explained in other chapters of this book.

– Prerequisites:

IBM Spectrum Virtualize nodes and the external storage systems are connected to the same SAN fabric.

Be sure to read “Prerequisites” on page 50 for other details about the Storage Migration prerequisites.

5. If all restrictions are satisfied and prerequisites are met, select all of the boxes and click **Next**, as shown in Figure 5-4 on page 53.

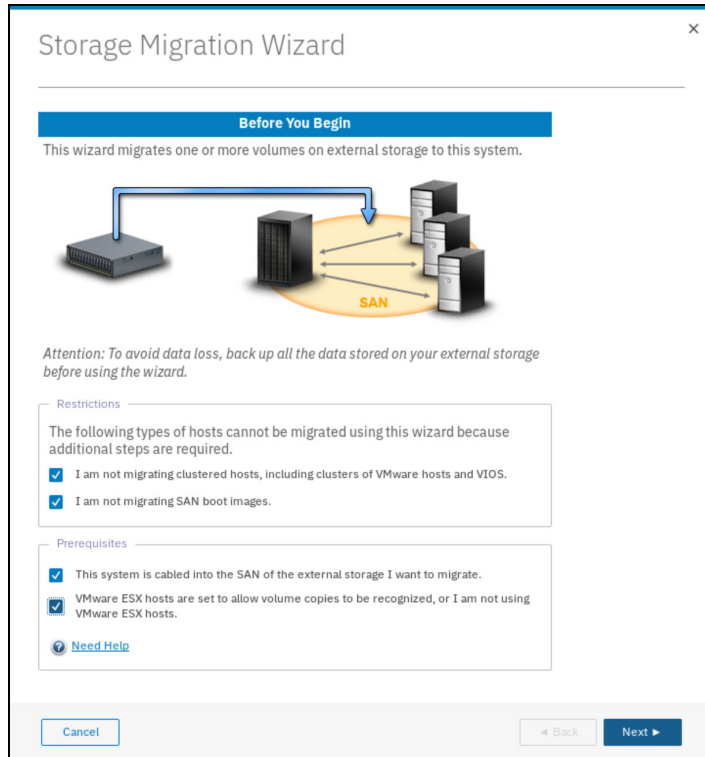


Figure 5-4 Restrictions and prerequisites confirmation

6. Prepare the migration by following the on-screen instructions that are shown in Figure 5-5.

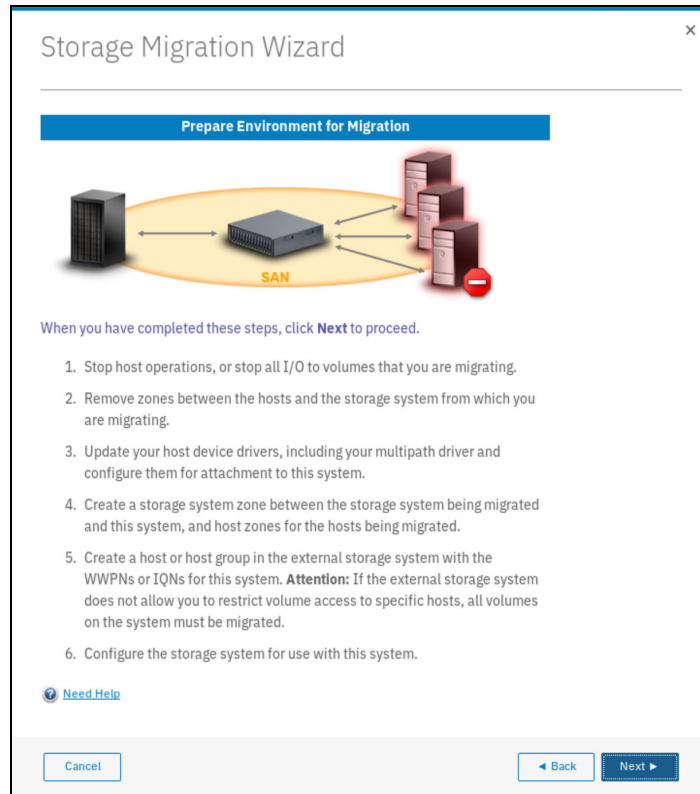


Figure 5-5 Preparing the environment for storage migration

The preparation phase includes the following steps:

- a. Before beginning the migration, ensure that all host operations are stopped and volumes are unmounted on the hosts to prevent applications from generating I/Os to the migrated system.
- b. Hosts may not support concurrent multipath drivers at the same time. Drivers that are not compatible with IBM Spectrum Virtualize might need to be removed from the hosts and recommended drivers installed.
- c. If migrating external storage systems that are connected to IBM Spectrum Virtualize using Fibre Channel, ensure that appropriate zoning changes have been conducted to simplify the migration. Use the following guidelines to ensure that Fibre Channel zones are configured correctly for migration:
 - Storage system zones
For each storage system being migrated, create a single zone that contains the IBM Spectrum Virtualize ports from all nodes and at least two ports from each controller on the external storage system unless otherwise stated by the zoning recommendations for that storage system.
 - Host zones
In the host zone, the host systems can identify and address the IBM Spectrum Virtualize nodes. You can have more than one host zone and more than one storage system zone. Create one host zone for each host which contains its Fibre Channel ports with one port zoned as single initiator.

- The IBM Spectrum Virtualize system will be seen as a host cluster from the external storage system being migrated. IBM Spectrum Virtualize must be configured as a host or host group using its WWPNs or IQNs, on the system to be migrated. Some established systems do not support LUN-to-host mapping, and therefore present all logical units (LUNs) to IBM Spectrum Virtualize. In that case, all the LUNs from that system should be migrated.
- Also, make sure that for your new or existing IBM Spectrum Virtualize system that the SAN configuration rules and zoning details are met. See this section of IBM Knowledge Center for IBM Spectrum Virtualize:

https://www.ibm.com/support/knowledgecenter/STSLR9_8.3.1/com.ibm.fs9200_831.doc/svc_configrulessummary_02171530.html

7. After the preparation steps have been completed, the IBM Spectrum Virtualize system is seen as a host by the external storage system. LUNs can be mapped to IBM Spectrum Virtualize. Map the external storage system by following the on-screen instructions that are shown in Figure 5-6.

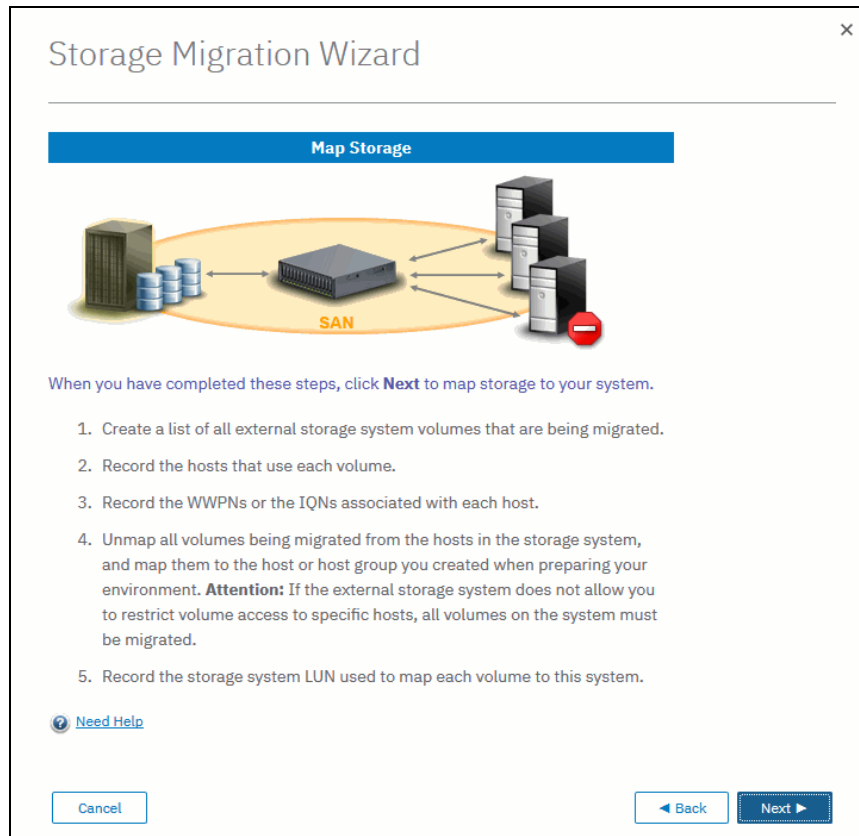


Figure 5-6 Steps to map the LUNs to be migrated to IBM Spectrum Virtualize

Before starting the migration of the storage system, record the hosts and their WWPNs or IQNs for each volume that is being migrated, in addition to the SCSI LUN ID when the volume or LUN is mapped to IBM Spectrum Virtualize.

Table 5-1 shows an example of the type of table that can be used to capture information that relates to the external storage system LUNs.

Table 5-1 Example table for recording external LUN information

Volume Name or ID	Hosts accessing this LUN	Host WWPNs or IQNs	SCSI / LUN ID when mapped
1 IBM Db2® logs	DB2server	21000024FF2...	0
2 Db2 data	DB2Server	21000024FF2...	1
3 file system	FileServer1	21000024FF2...	2

Note: Make sure to record the SCSI ID of the LUNs to which the host is originally mapped. Some operating systems do not support changing the SCSI IDs during data migration.

8. Click **Next** and wait for the system to discover external devices.
9. The next step displays all MDisks that were discovered after the last step. These MDisks are the original volumes or LUNs on the external storage system. If the MDisks or volumes being migrated are not in the list, verify zoning or IP configuration and as applicable the volume LUN mappings. Repeat the previous step to trigger the discovery procedure again.
10. Select the MDisks to be migrated, as shown in Figure 5-7. In this example, one MDisk has been found and will be migrated: `mdisk1`. Detailed information about an MDisk is visible by double-clicking it.

To select multiple elements or MDisks from within the pane, use the standard **Shift+left-click** or **Ctrl+left-click** actions.

Note: MDisks which were successfully discovered can be exported to a CSV file, for further use, by clicking **Export to CSV**.

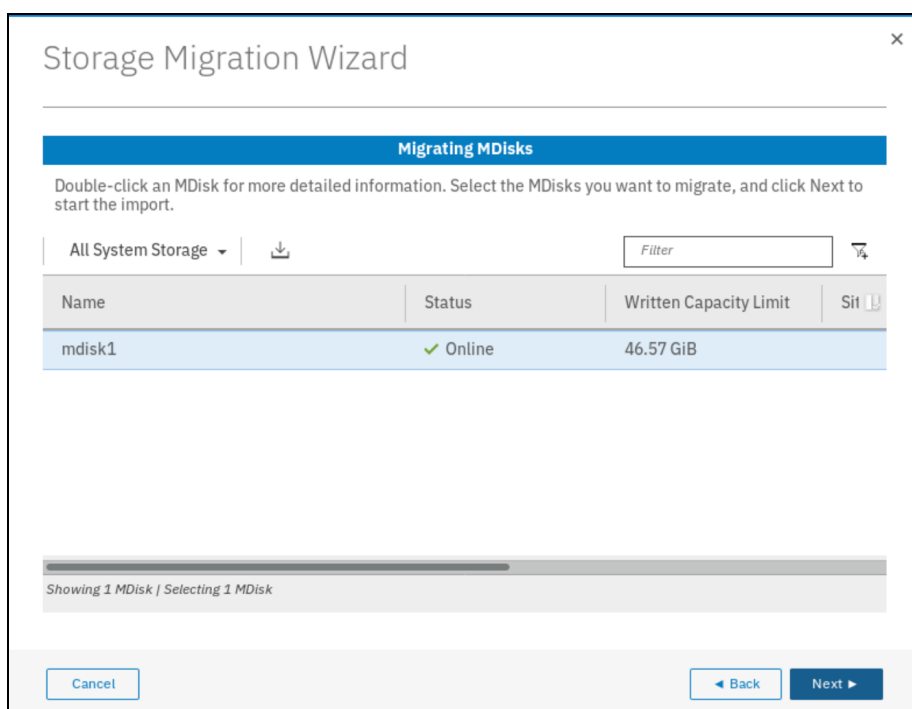


Figure 5-7 Discovering mapped LUNs from external storage

Note: Select only the MDisks that are applicable to the current migration plan. After 21 on page 61 of the current migration completes, another migration can be started to migrate any remaining volumes or LUNs.

11. Click **Next** and wait for the MDisk to be imported. During this task, the system creates a new storage pool called MigrationPool_XXXX (where XXXX stands for the extent site in MiB, 1024 MiB per default) and adds the imported MDisk to the storage pool as an image-mode volume. One MDisk corresponds to one logical volume on the storage system being migrated.

12. The next window lists all of the hosts that are configured on the system and enables you to configure new hosts. This step is optional and can be bypassed by clicking **Next**.

In this example, the host iSCSI_Host was already configured, as shown in Figure 5-8. If no host is selected, you can create a host after the migration completes and map the imported volumes to it.

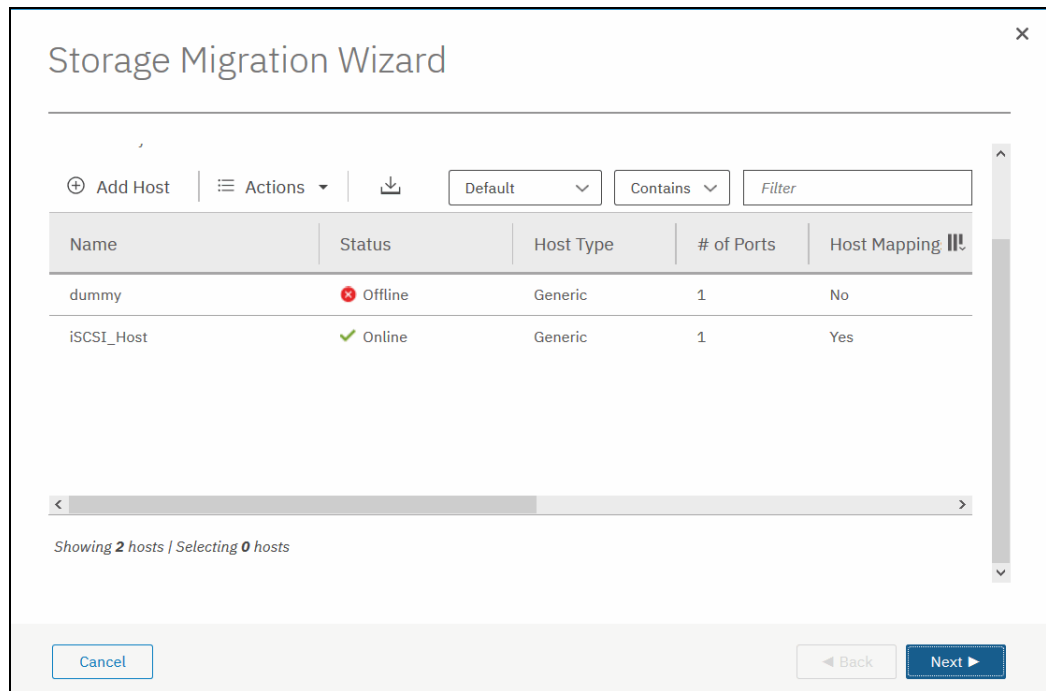


Figure 5-8 Listing of configured hosts to which to map the imported volume

13. If the host which needs access to the migrated data is not configured, select **Add Host** to begin the Add Host wizard. Enter the host connection type, name, and connection details in the appropriate fields. Optionally, click **Advanced** to modify the host type and I/O group assignment. Figure 5-9 shows the Add Host wizard with sample details completed.

The screenshot shows a window titled "Add Host" with a close button (X) in the top right corner. The window is divided into two sections: "Required Fields" and "Optional Fields".

Required Fields:

- Name: ISCSI HOST
- Host connections: iSCSI (SCSI)
- Site: None
- Host IQN: 24532

Optional Fields:

- CHAP authentication:
- CHAP secret: Enter 1 to 79 characters
- CHAP username: Enter 1 to 31 characters
- Host type: Generic
- I/O groups: All
- Host cluster: No Host Cluster Selected

At the bottom of the window, there are two buttons: "Cancel" and "Add".

Figure 5-9 If not already defined, you can create a host during the migration process

14. Click **Add**.

The host is created and is now listed in the Configure Hosts window, as shown in Figure 5-8 on page 57. Click **Next** to proceed.

15. The next window displays the new volumes and enables you to map them to hosts. The volumes are listed with names that were automatically assigned by the system during the IBM Spectrum Virtualize Migration wizard process. The names can be changed to reflect something more meaningful to the user by selecting the volume and clicking **Rename** in the **Actions** menu.

16. Map the volumes to hosts by selecting the volumes and clicking **Next**, as shown in Figure 5-10.

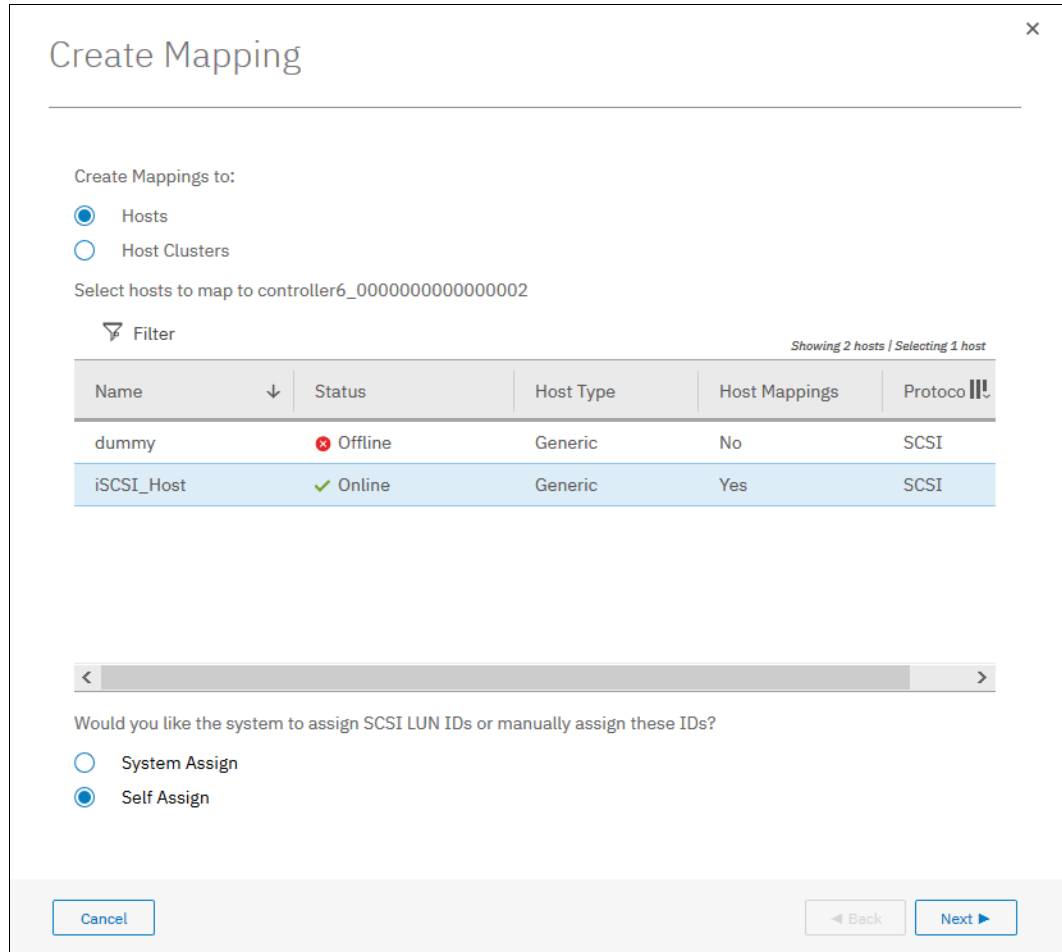


Figure 5-10 Select the host to map to the new Volume

17. SCSI IDs can be manually assigned to the LUNs being mapped. This technique is particularly useful when the host needs to have the same LUN ID for a LUN before and after it is migrated. To assign the SCSI ID manually, select the **Self Assign** option and follow the instructions as shown in Figure 5-11.

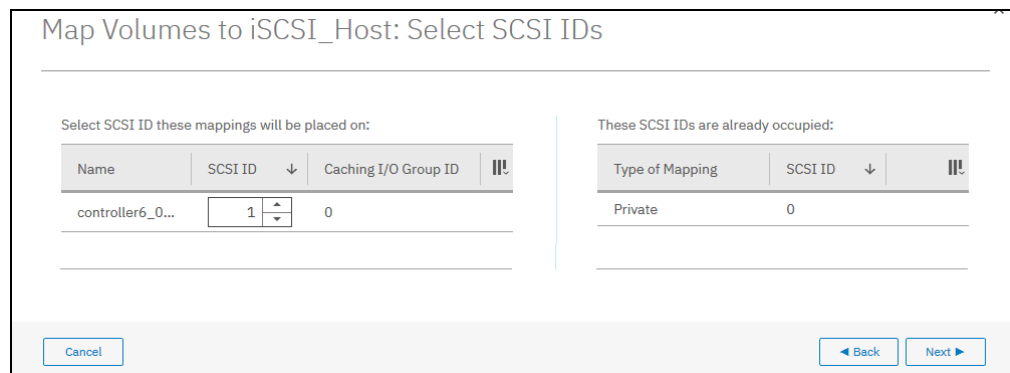


Figure 5-11 Manually assign a LUN SCSI ID to mapped Volume

18. When the LUN mapping is ready, click **Next**. A new window is displayed with a summary of the new and existing mappings, as shown in Figure 5-12.

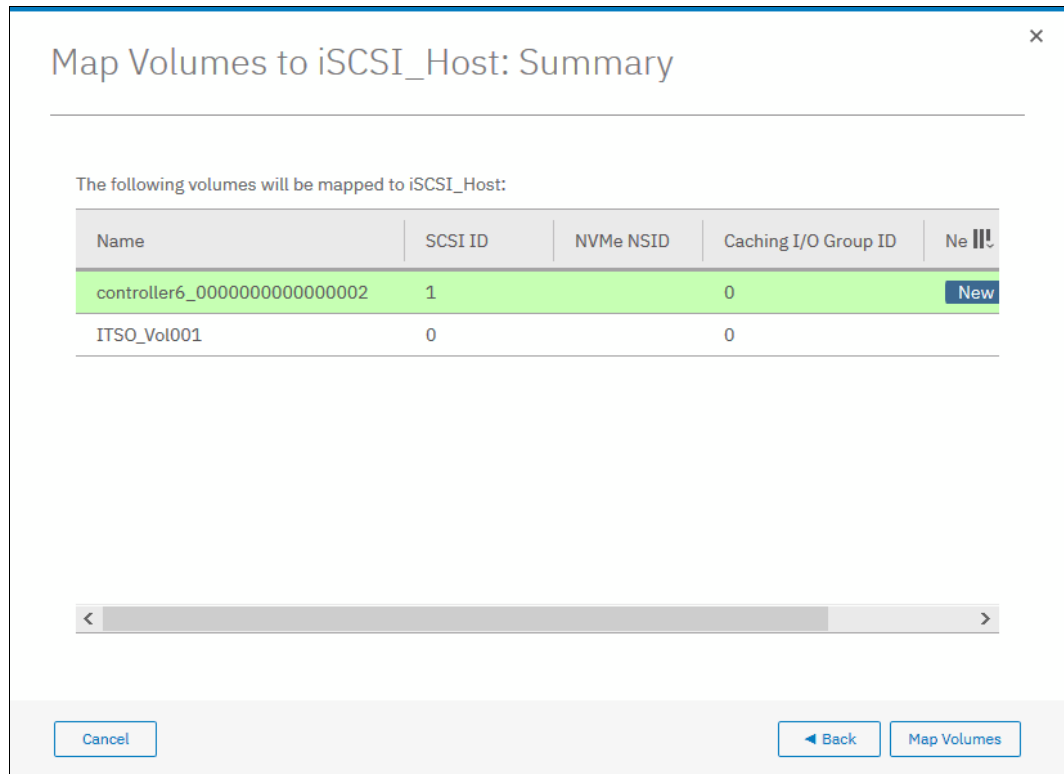


Figure 5-12 Volumes mapping summary before migration

19. Click **Map Volumes** and wait for the mappings to be created.
20. Select the storage pool that you want to migrate the imported volumes into. Ensure that the selected storage pool has enough space to accommodate the migrated volumes before continuing.

Note: You can decide not to migrate to a storage pool and to leave the imported MDisk as an image-mode volume. This technique is not recommended as no volume mirroring will be created.

With volume mirroring, the system creates two copies (Copy0 and Copy1) of a volume. Typically, Copy0 is located in the Migration Pool, and Copy1 is created in the target pool of the migration. When the host generates a write I/O on the volume, data is written at the same time on both copies.

Read I/Os are performed on the preferred copy. In the background, a mirror synchronization of the two copies is performed and runs until the two copies are synchronized. The speed of this background synchronization can be changed in the volume properties menu.

21. Click **Next**, as shown in Figure 5-13, to complete assignment of the image-mode VDisk or volume to the wanted pool.

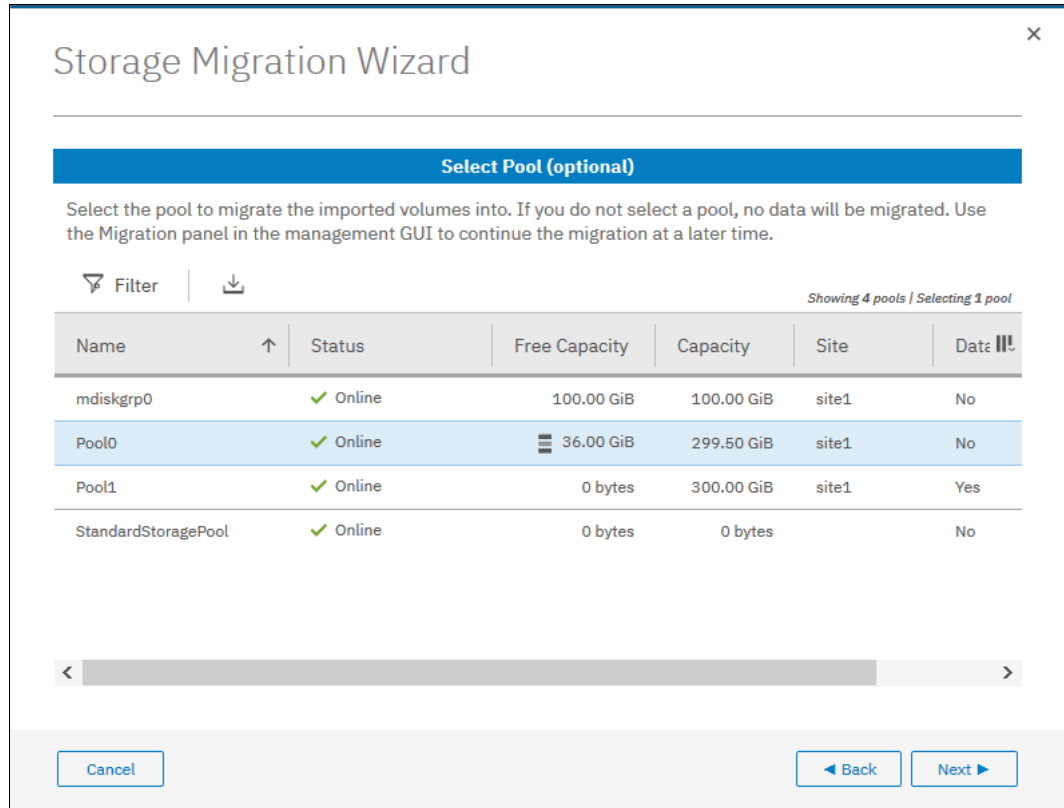


Figure 5-13 Select the pool to migrate the MDisk to

At this point the migration will begin. The migration task will continue running in the background and it uses the volume mirroring function to place a copy of the image-mode volumes in the selected storage pool.

22. Click **Finish** to close the IBM Spectrum Virtualize Migration wizard as shown in Figure 5-14.

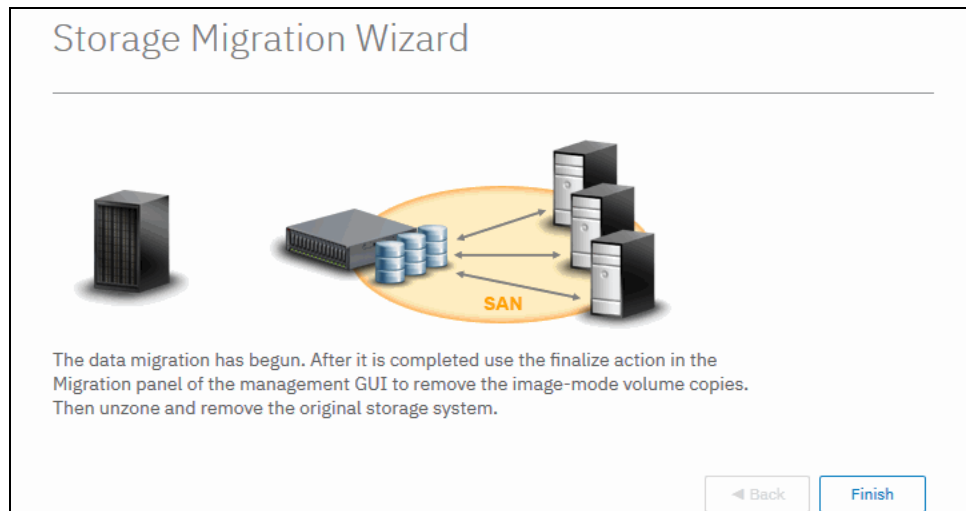


Figure 5-14 Migration is started

23. Migrations which were initiated using the IBM Spectrum Virtualize Migration wizard will continue running until they are complete. The migration progress is displayed in the Storage Migration window, as shown in Figure 5-15. The target storage pool and the progress of the volume copy synchronization is also displayed there.




Figure 5-15 Ongoing Migration is listed in the Storage Migration Window

24. To check the progress of the migration using the command-line interface, enter the `lsvdisksyncprogress` command, as shown in Figure 5-16.




Figure 5-16 CLI command showing migration progress

Note: Progress of the migration can also be monitored within the Running Tasks view in the IBM Spectrum Virtualize user interface.

25. After the migrations have completed, select all of the migrations that are ready and that you want to finalize.

Right-click the selection, and click **Finalize**, as shown in Figure 5-17.

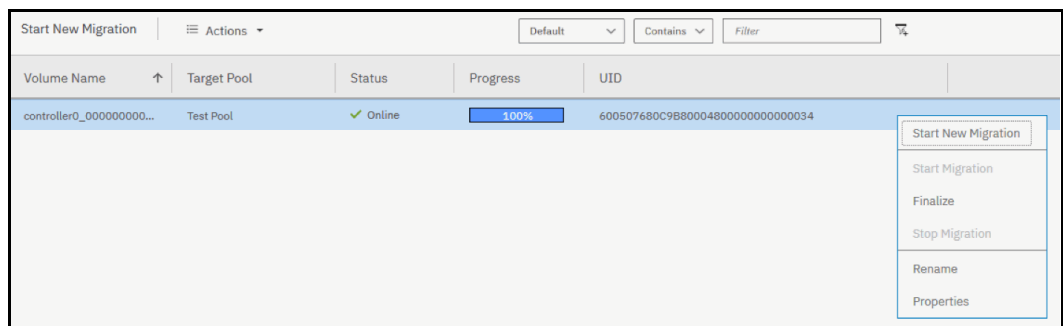


Figure 5-17 Finalizing the migration

26. During the Finalize process you are asked to confirm the action.

This is because finalizing a migration will remove the MDisk from the Migration Pool. The secondary MDisk remains in the destination pool and becomes the primary MDisk which will be used by the hosts.

Figure 5-18 is an example of the confirmation message.

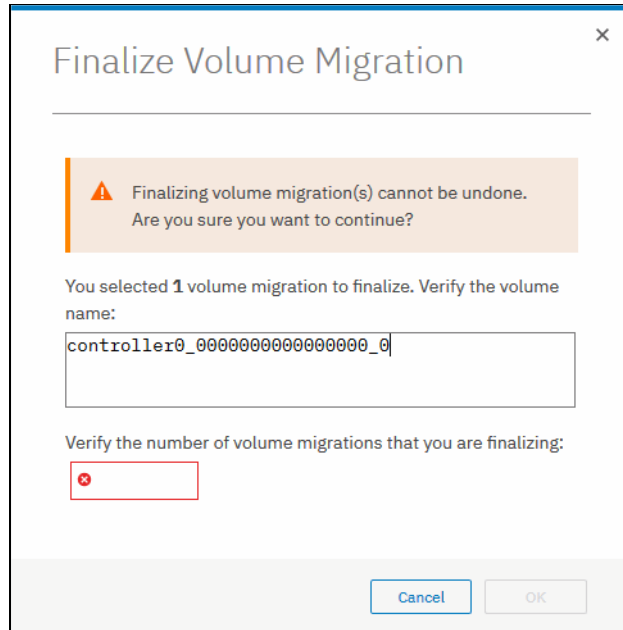



Figure 5-18 Migration finalization confirmation

27. When finalized, the image-mode copies of the volumes are deleted and the associated MDisks are removed from the migration pool. The status of those MDisks returns to unmanaged. You can verify the status of the MDisks by going to **Pools** → **External Storage**, as shown in Figure 5-19. In the example, mdisk3 has been migrated and finalized, and it displays as unmanaged in the external storage window.

Name	State	Capacity	Mode	Site	Pool	Storage System
controller7	✓ Online	IBM 2145	Serial Number: 2076	Site: site1	WWNN: 500507680000BECC	
controller2	✓ Online	IBM 2145	Serial Number: 2076	Site: site1	WWNN: 500507680300C1BD	
controller5	✓ Online	IBM 2145	Serial Number: 2076	Site: site2	WWNN: 50050768030026F0	
controller0	✓ Online	IBM 2145	Serial Number: 2076	Site: site1	WWNN: 50050768030026F1	
controller4	✓ Online	IBM 2145	Serial Number: 2076	Site: site2	WWNN: 500507680300C1BC	
controller1	✓ Online	IBM 2145	Serial Number: 2076	Site: site1	WWNN: 500507680300CB88	
controller6	✓ Online	IBM 2145	Serial Number: 2076	Site: site1	WWNN: 500507680000BECC	
mdisk3	✓ Online		300.00 GiB	Unmanaged	site1	controller6
mdisk2	✓ Online		300.00 GiB	Managed	site1	Test Pool controller6
mdisk1	✓ Online		200.00 GiB	Managed	site1	Pool0 controller6
site3_controller	✓ Online	IBM 2145	Serial Number: 2076	Site: site3	WWNN: 500507680300CB89	

Figure 5-19 External Storage MDisks window

All the steps that are described in the IBM Spectrum Virtualize Migration wizard can be performed manually by using the IBM Spectrum Virtualize GUI or command-line interface, but it is better to use the IBM Spectrum Virtualize Migration wizard as a guide if completing the migration using one of these other methods to ensure all necessary steps are completed correctly.



Migration alternatives in IBM Spectrum Virtualize environments

This chapter explains various methods and illustrated use cases for migrating data within an IBM Spectrum Virtualize Environment and, if necessary, can provide alternatives to the method presented in Chapter 5, “Using IBM Spectrum Virtualize Migration wizard” on page 49.

The methods explained in this chapter might be preferred when an IBM Spectrum Virtualize system is already implemented in your environment and you want to retain the virtualization layer provided by it on top of the new system(s) being brought in.

All of the functions used for the following scenarios presented in this chapter are standard components of the IBM Spectrum Virtualize software and available to all supported devices.

The chapter contains the following sections:

- ▶ Migrating back-end storage using MDisk migration
- ▶ Migrating volumes between pools using Volume Mirroring
- ▶ Migrating volume copies between pools

Finally, we have examined a scenario in which the goal is to export your data from the existing IBM Spectrum Virtualize system to other storage systems to get rid of the virtualization layer. This scenario would mostly apply when the IBM Spectrum Virtualize system in place is a SAN Volume Controller that you want to decommission. That particular use case is covered in the last section of this chapter:

- ▶ Migrating out of a IBM Spectrum Virtualize system through export to Image Mode

6.1 Migrating back-end storage using MDisk migration

One way to decommission an existing, already virtualized storage system is to first add the new storage system as back-end system to the IBM Spectrum Virtualize system and add the new MDisks to an existing MDisk group, then remove the original MDisks. The IBM Spectrum Virtualize software will take care of the movement of data to the new MDisks. This process is concurrent (host I/Os do not need to be stopped) and can be an easy way to migrate from one back-end storage system to another.

This method is particularly favorable when you have many volumes to move. In our example, we have an IBM FlashSystem A9000R (A9000R) virtualized by an SAN Volume Controller, and we have eight MDisks from the A9000R in a MDisk group. We want to decommission the A9000R and replace it with a FlashSystem 9100, while keeping the SAN Volume Controller in place.

The existing storage under SAN Volume Controller (A9000R MDisks, in this case) can be seen by selecting **Pools** → **MDisk per Pool** view in the SAN Volume Controller GUI, as shown in Figure 6-1.

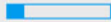
Name	State	Effective Capacity
Unassigned MDisks (0)		
Pool_A9000	✓ Online	 10.00 TiB / 64.00 TiB (16%)
ITSO_A9000_Mdisk_1	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_4	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_6	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_5	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_3	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_2	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_7	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_8	✓ Online	8.00 TiB

Figure 6-1 A9000R MDisks

Note: It is best practice to use a minimum of 8 MDisks (or increments of 8) from any storage that is presented to a SAN Volume Controller.

Preparation of the new storage system

Perform the following tasks on the new storage system (a FlashSystem 9100 in our scenario):

1. From the FlashSystem 9100 GUI, create eight *volumes* from your available pool;
2. Create appropriate zoning, cabling or iSCSI connections between SAN Volume Controller and FlashSystem 9100;
3. Create a *hostcluster* called SAN Volume Controller in FlashSystem 9100.
4. Create *hosts* of the *generic* type for all SAN Volume Controller nodes and assign them to the *hostcluster*;
5. Map the eight *volumes*, in this case from FlashSystem 9100, to the SAN Volume Controller;

6. The eight volumes will be seen in by the SAN Volume Controller as Unmanaged MDisks.

Discovering the new storage volumes in SAN Volume Controller

Go to the **External Storage** view in the IBM Spectrum Virtualize GUI and select **Discover Storage**, as shown in Figure 6-2



Figure 6-2 Discover Storage in SAN Volume Controller

The MDisks from the new back-end system (FlashSystem 9100) will be found under **Unassigned MDisks**, as shown in Figure 6-3.

Name	State	Effective Capacity
Unassigned MDisks (8)		
ITSO_FS9150_Mdisk_5	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_1	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_8	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_2	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_7	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_3	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_4	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_6	✓ Online	8.00 TiB

Figure 6-3 FlashSystem 9100 MDisks newly discovered

Click **Assign**, as shown in Figure 6-4.



Figure 6-4 Assigning MDisks

In the **Assign External Storage** window, select the pool to add the new MDisks to - the existing Pool_A9000 in this case, as shown in Figure 6-5. Select all MDisks, and make sure to assign them the correct **Tier** specification (Tier 0 Flash in this case).

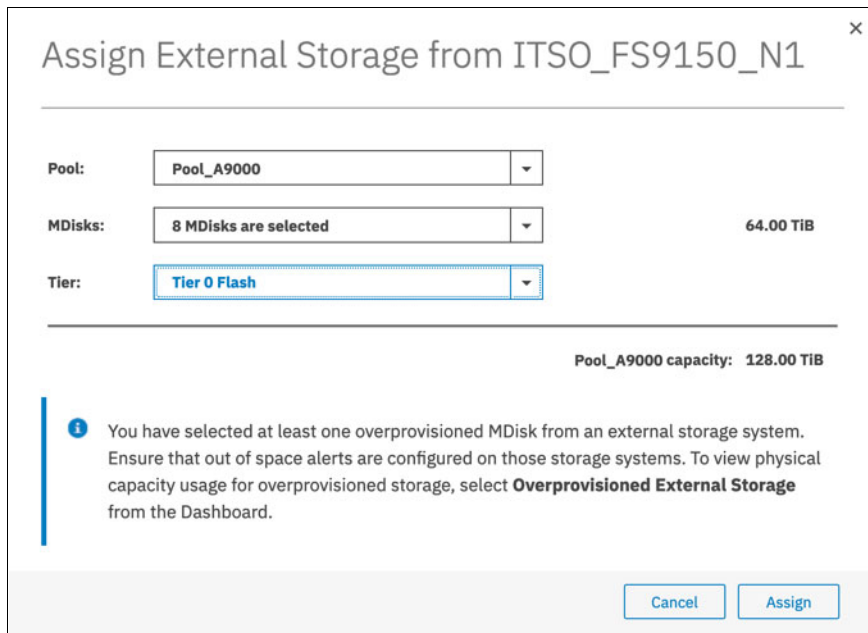


Figure 6-5 Assigning new MDisks to existing MDisk Pool

Now, both the old MDisks from A9000R and the new MDisks from FlashSystem 9100 are in the same MDisk pool, as shown in Figure 6-6.

Pool	Status	Capacity
Pool_A9000	Online	10.00 TiB / 128.00 TiB (8%)
ITSO_A9000_Mdisk_1	Online	8.00 TiB
ITSO_A9000_Mdisk_4	Online	8.00 TiB
ITSO_A9000_Mdisk_6	Online	8.00 TiB
ITSO_A9000_Mdisk_5	Online	8.00 TiB
ITSO_A9000_Mdisk_3	Online	8.00 TiB
ITSO_A9000_Mdisk_2	Online	8.00 TiB
ITSO_A9000_Mdisk_7	Online	8.00 TiB
ITSO_A9000_Mdisk_8	Online	8.00 TiB
ITSO_FS9150_Mdisk_8	Online	8.00 TiB
ITSO_FS9150_Mdisk_6	Online	8.00 TiB
ITSO_FS9150_Mdisk_1	Online	8.00 TiB
ITSO_FS9150_Mdisk_2	Online	8.00 TiB
ITSO_FS9150_Mdisk_4	Online	8.00 TiB
ITSO_FS9150_Mdisk_5	Online	8.00 TiB
ITSO_FS9150_Mdisk_3	Online	8.00 TiB
ITSO_FS9150_Mdisk_7	Online	8.00 TiB

Figure 6-6 All MDisks in same MDisk pool

The next step in the migration process is to remove the A9000R MDisks from the pool. The SAN Volume Controller will take care of automatically moving data to the new MDisks provisioned from FlashSystem 9100.

Important: Before removing the old A9000R MDisks, make sure the new MDisks are large enough to host the migrating data. Also consider that you might not have the same data reduction ratio on the new MDisks, so make sure extra physical space is available at the back-end storage level.

Select the MDisks from A9000R, right click and select **Remove**, as shown in Figure 6-7. This will start the removal of the MDisks from the Pool; the IBM Spectrum Virtualize software will move all data to the remaining MDisks in the Pool. This process is sequential (one MDisk at a time) and fully automatic, and cannot be stopped once started.

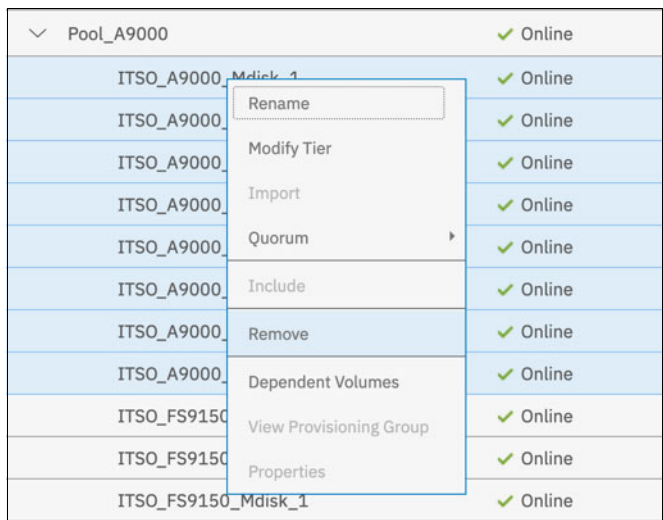


Figure 6-7 Remove MDisks from MDisk Pool

The process can be monitored from the **Running Tasks** view, as shown in Figure 6-8

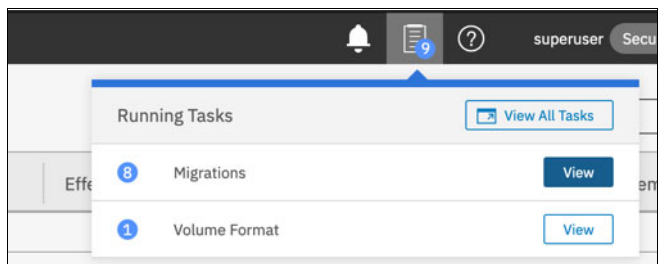


Figure 6-8 Running Tasks

All the MDisks removal operations are listed, with their progress, as shown in Figure 6-9 on page 70.

Name	Progress
MDisk ITSO_A9000_Mdisk_1 → Pool Pool_A9000	3%
MDisk ITSO_A9000_Mdisk_3 → Pool Pool_A9000	0%
MDisk ITSO_A9000_Mdisk_2 → Pool Pool_A9000	0%
MDisk ITSO_A9000_Mdisk_5 → Pool Pool_A9000	0%
MDisk ITSO_A9000_Mdisk_4 → Pool Pool_A9000	0%
MDisk ITSO_A9000_Mdisk_8 → Pool Pool_A9000	0%
MDisk ITSO_A9000_Mdisk_7 → Pool Pool_A9000	0%
MDisk ITSO_A9000_Mdisk_6 → Pool Pool_A9000	0%

Figure 6-9 Running Migration progress

When all the running tasks are completed, the only MDisks available in the Pool are the ones from the FlashSystem 9100, as shown in Figure 6-10. All the data of the volumes allocated from this pool now resides only on these MDisks.

Name	State	Effective Capacity
Unassigned MDisks (8)		
Pool_A9000	✓ Online	10.00 TiB / 64.00 TiB (16%)
ITSO_FS9150_Mdisk_8	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_6	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_1	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_2	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_4	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_5	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_3	✓ Online	8.00 TiB
ITSO_FS9150_Mdisk_7	✓ Online	8.00 TiB

Figure 6-10 Only FlashSystem 9100 MDisks in MDisk Pool

Rename the Pool so it reflects the new FlashSystem 9100 storage, as shown in Figure 6-11.

Rename Pool

Pool_A9000

Enter new name:

Pool_FS9150

Cancel
Rename

Figure 6-11 Renaming of the MDisk Pool

The removed MDisks enter the Unassigned state, and can be safely unmapped from the SAN Volume Controller; finally, the old backend system can be disconnected from the SAN Volume Controller and decommissioned.

6.2 Migrating volumes between pools using Volume Mirroring

This section shows how to use the *Volume Mirroring* feature in IBM Spectrum Virtualize to migrate data from one storage device to another. This is also a common scenario if you already have IBM Spectrum Virtualize system such as SAN Volume Controller in place in your environment, and want to non-disruptively migrate your data from one back-end storage system to another, for example.

The process is fully transparent to the application. Volume Mirroring is included in the base virtualization license, so you do not need to purchase any additional licenses. This function can also be used to migrate from a fully allocated volume to a thin provisioned volume in a regular pool, or to a thin provisioned, compressed or deduplicated volume in a Data Reduction Pool (DRP).

Using the IBM Spectrum Virtualize volume mirroring feature for data migration involves the following basic steps:

1. Create a new Pool and add capacity from the new storage system to it;
2. Add a target copy in the newly created Pool for every volume with a copy in the storage system to be replaced;
3. Run synchronization and wait for it to complete;
4. Remove the source volume copy/copies;
5. (Optional) Remove the old storage system's MDisks and Pool.

Tip: When using the IBM Spectrum Virtualize volume mirroring feature, the data migration can be stopped at any time without compromising data consistency on the primary copy of the volumes.

6.2.1 Use case

In the environment used for this example, a SAN Volume Controller is already in place with a FlashSystem A9000R back-end storage. A new back-end FlashSystem 9150 storage system will be added as data migration target. The procedure explained in this use case is independent of the application host operating system.

Figure 6-12 shows the initial configuration, where our example volume is contained in pool Poo1_A9000.

Name	State	Synchronized	Pool
ITSO_AIX_rootvg_001	✔ Online		Pool_A9000

Figure 6-12 Volume to be migrated

Pool Poo1_A9000 comprises several MDisks provided by the same A9000R backend storage system, as shown in Figure 6-13 on page 72.


Name	State	Effective Capacity
Unassigned MDisks (0)		
Pool_A9000	✓ Online	 10.00 TiB / 64.00 TiB (16%)
ITSO_A9000_Mdisk_1	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_4	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_6	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_5	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_3	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_2	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_7	✓ Online	8.00 TiB
ITSO_A9000_Mdisk_8	✓ Online	8.00 TiB

Figure 6-13 Pool composition

In the **Pools** window, select **Create Pool** to create the target pool, as shown in Figure 6-14.

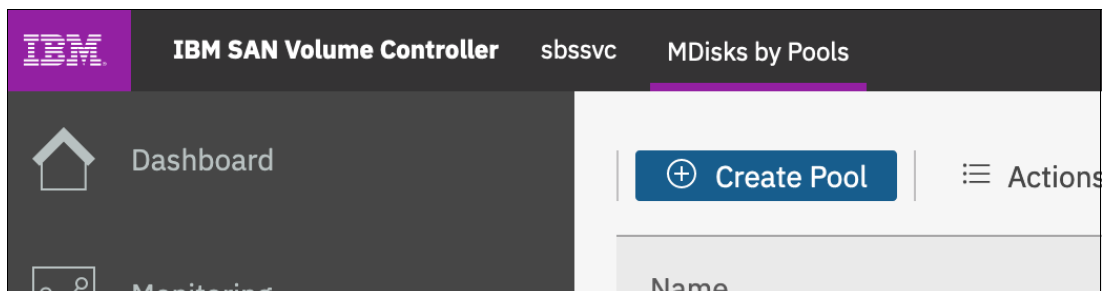


Figure 6-14 Create pool action button

The **Create Pool** page allows you to decide the extent size for the new pool; you can also choose to create a regular pool or a Data Reduction Pool by ticking the **Enable Data Reduction** check box (Figure 6-15).

Create Pool
×

Name:

Extent size: ▼

Maximum addressable capacity: 4.00 PiB

Data reduction: Enable

ⓘ Need Help
Cancel
Create

Figure 6-15 Create Pool page

Note: The migration procedure is the same regardless of extent size or type of target pool created.

For more information about extents, pool capacity and Data Reduction Pools see the IBM Redbooks publication *Implementing the IBM System Storage SAN Volume Controller*, SG24-7933.

Populate the pool with the new MDisks provided by the FlashSystem 9150 storage system by simply clicking the **Assign** button at the bottom of the **Pools** → **MDisks by Pool** view, as shown in Figure 6-16.

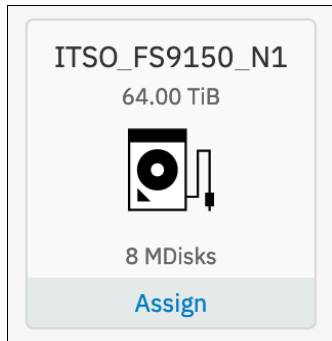


Figure 6-16 Assign MDisks to the newly created pool

In the **Assign External Storage** page, select the newly created pool and the MDisks you want to use for it (Figure 6-17); make sure to select the correct tier for the MDisks being assigned. In this example, MDisks provided by FlashSystem 9150 are marked as `tier0_flash`.

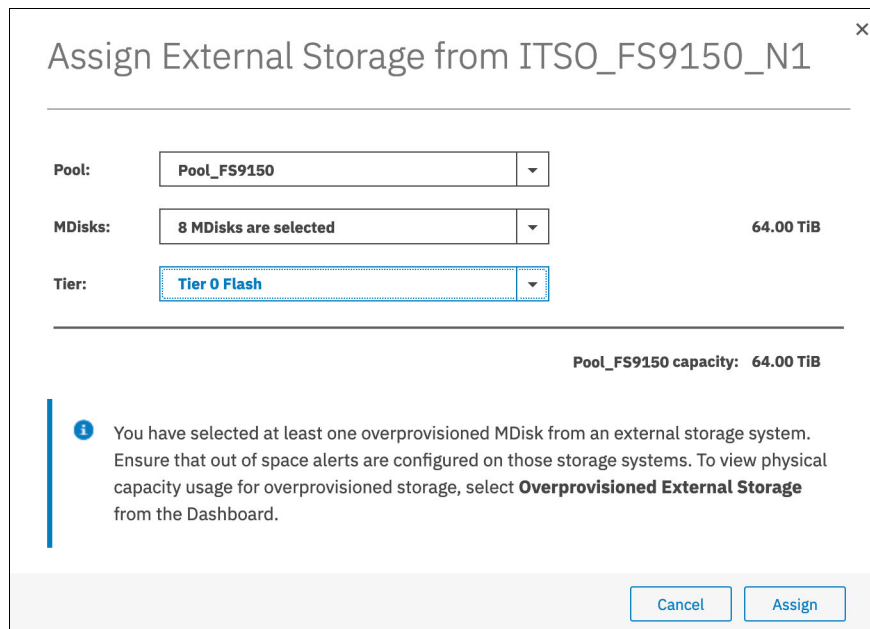


Figure 6-17 Assign External Storage page

Figure 6-18 shows the pool list after the MDisk assignment.

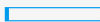
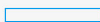
Name	State	Effective Capacity
Unassigned MDisks (0)		
> Pool_A9000	✓ Online	 50.00 GiB / 64.00 TiB (0%)
> Pool_FS9150	✓ Online	 0 bytes / 64.00 TiB (0%)

Figure 6-18 Pool list

In the **Volumes** view, select the volume you want to migrate. Right-click that volume and select **Add Volume Copy...** from the pop-up menu, as shown in Figure 6-19.

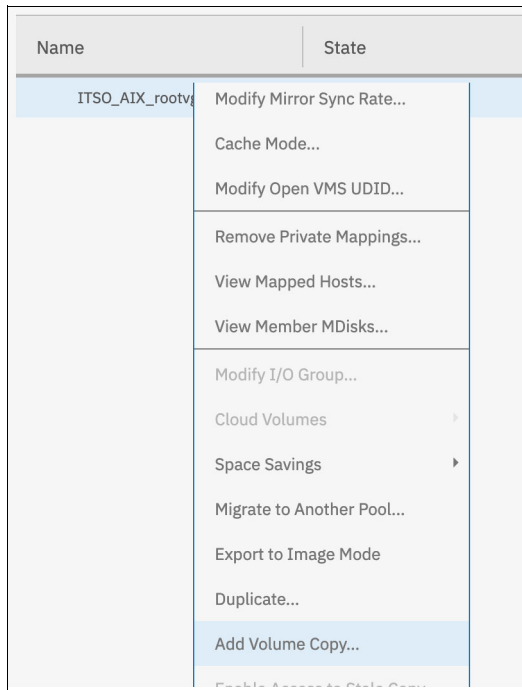


Figure 6-19 Right-click volume menu

The **Add Volume Copy** page displays (Figure 6-20). Select the newly created pool as target for Copy 1.

Note: While performing the migration, you can also decide to apply Capacity Savings settings such as thin provisioning, compression and deduplication to the new Volume Copy.

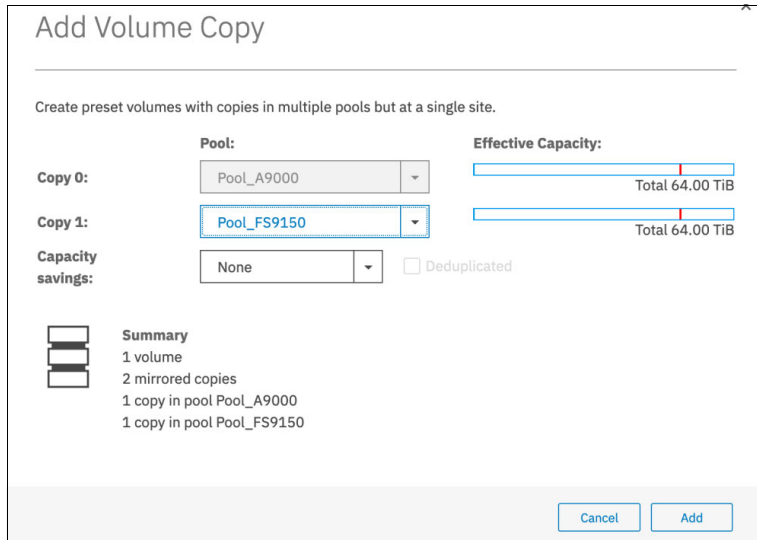


Figure 6-20 Add Volume Copy page

Click **Add** to create the new mirrored volume copy. When the command finishes, the copy synchronization task starts in the background. You can check the status of the synchronization by clicking the **Tasks** menu and selecting **View** (Figure 6-21), or by going to **Monitoring** → **Background Tasks**.




Figure 6-21 Tasks quick-menu

The volume synchronization task progress is displayed, along with the estimated time remaining, as shown in Figure 6-22.



Figure 6-22 Volume Synchronization task details

You can decide the speed at which data is copied from one Volume Copy to the other by modifying the *Mirror Sync Rate* attribute of the volume, as illustrated in Figure 6-23 on page 76. The corresponding right-click menu option allows you to choose between the possible rates, ranging 0 - 64 MBps (with a default of 2 MBps), as shown in Figure 6-24 on page 76.

Important: always plan carefully when and how to run volume copy synchronizations. The synchronization process adds additional strain to your back-end storage. Select an appropriate Sync Rate based on the number of volume synchronizations you are running in parallel and the load on your IBM Spectrum Virtualize system, and monitor closely its performances until the migration is completed.

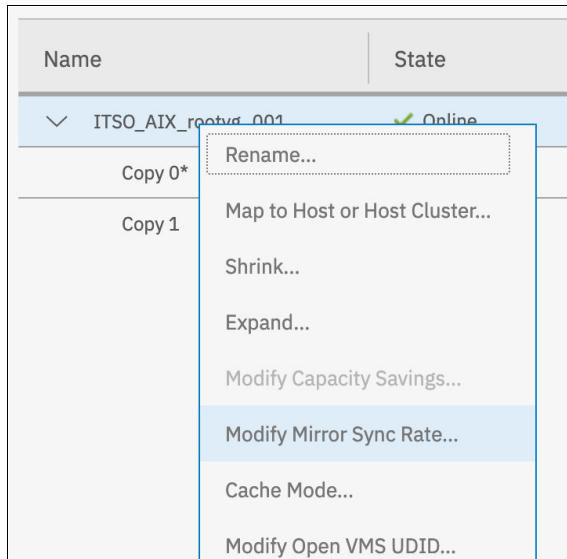


Figure 6-23 Volume right-click to Modify Sync Rate

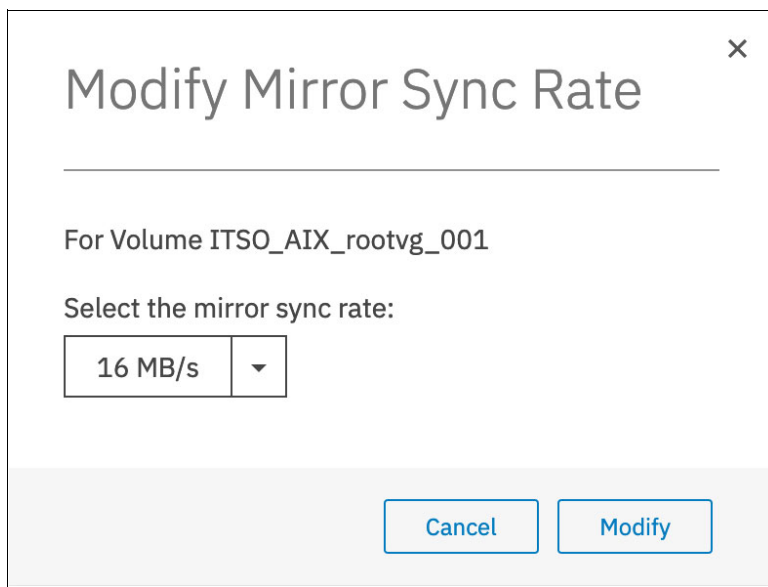


Figure 6-24 Modify Mirror Sync Rate page

While data is being copied in the background, the newly created volume copy will be displayed as **Synchronized: No**, as shown in Figure 6-25.

Name	State	Synchronized
ITSO_AIX_rootvg_001	✓ Online	
Copy 0*	✓ Online	Yes
Copy 1	✓ Online	No

Figure 6-25 Volume copy state while synchronization is running

After the background synchronization process is completed, both copies will be displayed as **Synchronized: Yes**, as shown in Figure 6-26.

Name	State	Synchronized
ITSO_AIX_rootvg_001	✓ Online	
Copy 0*	✓ Online	Yes
Copy 1	✓ Online	Yes

Figure 6-26 Volume synchronization state after the background task has completed

You can proceed with the deletion of the source copy of the volume from the old back-end storage. Right-click the volume Copy (marked with a *) and select **Delete**, as shown in Figure 6-27.

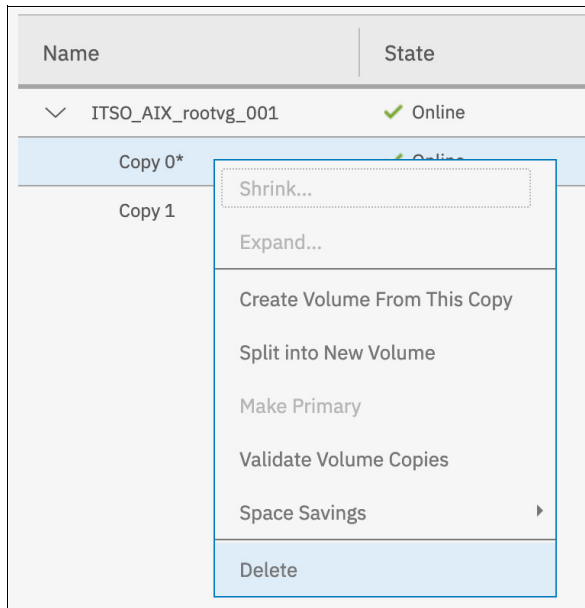


Figure 6-27 Volume copy right-click delete

Click **Yes** to confirm the deletion (Figure 6-28).




Figure 6-28 Volume copy deletion confirmation page

The volumes view updates to display the only remaining copy of the volume in the pool, as shown in Figure 6-29.

Name	State	Synchronized	Pool
ITSO_AIX_rootvg_001	✓ Online		Pool_FS9150

Figure 6-29 Final state of the volume

When you have completed the migration of all volumes from one pool to the other, you can proceed with the removal of all the MDiskS from the A9000R back-end storage system and the deletion of the pool.

6.3 Migrating volume copies between pools

An alternative way of migrating volume data from one pool to another in IBM Spectrum Virtualize is to use the Migrate Volume Copy feature. As opposed to copying data from one pool to the other, the Migrate Volume Copy function *moves* extents between the pools. This action means that volume data is physically removed from one pool, and placed into the other.

The mirrored volumes approach, as explained in section 6.2, “Migrating volumes between pools using Volume Mirroring” on page 71, is generally preferred because at any time you are assured of having at least one non-impacted volume copy, which can only be affected by failures in the source pool. Using mirrored volumes can also allow you to speed up, down, stop and restart the migration at will. However, there are some cases when the mirrored volumes approach is not applicable - for example, if you want to migrate a volume that already has two copies and you want to maintain high-availability during the process. In that case, the mirrored volumes approach would require you to delete one of the copies before creating the new one in the target pool, and leave you with only one valid copy of the data until the synchronization process is complete.

6.3.1 Use case

In this use case, we assume that we have a SAN Volume Controller in Enhanced Stretched Cluster configuration, with two back-end FlashSystem A9000 storage systems (one in each site of the stretched architecture) providing capacity to two different pools, Pool_A9000_1 and Pool_A9000_2. As part of a technology refresh project, we want to migrate all data contained in these pools to two corresponding new pools, with the same extent size, receiving capacity from two FlashSystem 9150 systems - Pool_FS9150_1 and Pool_FS9150_2. We assume also that all back-end storage has been correctly connected and configured, and all pools are available for use.

In our example volume, ITSO_AIX_rootvg_001, already has two copies in this stretched configuration for high availability purposes (Figure 6-30 on page 79), which we want to maintain during the migration process from the old back-end storage to the new.

Name	State	Synchronized	Pool
ITSO_AIX_rootvg_001	Online		Pool_A9000_1
Copy 0*	Online	Yes	Pool_A9000_1
Copy 1	Online	Yes	Pool_A9000_2

Figure 6-30 Initial volume configuration

The migration process works on one copy at a time; First, select **Migrate to Another Pool...** from the volume's right click menu, as shown in Figure 6-31.

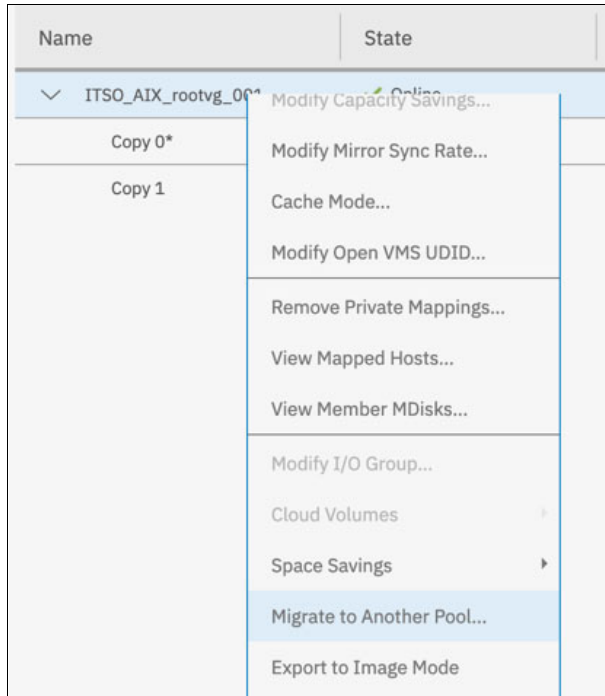


Figure 6-31 Volume right-click menu

The **Migrate Volume Copy** page displays (Figure 6-32 on page 80). Here, select the first volume copy to migrate (Copy 0, for example), and the target Pool to which you want to migrate that copy. In our example, we choose Pool_FS9150_1. Finally, click **Migrate** to start the migration process.

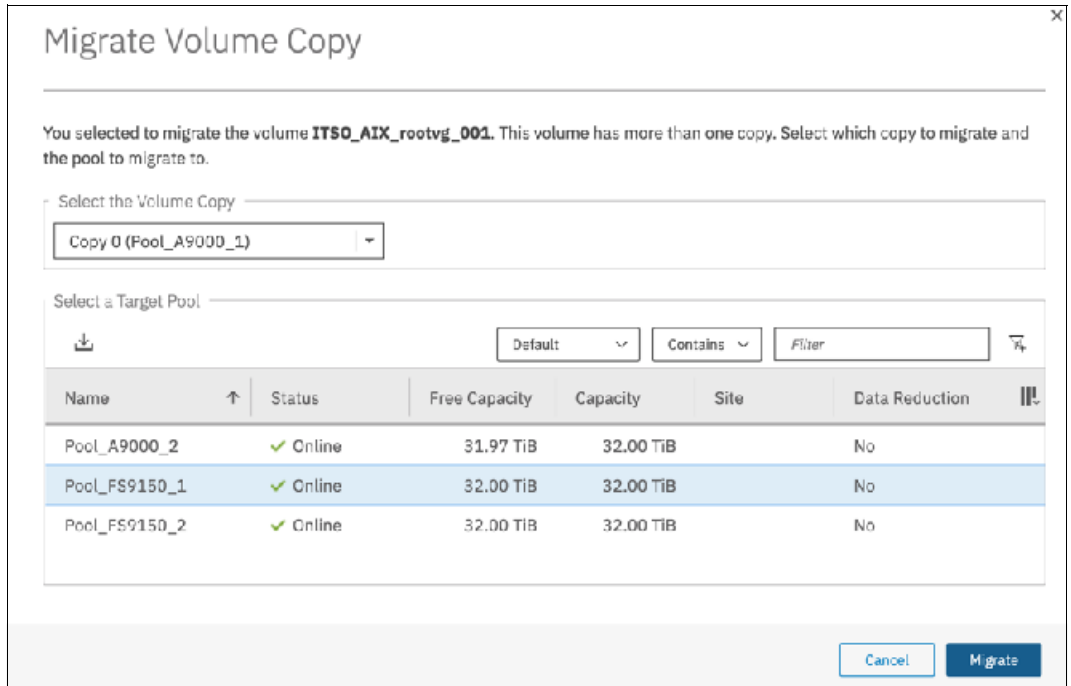


Figure 6-32 Migrate Volume Copy page

You can check the status of the migration by clicking the **Tasks** menu and selecting **View** (Figure 6-33), or by going to **Monitoring** → **Background Tasks**.

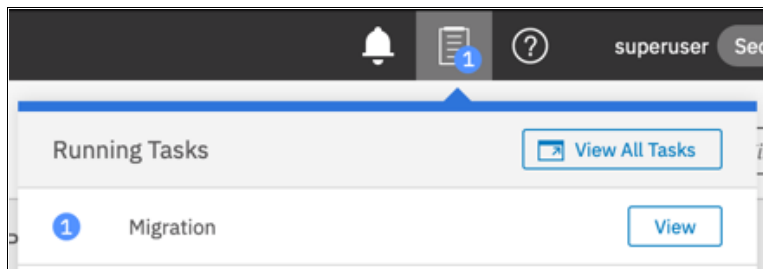


Figure 6-33 Tasks menu

The migration activity is displayed alongside its completion percentage (Figure 6-34).




Figure 6-34 Migration task progress

When the migration has completed, the volume copy information is updated to show the new pool for volume Copy 0, as shown in Figure 6-35 on page 81.

Name	State	Synchronized	Pool
ITSO_AIX_rootvg_001	✓ Online		Pool_FS9150_1
Copy 0*	✓ Online	Yes	Pool_FS9150_1
Copy 1	✓ Online	Yes	Pool_A9000_2

Figure 6-35 Volume status after the first copy migration

Repeat the process for the other volume copy. In our example, we now migrate Copy 1 from Pool_A9000_2 to Pool_FS9150_2 (Figure 6-36).

Migrate Volume Copy

You selected to migrate the volume **ITSO_AIX_rootvg_001**. This volume has more than one copy. Select which copy to migrate and the pool to migrate to.

Select the Volume Copy

Copy 1 (Pool_A9000_2)

Select a Target Pool

Default Contains Filter

Name	↑	Status	Free Capacity	Capacity	Site	Data Reduction
Pool_A9000_1		✓ Online	32.00 TiB	32.00 TiB		No
Pool_FS9150_1		✓ Online	31.97 TiB	32.00 TiB		No
Pool_FS9150_2		✓ Online	32.00 TiB	32.00 TiB		No

Cancel Migrate

Figure 6-36 Migrate Volume Copy page

When this second migration is finished, the volume has both copies in the new pools (Figure 6-37). The procedure can be repeated or run in parallel for other volumes in the pools; when completed for all volumes in the old pools, you can proceed with the removal of all the MDisks from the A9000 back-end storage systems and the deletion of the old pools.

Name	State	Synchronized	Pool
ITSO_AIX_rootvg_001	✓ Online		Pool_FS9150_1
Copy 0*	✓ Online	Yes	Pool_FS9150_1
Copy 1	✓ Online	Yes	Pool_FS9150_2

Figure 6-37 Final configuration of the migrated volume

6.4 Migrating out of a IBM Spectrum Virtualize system through export to Image Mode

The Image Mode functionality can be used to migrate out of a IBM Spectrum Virtualize storage system to any other storage system, excluding the ones that cannot be virtualized, as explained in the Restriction note that follows.

Restriction: IBM SAN Volume Controller and IBM FlashSystem V9000 can only be used in **replication** layer. This means that they cannot be virtualized by other IBM Spectrum Virtualize systems. FlashSystem 7200, 9100 and 9200, however, can also be used in **storage** layer and virtualized by other IBM Spectrum Virtualize systems. For more information about layers, see:

https://www.ibm.com/support/knowledgecenter/STSLR9_8.3.1/com.ibm.fs9200_831.doc/svc_systemlayer.html

One common example of this scenario is the migration out of a SAN Volume Controller system to a non-SAN Volume Controller system, such as FlashSystem 9100. The scenario is depicted in Figure 6-38. This activity might be carried out for one of these reasons:

- ▶ You purchased a new storage system, such as a FlashSystem 9200, and are using the SAN Volume Controller to migrate data to it using the export to Image Mode functionality.
- ▶ You want to move a host and its data currently connected to the SAN Volume Controller system to a site where there is no SAN Volume Controller.

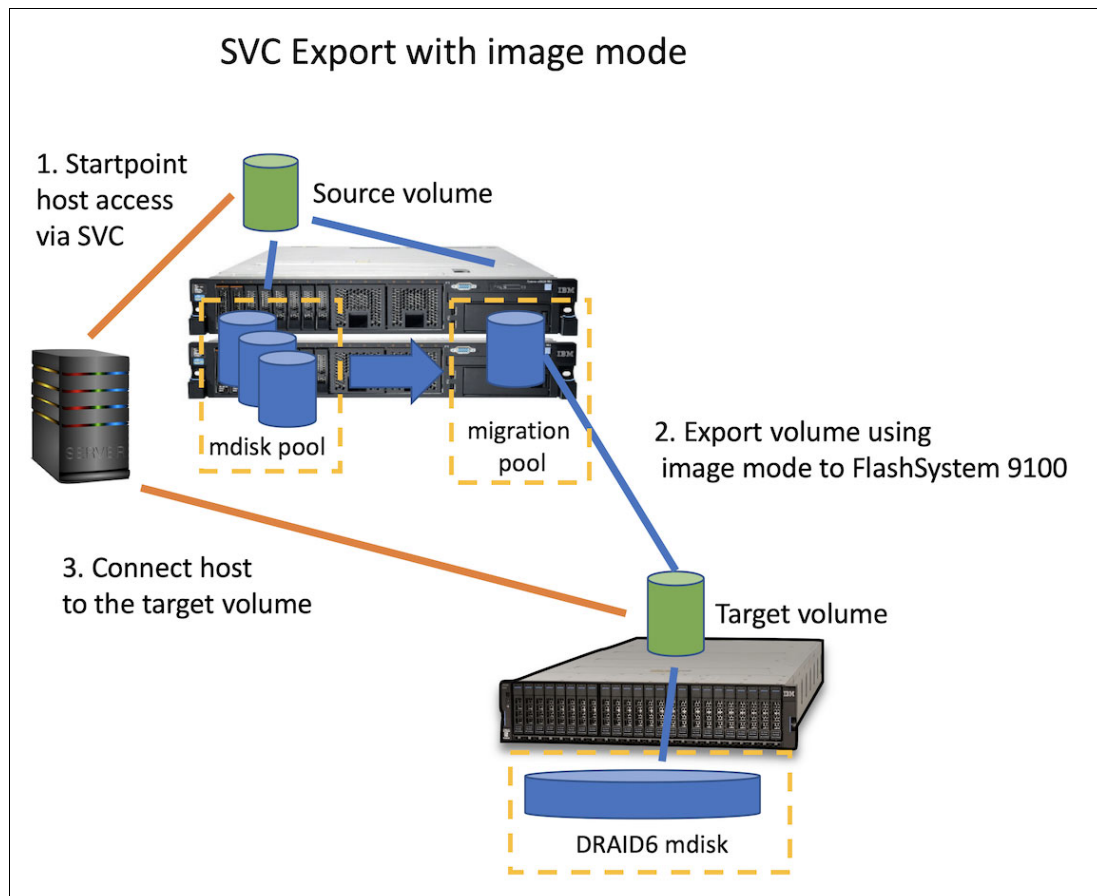


Figure 6-38 Scenario overview

To virtualize an IBM FlashSystem 9100 behind IBM SAN Volume Controller, both systems must be connected via SAN fabric, iSCSI network or direct attachment.

After connecting the two systems to your SAN fabric, make sure that the FlashSystem 9100 is discovered under the **Pools** → **External Storage** view in IBM SAN Volume Controller GUI. If that is not the case, re-check your cabling, SAN zoning or iSCSI connectivity.

6.4.1 Use case

In this example, we want to migrate one volume currently in use by a Linux host from IBM SAN Volume Controller to IBM FlashSystem 9100.

The 50 GiB disk is defined as a mirrored volume on the SAN Volume Controller, mapped to the host, which sees it as a multipath device; the device was formatted, partitioned and configured with an ext4 file system, and mounted at /itso. It contains some sample files as seen in Example 6-1.

Example 6-1 SAN Volume Controller volume seen as multipath device by the Linux host

```
[root@comp-1-31 itso]# multipath -ll
mpathc (3600507680c80011a80000000000000a) dm-5 IBM      ,2145
size=50G features='1 queue_if_no_path' hwhandler='0' wp=rw
|-+- policy='service-time 0' prio=50 status=active
|  ~- 6:0:3:0 sdi 8:128 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
   ~- 6:0:2:0 sdj 8:144 active ready running

[root@comp-1-31 itso]# df -h
Filesystem              Size  Used Avail Use% Mounted on
...
/dev/mapper/mpathc1     50G   53M   47G   1% /itso

[root@comp-1-31 itso]# ls /itso
bert.txt  chris.txt  duncan.txt  francesco.txt  harry.txt  lost+found  markus.txt
matt.txt  roger.txt  steve.txt
```

Tip: If the volumes you plan to migrate are not mirrored volumes, we advise you to use the Volume Mirroring feature of IBM Spectrum Virtualize to temporarily create a second synchronous copy, as explained in 6.2, “Migrating volumes between pools using Volume Mirroring” on page 71. As the export to Image Mode functionality physically moves extents of data from one back-end storage to the other, the additional copy can be retained as a backup strategy in case anything goes wrong in the process.

First, we need to present a target volume from the new Storage System (FlashSystem 9100) to the SAN Volume Controller. From the **Volumes** → **Volumes** view in FlashSystem 9100, select **Create Volumes** (Figure 6-39).




Figure 6-39 Create Volumes action button

In the **Create Volumes** page, make sure to set a volume capacity of equal or greater capacity than the source volume you want to migrate (in this example, 50 GiB), as shown in Figure 6-40.

The screenshot shows the 'Create Volumes' interface. The 'Basic' tab is active. The 'Pool' is set to 'Pool0'. The 'Capacity Details' section shows a progress bar and 'Total 66.54 TiB'. The 'Volume Details' section includes the following fields:

Quantity:	Capacity:	Name:
1	50 GiB	ITSO_LNX_Vol_001_ImageMode_Target

Below the 'Volume Details' section, there is a 'Capacity savings' dropdown set to 'None' and a 'Deduplicated' checkbox. At the bottom, there is an 'I/O group' dropdown set to 'Automatic', a 'Summary' button, and a 'Need Help' link. The footer contains 'Cancel', 'Create and Map', and 'Create' buttons.

Figure 6-40 Create Volumes page

Next, map the newly created volume on the FlashSystem 9100 to the SAN Volume Controller. If you have not done it before, you must first define the SAN Volume Controller as a Host or Host Cluster in the FlashSystem 9100.

To that end, right-click the volume and select **Map to Host or Host Cluster...**, then select the SAN Volume Controller (ITSO_SVC in our example) in the **Create Mapping** page, as shown in Figure 6-41 on page 85.

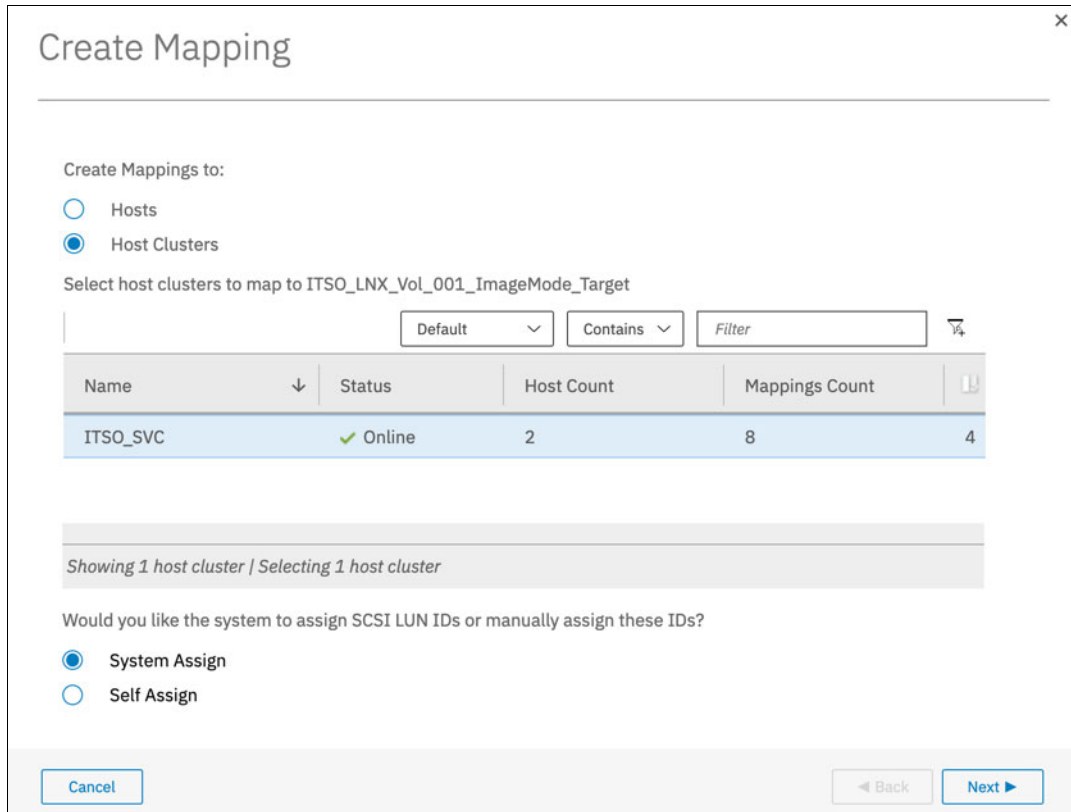


Figure 6-41 Map the target Volume to the SAN Volume Controller

Confirm the mapping, as shown in Figure 6-42.

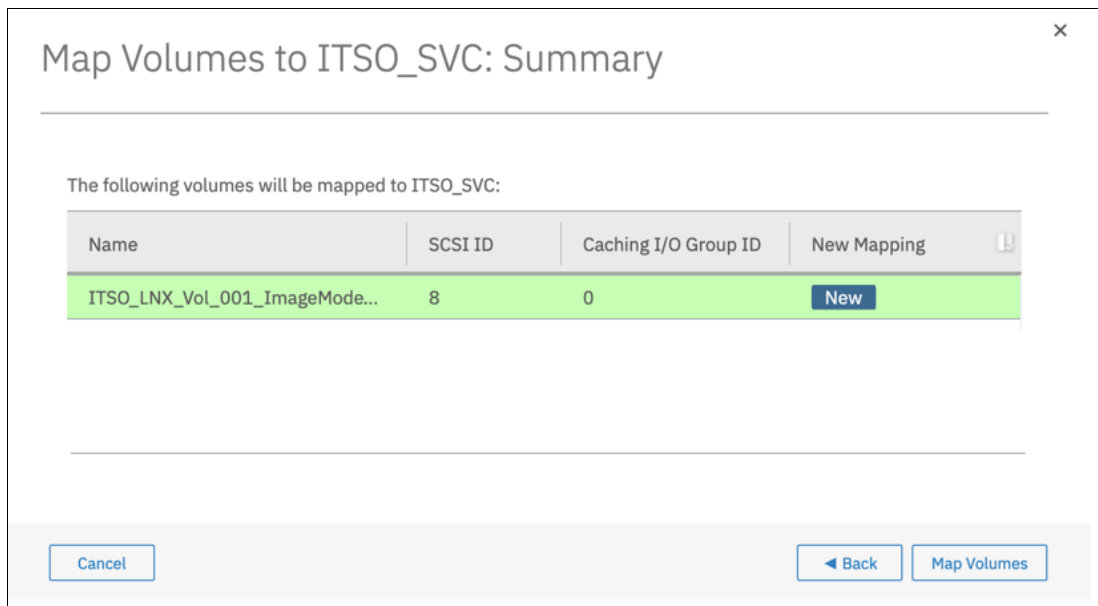


Figure 6-42 Confirm the mapping

Discover and configure the target volume in SAN Volume Controller

In the SAN Volume Controller GUI, go to the **External Storage** view and select **Discover Storage** from the **Actions** menu to rescan for new MDisk, as shown in Figure 6-43.

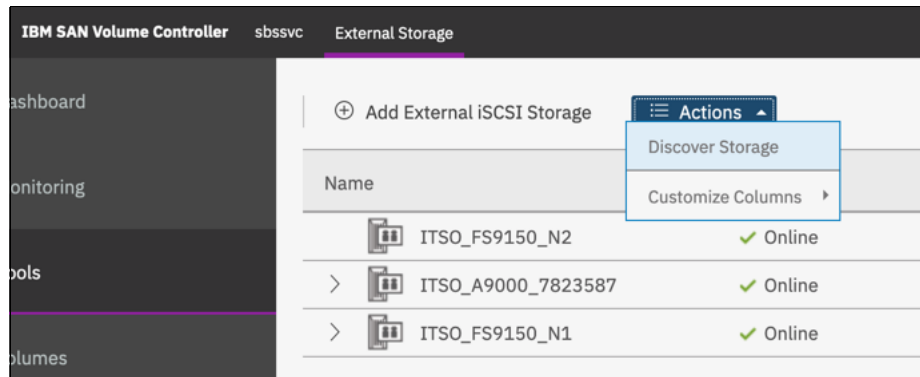


Figure 6-43 Discover external storage

A new Managed Disk (MDisk) should appear under the FlashSystem 9100 controller.

To make sure you are using the correct target disk, verify the UID of the volume as defined on the FlashSystem 9100 (via the volume **Properties** page - Figure 6-44 and Figure 6-45 on page 87).

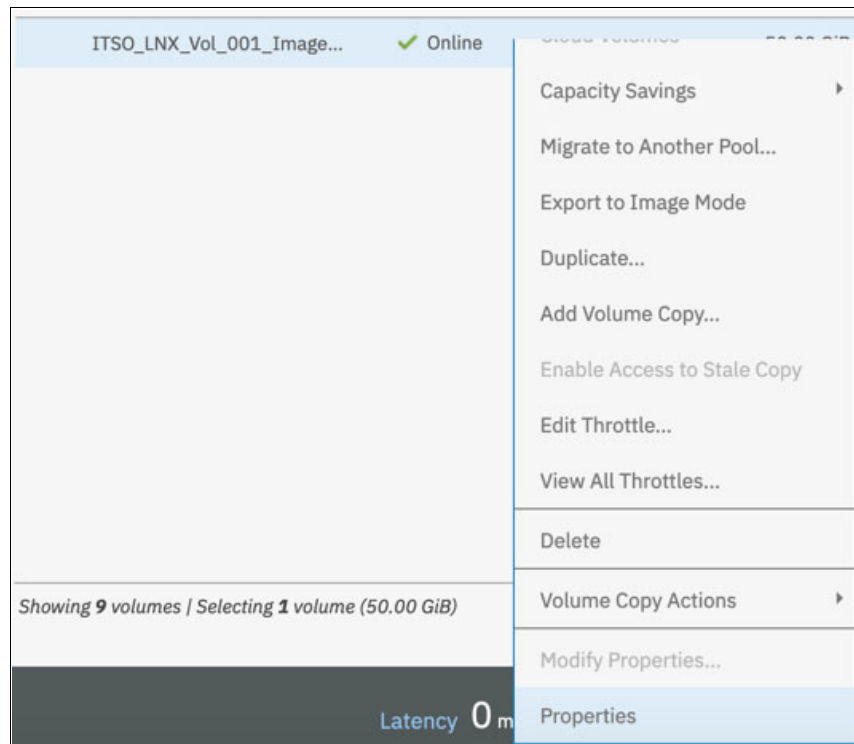


Figure 6-44 Volume right-click menu on the FlashSystem 9100

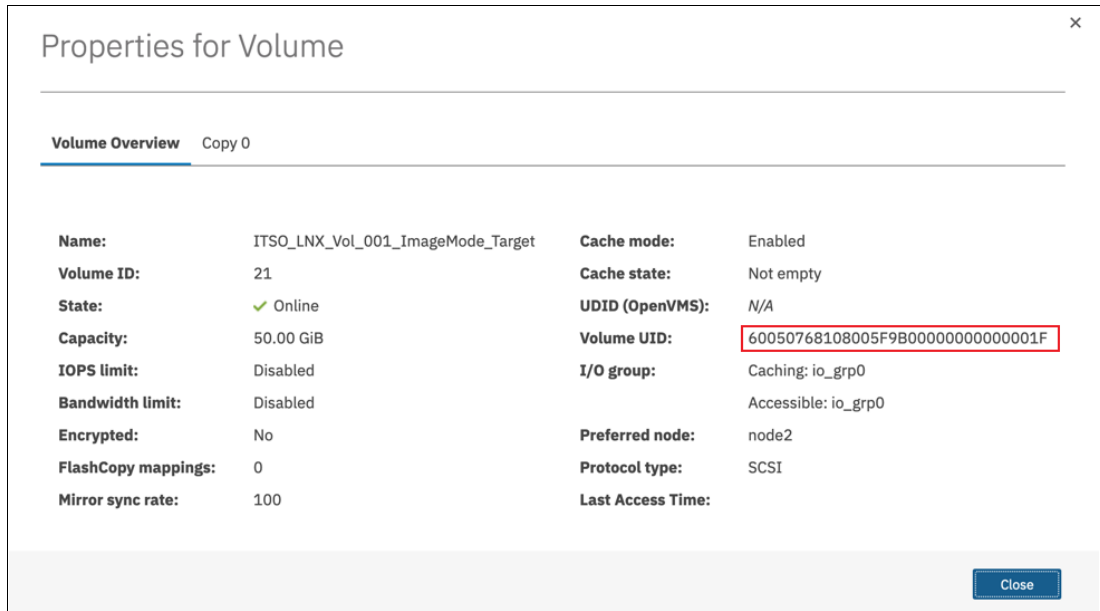


Figure 6-45 FlashSystem 9100 volume Properties page

Cross-check the **Volume UID** as shown in the FlashSystem 9100 to the MDisk **UID** as shown in the SAN Volume Controller GUI (from the MDisk **Properties** page - Figure 6-46 on page 87 and Figure 6-47 on page 88).

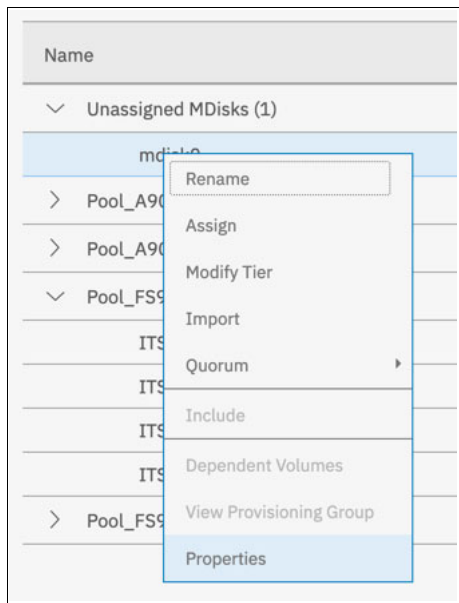


Figure 6-46 SAN Volume Controller MDisk right-click menu

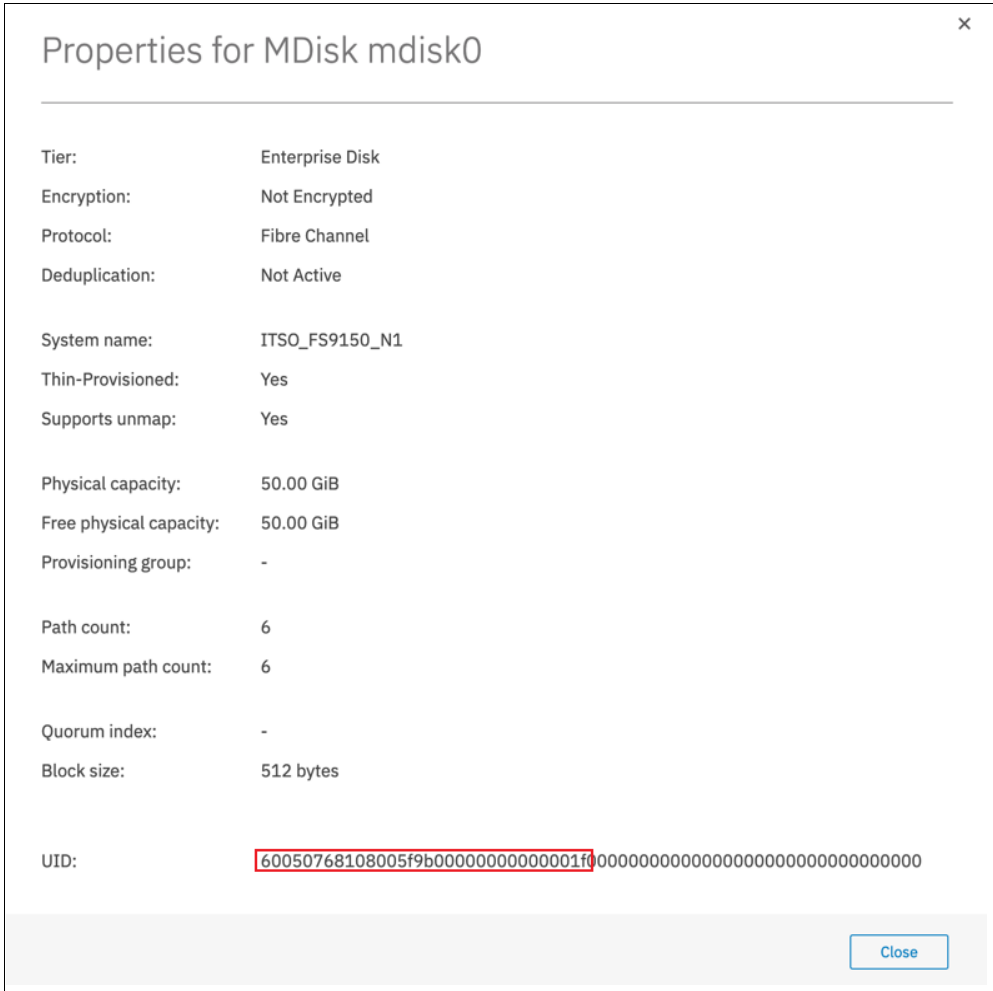


Figure 6-47 SAN Volume Controller MDisk Properties page

For convenience, we are going to rename mdisk0 to be able to quickly identify it in the following steps. From the volume right-click menu, select **Rename** (Figure 6-48).

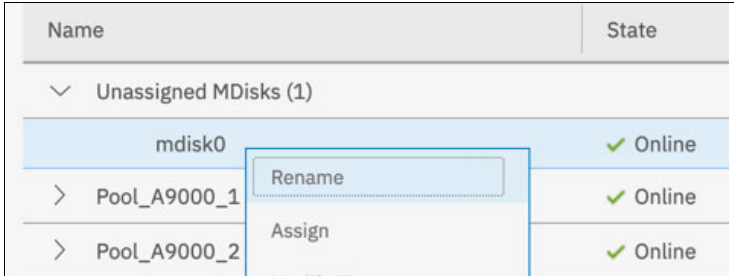


Figure 6-48 MDisk right-click menu

Then enter a new name in the Rename page as shown in Figure 6-49.

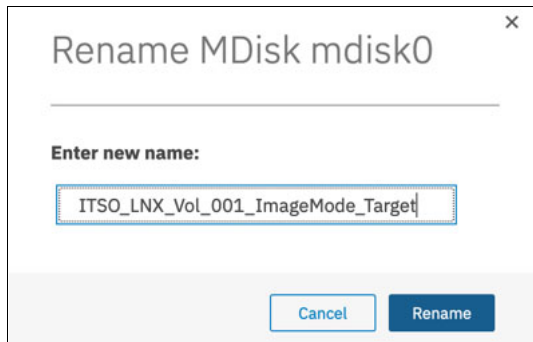


Figure 6-49 MDisk Rename page

Creating a temporary migration pool

The procedure requires to select one pool where the target MDisk will be placed while the extents migration is in progress. You can use any already defined pool on your SAN Volume Controller, or optionally decide to create a new, temporary one.

To create a new temporary pool, simply select **Create** in the **Pools** view, as shown in Figure 6-50.




Figure 6-50 Create Pool action button

Use the **Create Pool** page to give the pool a name and select the same extent size as in source pool.

Note: at this point, it is not recommended to enable the Data Reduction feature for the Migration Pool. As you might want to migrate to an already compressed/deduplicated back-end storage (such as FlashSystem 9100), the use of Data Reduction Pools over Data Reduction Pools is strongly discouraged.

Create Pool

Name:

Extent size: ▼

Maximum addressable capacity: 4.00 PiB

Data reduction: Enable

Figure 6-51 Create Pool page

Exporting one volume copy to Image Mode

We are now ready to export volume data from the SAN Volume Controller to the FlashSystem 9100 using the Image Mode functionality. To start, right-click the volume and select **Export to Image Mode**, as shown in Figure 6-52.

Name	State	Synchronized
<div style="display: flex; align-items: center;"> ▼ ITSO_LNX_Vol_001 </div>	<div style="display: flex; align-items: center;"> ✓ Online </div>	
Copy 0*	<div style="display: flex; align-items: center;"> ✓ Online </div>	
Copy 1	<div style="display: flex; align-items: center;"> ✓ Online </div>	

- Rename...
- Map to Host or Host Cluster...
- Shrink...
- Expand...
- Modify Capacity Savings...
- Modify Mirror Sync Rate...
- Cache Mode...
- Modify Open VMS UDID...
- Remove Private Mappings...
- View Mapped Hosts...
- View Member MDisks...
- Modify I/O Group...
- Cloud Volumes ▶
- Space Savings ▶
- Migrate to Another Pool...
- Export to Image Mode

Figure 6-52 Export to Image Mode

The **Export to Image Mode** page displays (Figure 6-53).

From the **Select Volume Copy** drop-down selector, choose which of the two volume copies to migrate. Extents from this copy will be moved from their original pool to the target MDisk, while the other copy will remain untouched.

Next, select the target MDisk you want to use for the migration. The page will only display unassigned MDisks with a capacity equal or greater to the capacity of the source volume you are trying to migrate.

Select the wanted one (ITSO_LNX_Vol_001_ImageMode_Target, in this case) and click **Next**. See Figure 6-53.

Export to Image Mode

You selected to export the volume **ITSO_LNX_Vol_001** to an image mode MDisk. This volume has more than one copy. Choose which copy to export and the MDisk to export to.

Select the Volume Copy

Copy 1 (Pool_A9000_2)

All System Storage | Filter

Name	Status	Capacity	Site
ITSO_LNX_Vol_001_ImageMode_Target	✓ Online	50.00 GiB	

Showing 1 MDisk / Selecting 1 MDisk

Cancel Back Next

Figure 6-53 Export to Image Mode page

To migrate the volume copy to Image Mode, SAN Volume Controller needs to place the target MDisk in a pool. Select the pool you want to use for this purpose - in our case, it is the temporary pool that we created earlier.

Finally, click **Finish** to start the migration, as shown in Figure 6-54.

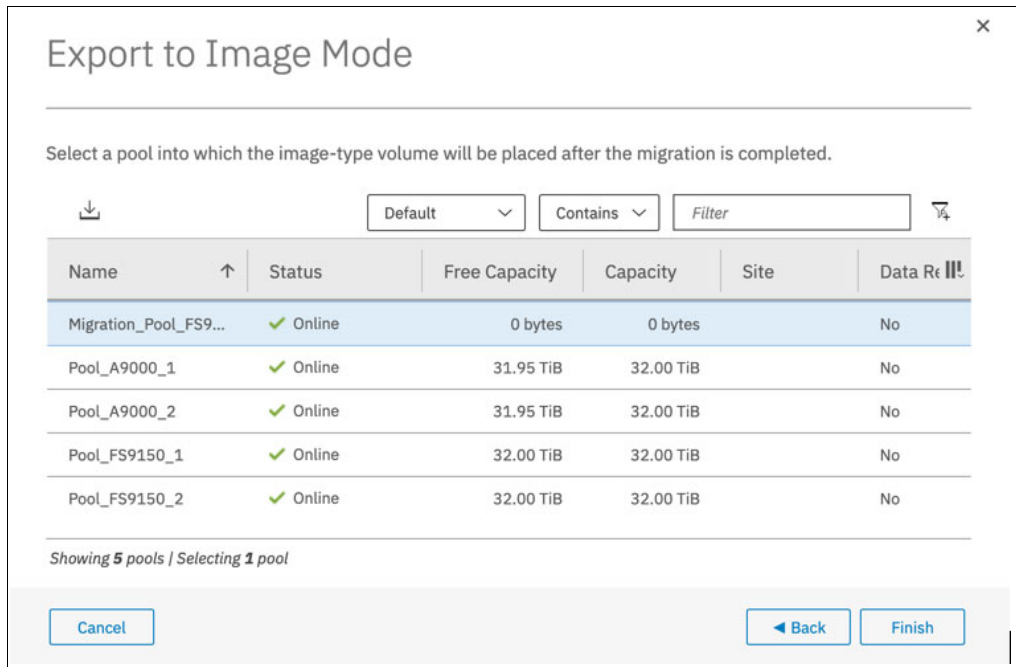


Figure 6-54 Selecting a Migration Pool

You can monitor the progress of the migration via the **Running Tasks** menu in the upper right corner of the IBM Spectrum Virtualize GUI - Figure 6-55.




Figure 6-55 Running Tasks quick view

If you click **View**, you are redirected to the **Monitoring** → **Background Tasks** page, where you can monitor the progress of the migration, as shown in Figure 6-56.




Figure 6-56 Migration task details

When the migration completes, the migrated copy volume will switch to **Virtualization Type: Image** (Figure 6-57).

Name	State	Synchronized	Virtualization Type
ITSO_LNX_Vol_001	✓ Online		Many
Copy 0*	✓ Online	Yes	Striped
Copy 1	✓ Online	Yes	Image

Figure 6-57 Volume copies types after the migration to Image Mode has completed

At this point, the host can continue using this volume and any subsequent change to the data will be kept in sync between the two copies by the Volume Mirroring feature of IBM Spectrum Virtualize.

Removing access to the migration source volume

At this point of the migration procedure, volume data is being stored both by the SAN Volume Controller (in one of the already established pools) and on the FlashSystem 9100 (via the Image Mode copy in our temporary Migration pool). However, host IOs are still being serviced by SAN Volume Controller. To be able to move completely to the new storage system, we need to divert host IOs directly to the target volume on FlashSystem 9100. The only way to do this is to unmap the SAN Volume Controller-provided volume from the host, and map the new one from the FlashSystem 9100. This means that we need to plan a certain amount of downtime for the application(s) using the volume, as this operation cannot be done concurrently. Therefore we need to make sure to plan an adequate cutover window to carry out the next steps.

First, make sure to quiesce all IOs on the volume. Stop any applications that are accessing the volumes, unmount file systems, unconfigure logical volumes/volume groups using the volume, remove devices, and so on. Linux commands for our example are shown in Example 6-2.

Example 6-2 Unmount file system and remove device on the Linux host

```
# Unmount /itso file system
[root@comp-1-31 /]# umount /itso
# Remove multipath device
[root@comp-1-31 /]# multipath -f mpathc
```

After you are certain that the volume does no longer receive any IO from the host, you can proceed with the unmapping of the volume on the SAN Volume Controller. From the volume right-click menu select **Remove Private Mappings...** (Figure 6-58).

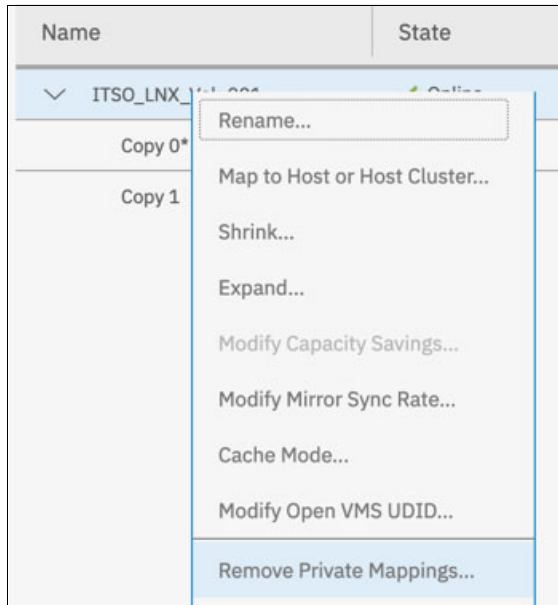


Figure 6-58 Remove Private Mappings

In the **Remove Private Mappings** page enter the number of mappings (1 in this case) and confirm by clicking **Remove**, as seen in Figure 6-59.

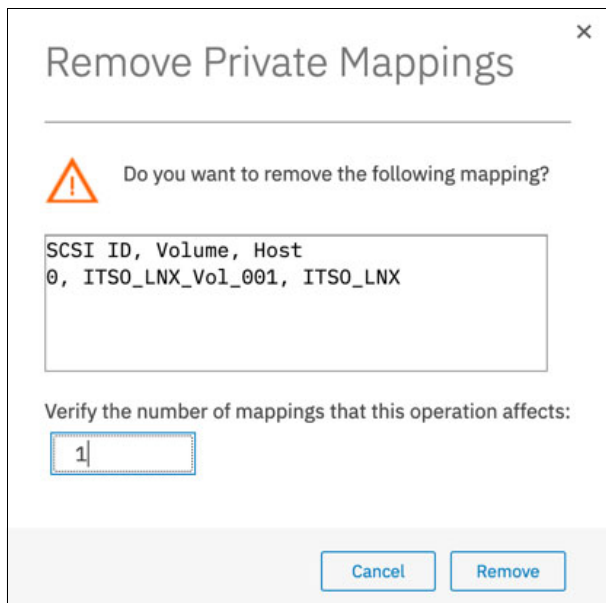


Figure 6-59 Remove Private Mapping page

To release the Image Mode MDisk, it is now necessary to delete the corresponding volume copy. This action will not actually delete any data on the MDisk, but only its association to the SAN Volume Controller volume.

To perform this step, right-click the Image Mode volume copy and select **Delete**, as show in Figure 6-60.

Name	State	Synchronized	Virtualization Type
ITSO_LNX_Vol_001	Online		Many
Copy 0*	Online	Yes	Striped
Copy 1	Online	Yes	Image

Shrink...
 Expand...
 Create Volume From This Copy
 Split into New Volume
 Make Primary
 Validate Volume Copies
 Space Savings ▶
 Delete

Figure 6-60 Volume Copy right-click menu

Confirm the volume copy removal (Figure 6-61).

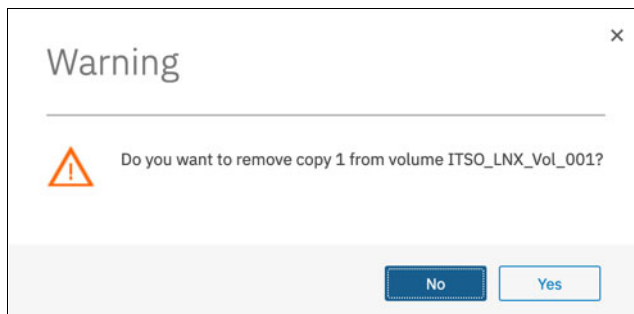


Figure 6-61 Confirm the copy removal

After the removal process, verify that the target MDisk has returned to the **Unassigned MDisks** group, as it is not part of the Migration Pool anymore.

See Figure 6-62.

Name	State	Effective Capacity
Unassigned MDisks (1)		
ITSO_LNX_Vol_001_ImageMode_Target	Online	50.00 GiB
Migration_Pool_FS9150	No Storage	

Figure 6-62 Target MDisk becomes is unassigned after Image Mode copy deletion

At this point, you can unmap the target MDisk from the SAN Volume Controller.

In the FlashSystem 9100 **Host Cluster** view, right-click the SAN Volume Controller Host Cluster object and select **Modify Shared Volume Mappings**, as shown in Figure 6-63.

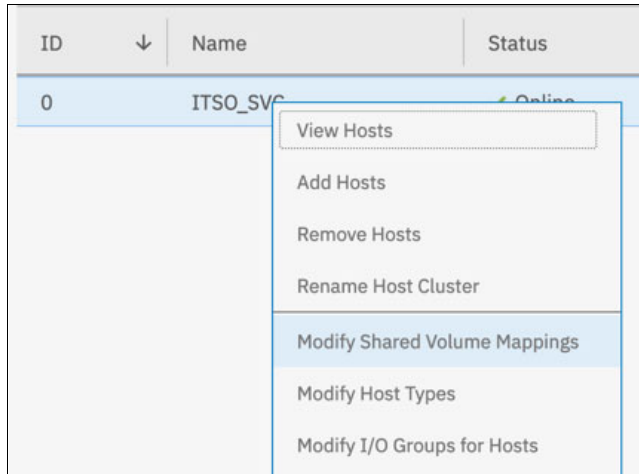


Figure 6-63 Host Cluster right-click menu

Select the Image Mode target volume and proceed to the next steps by clicking **Remove Volume Mappings** (Figure 6-64).

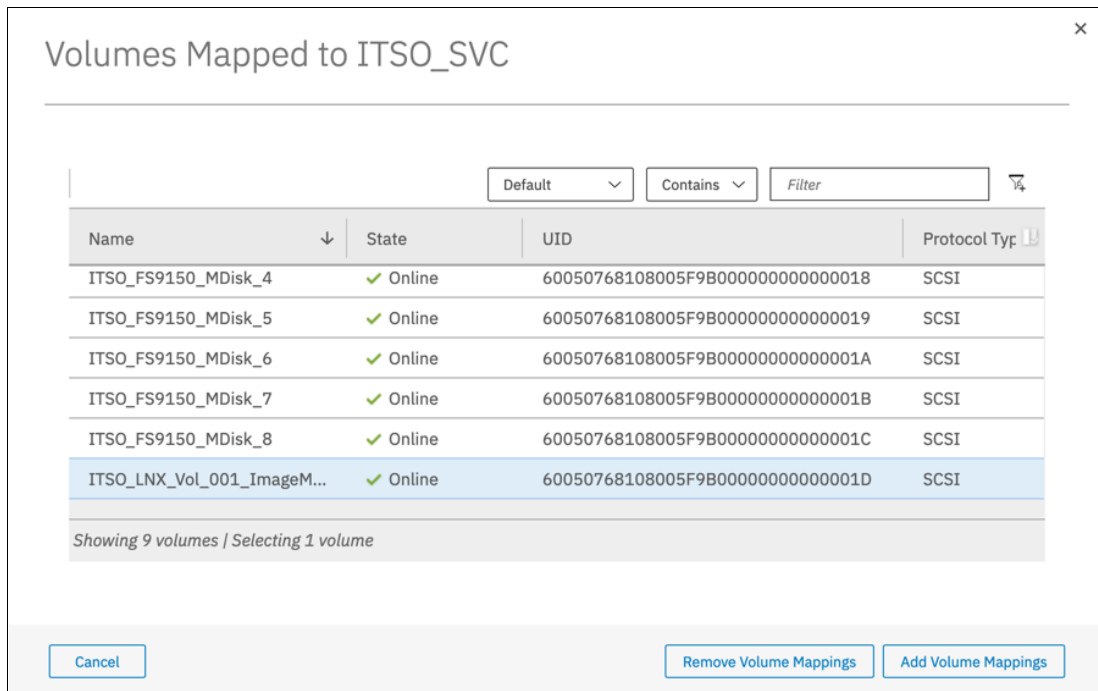


Figure 6-64 Volume Mappings edit page

In the **Summary** page, do not select any host to avoid retaining the mapping as a private mapping, and simply click **Next**, as shown in Figure 6-65.

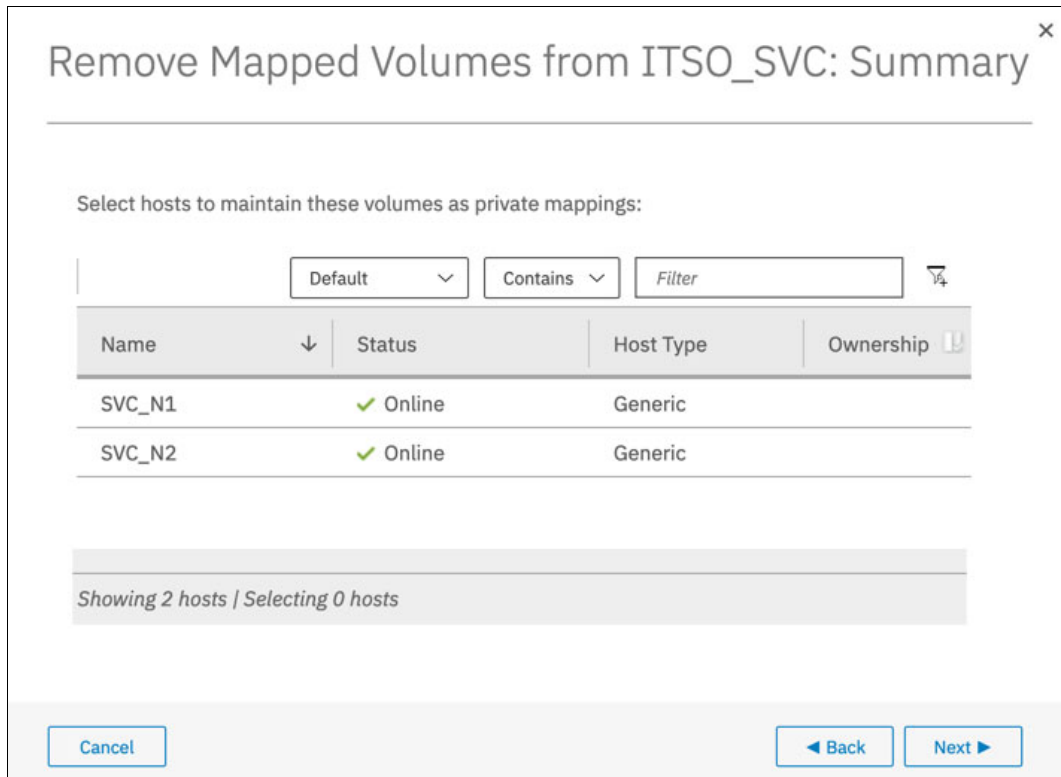


Figure 6-65 Volume Mapping removal summary page

Finally, confirm the mapping removal. See Figure 6-66.

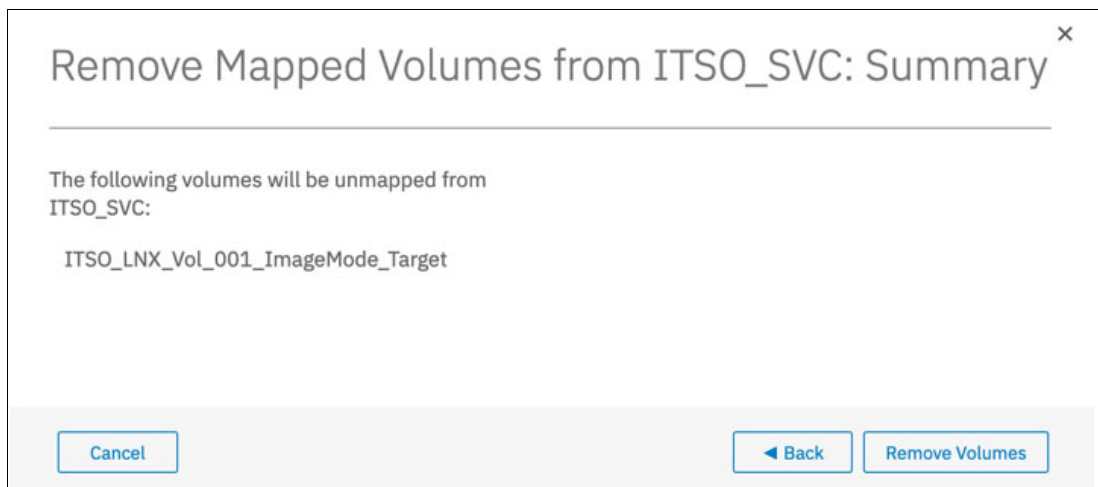


Figure 6-66 Volume Mapping removal confirmation page

Presenting the migration target volume to the host

You can now assign the migrated volume on the FlashSystem 9100 to the host, to resume application data access on the new storage system.

If you have not already done so, define a new **Host** object on the FlashSystem 9100. From the **Hosts** view, select **Add Host** (Figure 6-67).




Figure 6-67 Add Host action button

In the **Add Host** page, specify host details to allow the connection (Figure 6-68).

Verify that the connectivity between the host and the storage system is working as expected; if the host is reported as **offline** or **degraded** in the FlashSystem 9100 GUI, inspect SAN zoning, cabling or iSCSI connectivity to resolve the issue.

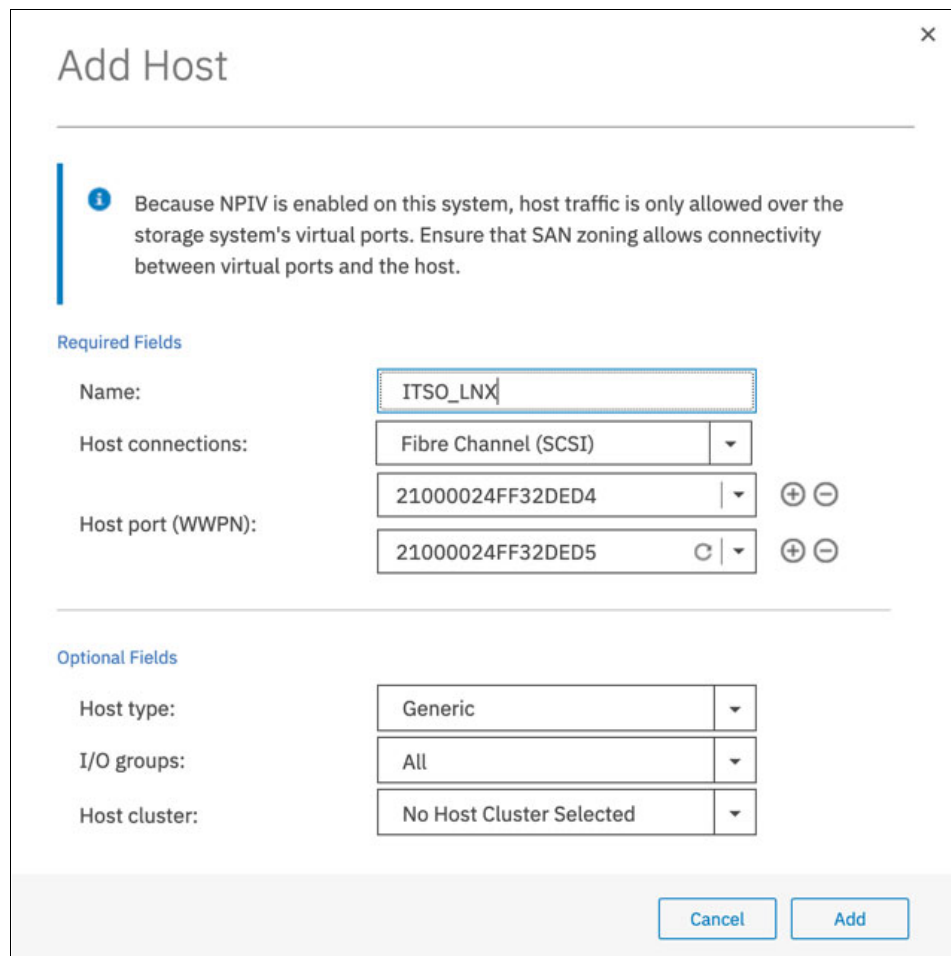
A screenshot of the 'Add Host' configuration page. The page title is 'Add Host'. An information icon and text state: 'Because NPIV is enabled on this system, host traffic is only allowed over the storage system's virtual ports. Ensure that SAN zoning allows connectivity between virtual ports and the host.' Below this, there are two sections: 'Required Fields' and 'Optional Fields'.
Required Fields:
- Name: ITSO_LNX
- Host connections: Fibre Channel (SCSI)
- Host port (WWPN): 21000024FF32DED4 (with a plus/minus icon)
- Host port (WWPN): 21000024FF32DED5 (with a refresh and plus/minus icon)
Optional Fields:
- Host type: Generic
- I/O groups: All
- Host cluster: No Host Cluster Selected
At the bottom right, there are 'Cancel' and 'Add' buttons.

Figure 6-68 Add Host page

Proceed with mapping the volume to the newly created host, via right-click menu (**Map to Host or Host Cluster...** option, Figure 6-69).




Figure 6-69 Volume right-click menu

In the **Create Mapping** page, select the host and click **Next**, as shown in Figure 6-70.

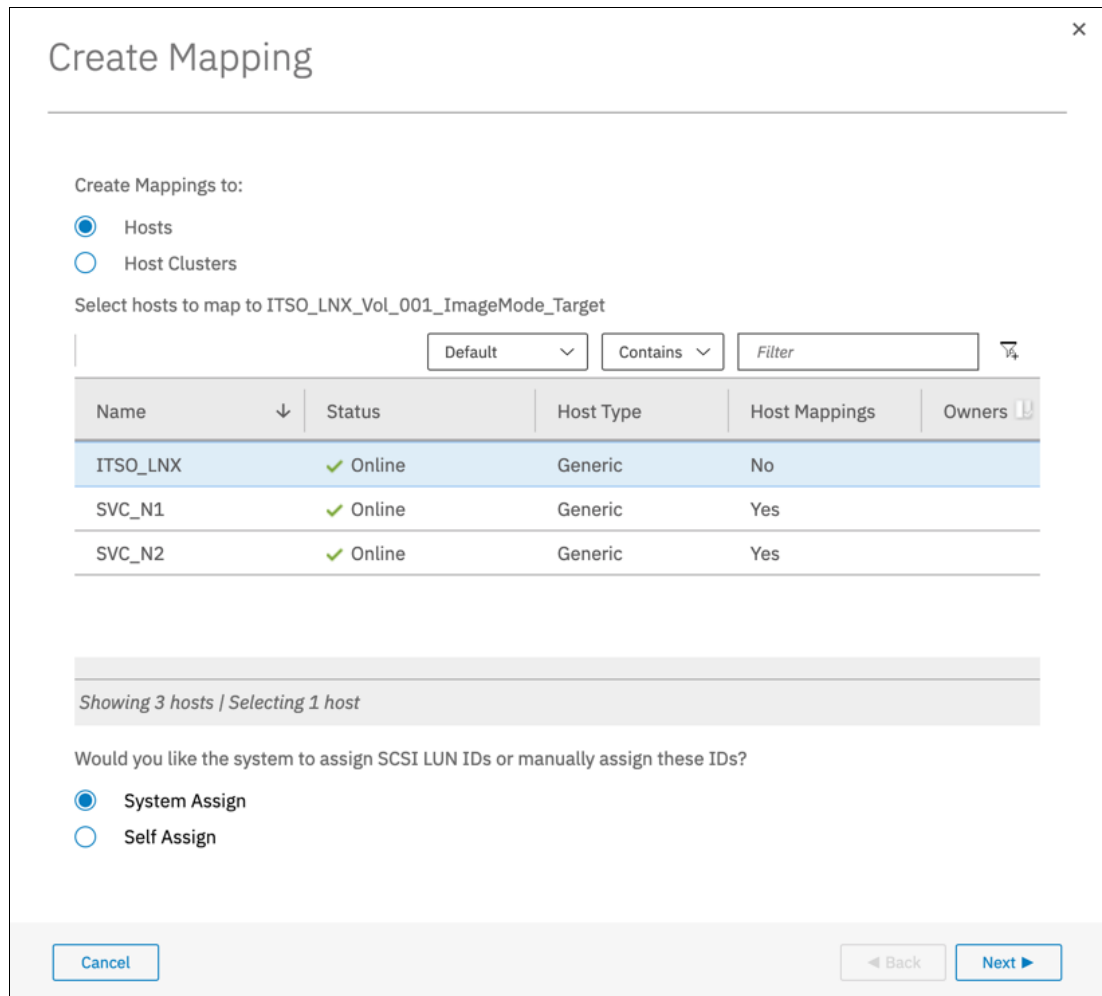


Figure 6-70 Create Mapping page

Confirm the mapping (Figure 6-71 on page 100).

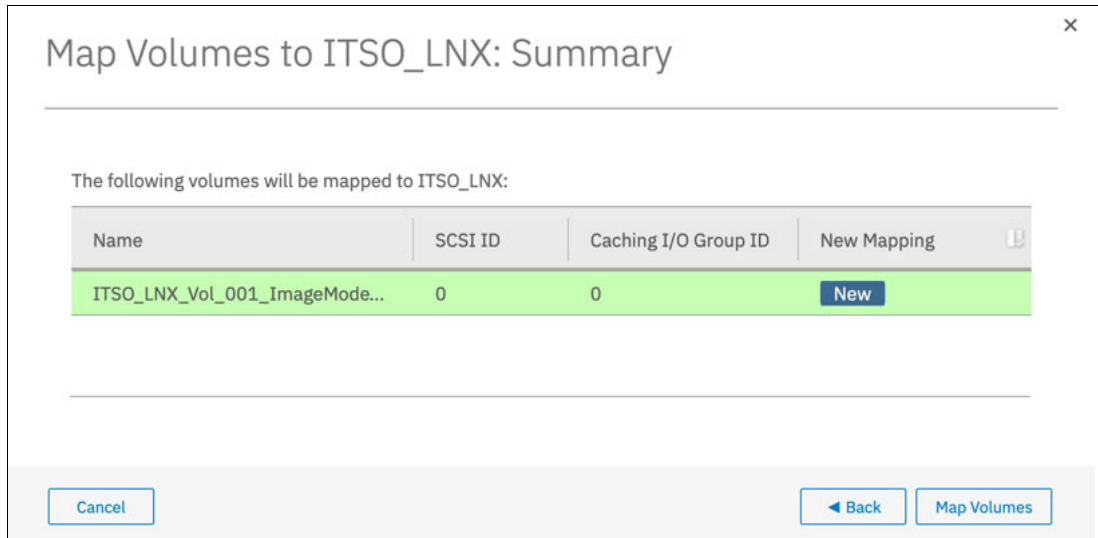


Figure 6-71 Volume mapping summary

Your migrated volume should now be visible on your host after a rescan. Reconfigure any logical device mapping/volume grouping, mount the file system(s) and verify that everything is correctly configured to restart your application. See Example 6-3 for our Linux host.

Example 6-3 Rescanning and reconfiguring device from FlashSystem 9100 on the Linux host


```
# Rescan devices
[root@comp-1-31 /]# /usr/bin/rescan-scsi-bus.sh
Scanning SCSI subsystem for new devices
...
# Show newly discovered multipath device
[root@comp-1-31 /]# multipath -ll
mpathd (360050768108005f9b0000000000001f) dm-3 IBM ,2145
size=50G features='1 queue_if_no_path' hwhandler='0' wp=rw
|-+- policy='service-time 0' prio=50 status=active
|  |- 5:0:5:1 sdl 8:176 active ready running
|  `-- 6:0:5:1 sdn 8:208 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
   |- 5:0:4:1 sdk 8:160 active ready running
   `-- 6:0:4:1 sdm 8:192 active ready running

# Mount the file system
[root@comp-1-31 /]# mount /dev/mapper/mpathd1 /itso

# Verify file system
[root@comp-1-31 /]# df -h
Filesystem              Size  Used Avail Use% Mounted on
...
/dev/mapper/mpathd1     50G   53M   47G   1% /itso

# Verify that all data has been migrated correctly
[root@comp-1-31 /]# ls /itso/
bert.txt  chris.txt  duncan.txt  francesco.txt  harry.txt  lost+found  markus.txt
matt.txt  roger.txt  steve.txt
```

After this step, resumed IO on the volume will be serviced by the FlashSystem 9100. You can retain the old volume copy on the SAN Volume Controller for backup purposes, or proceed with its deletion.



Data Migration with IBM Spectrum Virtualize Remote Copy Services

This chapter explains how to use the Remote Copy feature to migrate data from a IBM Spectrum Virtualize system to another IBM Spectrum Virtualize system.

The example used for this chapter shows a migration between a SAN Volume Controller volume and a FlashSystem 9100 volume, attached to a Linux host.

7.1 Data migration using IBM Spectrum Virtualize Global Mirror

The IBM Spectrum Virtualize Remote Copy functions can be used to migrate data while the application server is still active on the source storage system.

IBM Spectrum Virtualize allows three types of Remote Copy relationships:

- ▶ Metro Mirror
- ▶ Global Mirror
- ▶ Global Mirror with Change Volumes

You can use any of the three types of Remote Copy functions to migrate your data from one source IBM Spectrum Virtualize system, called the *Master Spectrum Virtualize system*, to one target IBM Spectrum Virtualize system, called the *Auxiliary Spectrum Virtualize system*.

See the following IBM Support page for Inter-System Metro Mirror and Global Mirror Compatibility Cross Reference:

<https://www.ibm.com/support/pages/node/706911>

This section outlines the steps to set up a Remote Copy partnership between two IBM Spectrum Virtualize systems, and migrate data from one to the other using Global Mirror relationships and consistency groups.

Our example is based on a Global Mirror relationship between a SAN Volume Controller volume and FlashSystem 9100 volume to migrate data between two physically separate locations. You can also use Global Mirror as a data migration solution within a single IBM Spectrum Virtualize system.

Figure 7-1 shows the test environment.

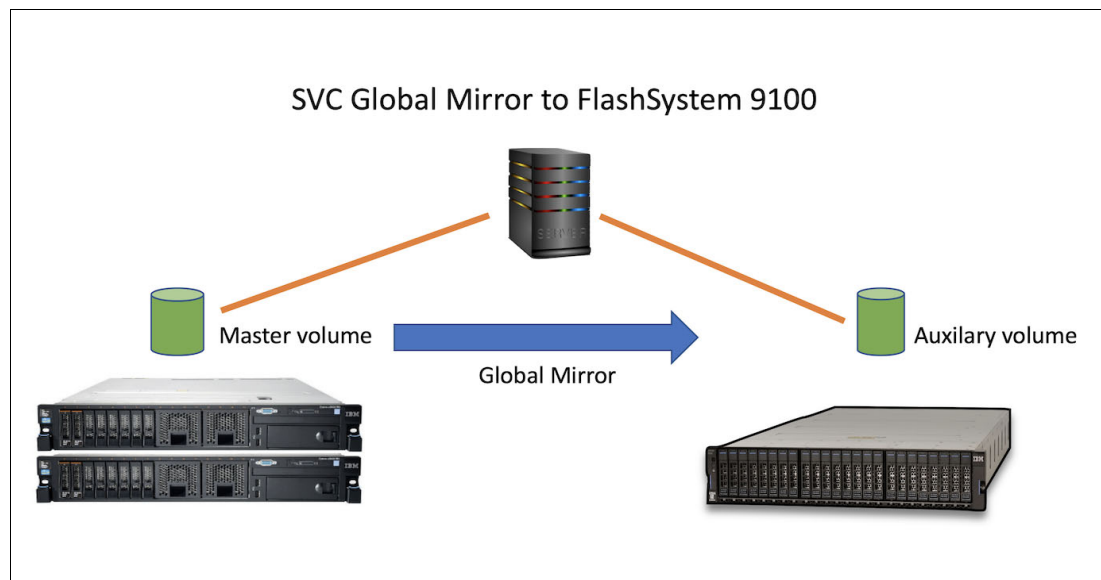


Figure 7-1 SAN Volume Controller Global Mirror to FlashSystem 9100 Lab Setup

Our starting point is the `ITS0_LNX_Vo1_001` volume defined on the `sbssvc` SAN Volume Controller system.

The volume is currently mapped to a Linux host; on the host, the volume is configured as a multipath device; an ext4 file system is defined on a partition of the device, and some files have been written to it.

Example 7-1 shows the volume data from the host point of view.

Example 7-1 Device and file system configuration example

```
# Device Mapper Multipath entry for the ITSO_LNX_Vol_001 volume
[root@comp-1-31 ~]# multipath -ll
mpathc (3600507680c80011a800000000000000a) dm-5 IBM      ,2145
size=50G features='1 queue_if_no_path' hwhandler='0' wp=rw
|-+- policy='service-time 0' prio=50 status=active
|  |- 5:0:5:0 sds 65:32 active ready running
|  |-- 6:0:3:0 sdu 65:64 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
|  |-- 5:0:4:0 sdt 65:48 active ready running
|  |-- 6:0:7:0 sdv 65:80 active ready running

# File system details
[root@comp-1-31 ~]# df -h
Filesystem                Size  Used Avail Use% Mounted on
...
/dev/mapper/mpathc1        50G   53M   47G   1% /itso

# File system contents
[root@comp-1-31 ~]# ls /itso
bert.txt  chris.txt  duncan.txt  francesco.txt  harry.txt  lost+found  markus.txt
matt.txt  roger.txt  steve.txt
```

7.1.1 Set up a IBM Spectrum Virtualize Remote Copy partnership

Create a Remote Copy partnership between the source and target systems of the migration, using the following steps:

1. In the target FlashSystem 9100 GUI, from the **Copy Services** → **Partnerships**, select **Create Partnership** (Figure 7-2).




Figure 7-2 Create Partnership button in the FlashSystem 9100 GUI

2. In the **Create Partnership** page, select the type of connectivity towards your Master system (Fibre Channel or IP), select the remote cluster in the drop-down menu (sbssvc in this example), and set appropriate **Link Bandwidth** and **Background Copy Rate** - Figure 7-3. Then click OK to create the partnership.

Important: Set the bandwidth used by the background copy process to less than or equal to the bandwidth that can be sustained by the communication link between the systems. The link must be able to sustain any host requests and the rate of background copy.

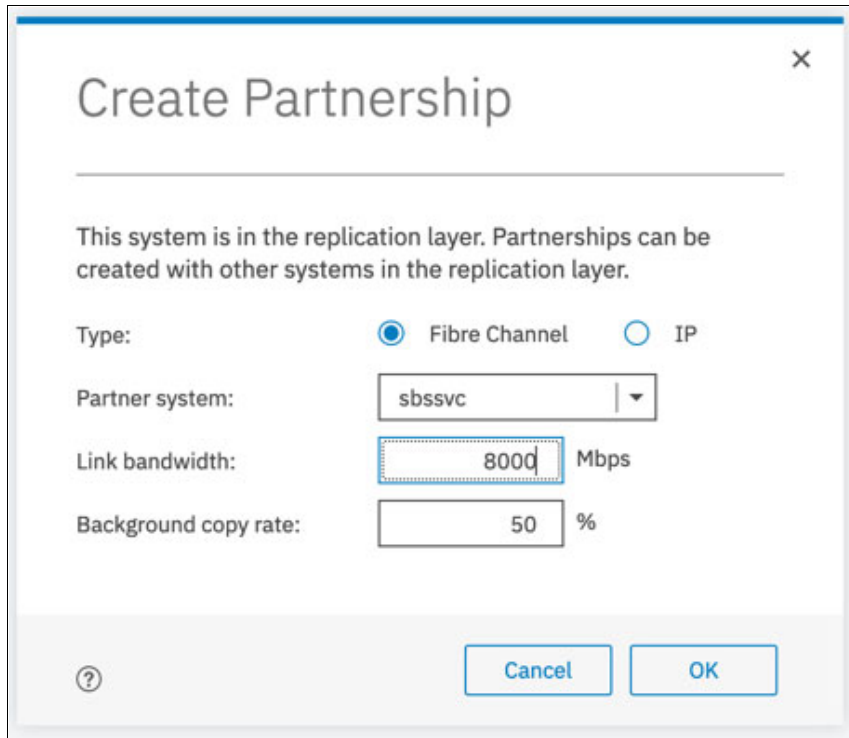


Figure 7-3 Create Partnership page

3. The **Partnerships** view now shows the newly created partnership. The state of the partnership is **Partially Configured** because the partnership must be set up in the opposite direction as well (Figure 7-4).

Name	Location	State
ITS09100	Local	-
sbssvc	Remote	▲ Partially Configured: Local

Figure 7-4 Newly created partnership in the FlashSystem 9100 GUI

4. Repeat the process in the source SAN Volume Controller GUI: from the **Copy Services** → **Partnerships** view, select **Create Partnership** (Figure 7-5).




Figure 7-5 Create Partnership button in the SAN Volume Controller GUI

- In the **Create Partnership** page, again identify your target cluster and input the **Link Bandwidth** and **Background Copy Rate** values - Figure 7-6.

Figure 7-6 Create Partnership page

- The partnership state changes to **Fully Configured** on both the SAN Volume Controller (Figure 7-7) and the FlashSystem 9100 (Figure 7-8).

Name	↑	Location	State
sbssvc		Local	-
ITSO9100		Remote	✓ Fully Configured

Figure 7-7 Partnership state on the SAN Volume Controller system

Name	↑	Location	State
ITSO9100		Local	-
sbssvc		Remote	✓ Fully Configured

Figure 7-8 Partnership state on the FlashSystem 9100 system

7.1.2 Creating consistency groups and relationships

After you established the Remote Copy partnership, you are ready to create the Remote Copy consistency groups and relationships.

Identify if volumes can be replicated independently. Some applications, such as databases, use multiple volumes and require the order of writes to these volumes to be preserved in the remote site. If an application requires write order to be preserved for the set of volumes that it uses, create a consistency group for these volumes

Creating Remote Copy consistency groups

To create a Remote Copy consistency group using the IBM Spectrum Virtualize GUI, perform the following steps:

1. In the Master system (SAN Volume Controller in this example) **Copy Services** → **Remote Copy** view, select **Create Consistency Group** (Figure 7-9).




Figure 7-9 Create Consistency Group button

2. Give the Consistency Group a name (Figure 7-10).

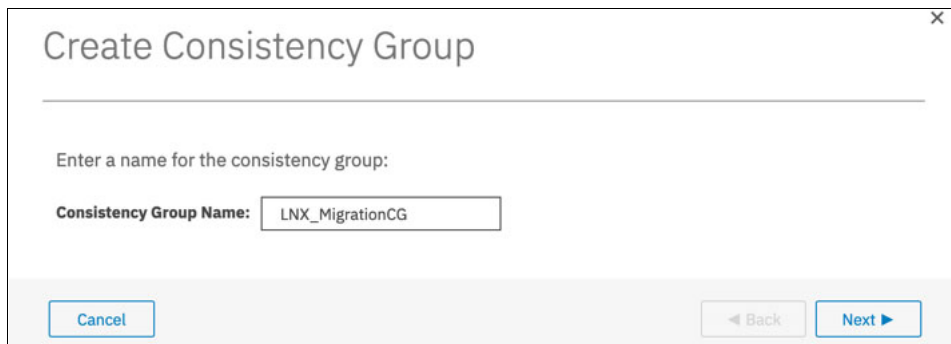


Figure 7-10 Create Consistency Group page

3. Select where the auxiliary volumes are located - in this case we are targeting the FlashSystem 9100, so we choose **On Another System** and select the ITS0FS9100 system in the drop-down menu, as shown in Figure 7-11.

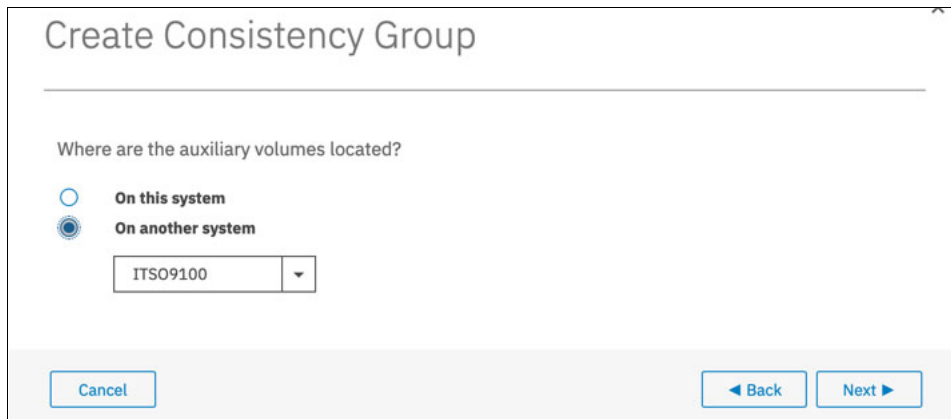


Figure 7-11 Selecting the auxiliary system for the Consistency Group

- Finally, specify if you want to create an empty consistency group, or you want to insert some Remote Copy relationship into it already. We do not have any relationship defined at the moment, so we select **No, create an empty consistency group** - Figure 7-12. Click **Finish** to create the Consistency Group.

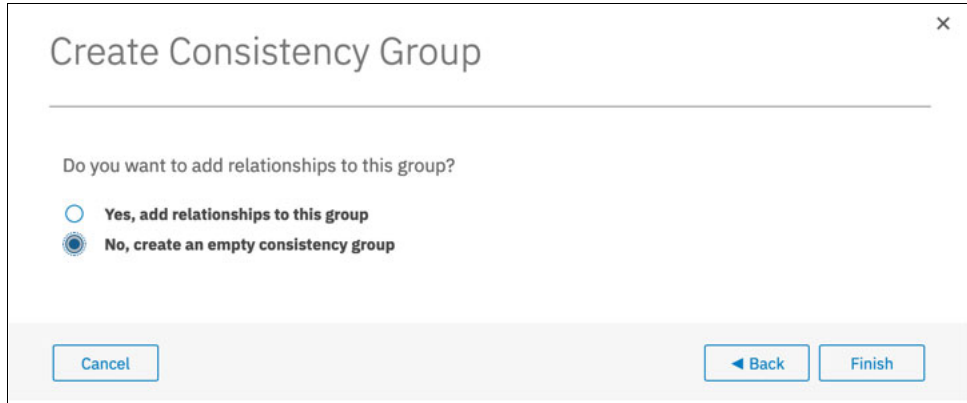


Figure 7-12 Create empty Consistency Group

- The newly created consistency group appears in the **Copy Services Remote** → **Copy** view (Figure 7-13).

Name	State	Master Volume	Auxiliary Volume
Not in a Group			
LNX_MigrationCG	State: Empty	Master System: sbssvc	Auxiliary System: ITS09100

Figure 7-13 Newly created consistency group is displayed

Creating Remote Copy relationships

To create the Global Mirror Remote Copy relationships using the GUI, perform the following steps:

- In the Master system **Copy Services** → **Remote Copy** view, right-click the newly created consistency group and select **Create Relationship** (Figure 7-14).

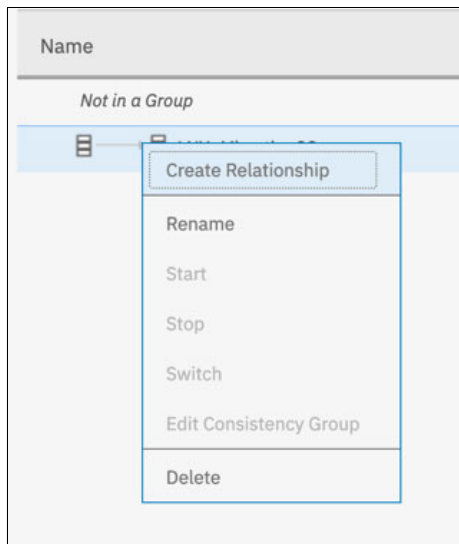


Figure 7-14 Create Relationship menu

2. On the **Create Relationship** page, select **Global Mirror** and click **Next** (Figure 7-15);

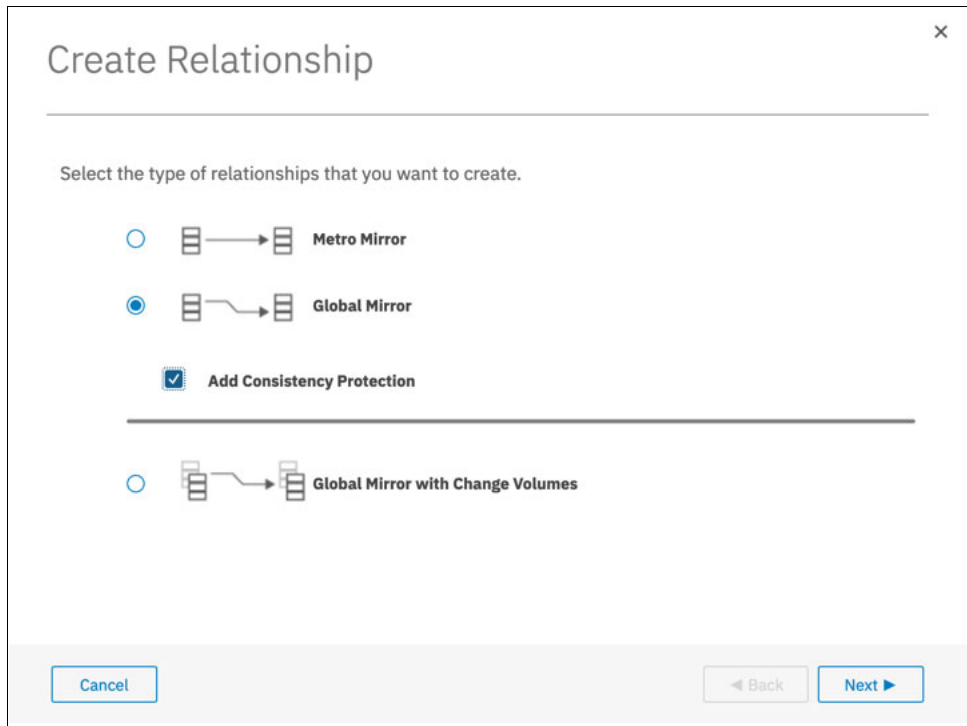


Figure 7-15 Remote Copy relationship type selection

3. Select the Master and Auxiliary volumes for the relationship. In this example, our Master volume is `ITSO_LNX_Vol_001`, defined on the system we want to migrate from (SAN Volume Controller). Our auxiliary volume is `ITSO_LNX_Vol_001_GM_Target`, defined on the system we want to migrate to (FlashSystem 9100), as shown in Figure 7-16 on page 110. Click **Add** to add the volume pair.

Tip: If after selecting your Master volume in the first drop-down you do not see the wanted target volume in the second, ensure that the two volumes have the exact same size, down to the number of bytes.

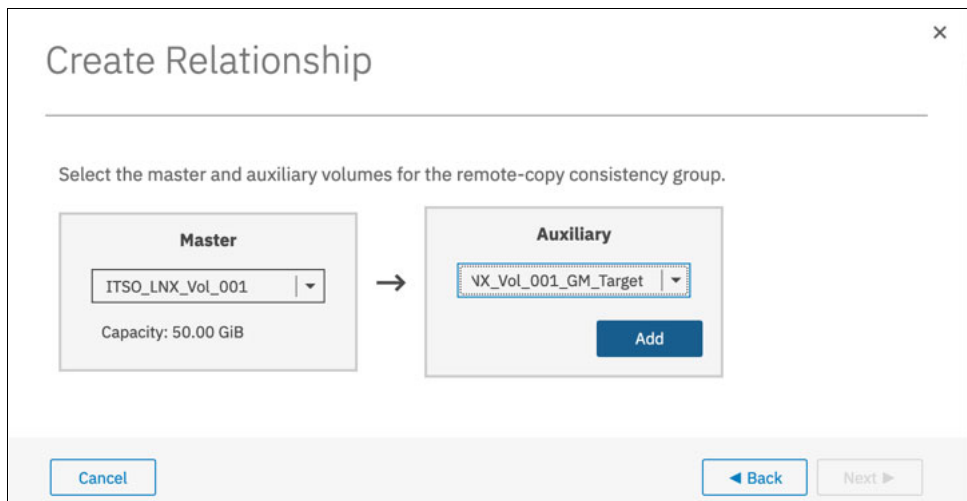


Figure 7-16 Relationship volume pair selector page

4. In our example we want to implement a regular Global Mirror relationship, so we select **No, do not add a master change volume** and click **Finish** (Figure 7-17).

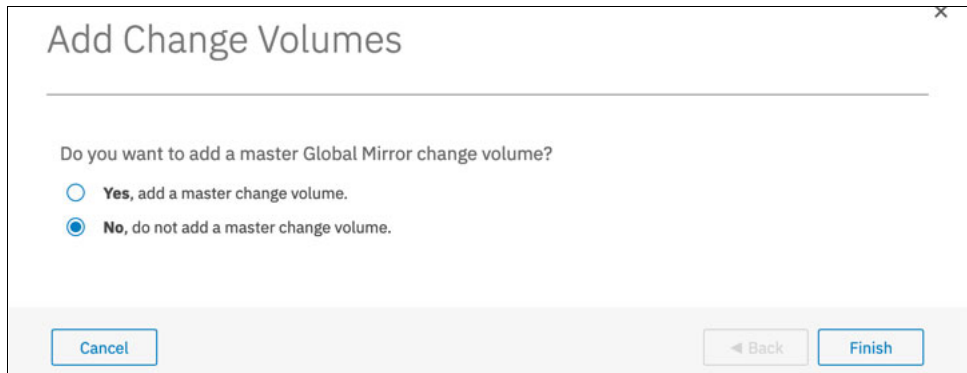


Figure 7-17 Add Change Volumes page

5. The volume pair for the relationship we want to create is added, as shown in Figure 7-18. Click **Next** to proceed to the next step.

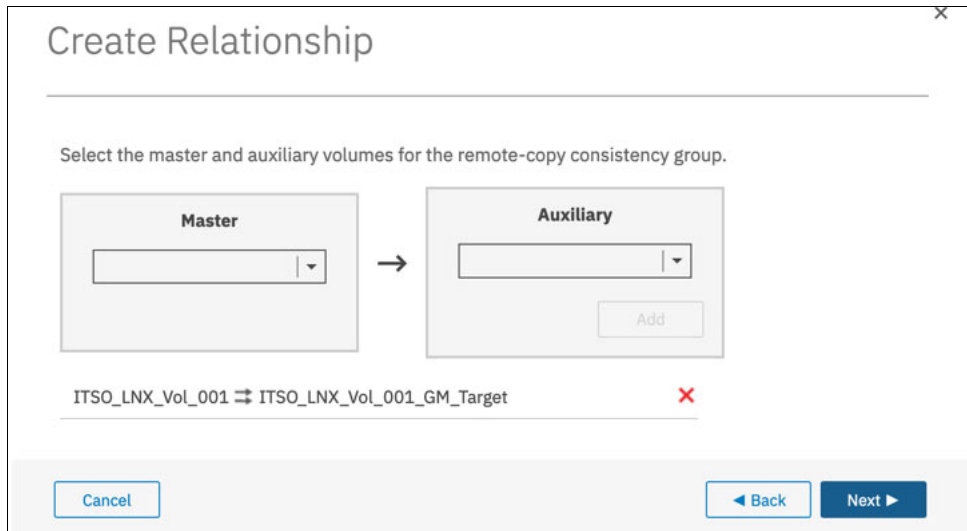


Figure 7-18 Relationship volume pair added

6. Select **No, the volumes are not synchronized** in order to make sure a full copy of the data from the Master to the Auxiliary system is performed upon launch of the relationship (Figure 7-19). Then click **Finish**.

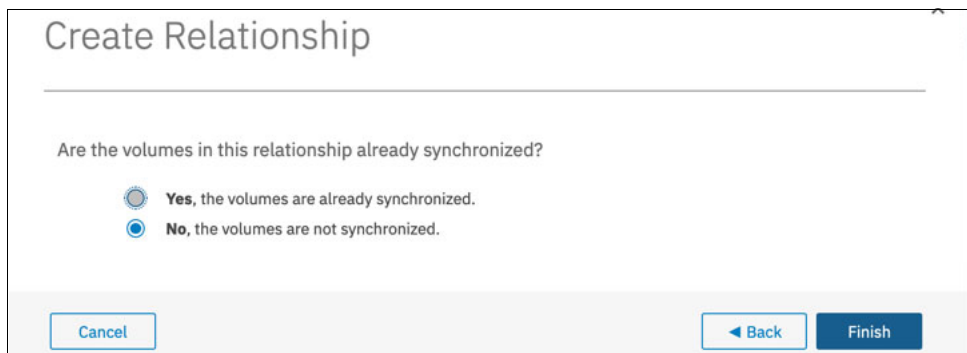


Figure 7-19 Create Relationship confirmation page

- The newly created relationship (rcre10) is displayed under its consistency group in **Inconsistent Stopped** state. The consistency group inherits the same state.

Name	State	Master Volume	Auxiliary Volume
<i>Not in a Group</i>			
LNX_MigrationCG	State: Inconsistent Stopped	Master System: sbssvc	Auxiliary System: ITS09100
rcre10	Inconsistent Stopped	ITSO_LNX_Vol_001	ITSO_LNX_Vol_001_GM_Target

Figure 7-20 Newly created relationship

- At this point, data is not being copied to the Auxiliary system yet. When you are ready to start the initial synchronization of the volumes, follow the instructions in the next section.

7.1.3 Starting and monitoring IBM Spectrum Virtualize Mirror Copy process

To copy the data from one system to the other using the IBM Spectrum Virtualize GUI, perform the following steps:

- Right-click the consistency group, as shown in Figure 7-21 and select **Start**.

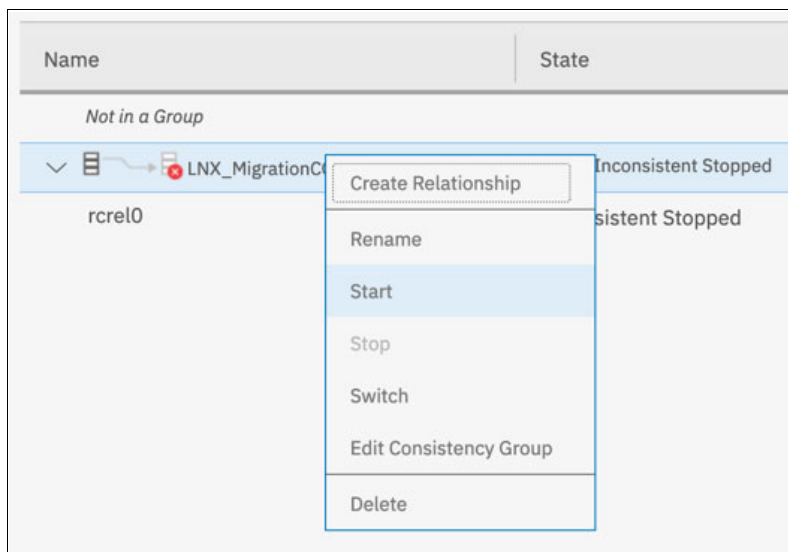


Figure 7-21 Start relationship menu

- After the command completes, the consistency group enters the **Inconsistent Copying** state (Figure 7-22). This means that data is being copied from the Master volume to the Auxiliary volume, but the initial synchronization has not been completed yet.

Name	State	Master Volume	Auxiliary Volume
<i>Not in a Group</i>			
LNX_MigrationCG	State: Inconsistent Copying	Master System: sbssvc	Auxiliary System: ITS09100
rcre10	Inconsistent Copying	ITSO_LNX_Vol_001	ITSO_LNX_Vol_001_GM_Target

Figure 7-22 Relationship in Inconsistent Copying

- After the consistency group is started, one or more background activities displays in **Running Tasks** quick menu (Figure 7-23).

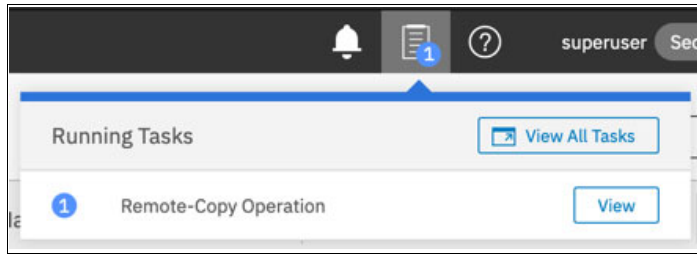


Figure 7-23 Remote Copy operation in the Running Task menu

- Clicking **View** takes you to the **Monitoring** → **Background Tasks** view, where the synchronization progress is displayed, as shown in Figure 7-24.




Figure 7-24 Remote Copy background tasks details

- When the volumes are synchronized, the consistency group state changes to **Consistent Synchronized** (Figure 7-25). While the group is in this state, all subsequent changes issued by the host(s) to the volumes in the consistency group are replicated in the same order at the Auxiliary site, maintaining consistency between the Master and the Auxiliary volumes.

Name	State	Master Volume	Auxiliary Volume
Not in a Group			
LNX_MigrationCG	State: Consistent Synchronized	Master System: sbssvc	Auxiliary System: ITSO9100
rcrel0	Consistent Synchronized	ITSO_LNX_Vol_001	ITSO_LNX_Vol_001_GM_Target

Figure 7-25 Consistency Group in Consistent Synchronized state

7.1.4 Stopping the Remote Copy and switching to the new storage system

At any point in time after the initial synchronization of data between the migration source system and the migration target system, you can switch your application workload to the new storage system.

This action requires the application to be shut down because you must switch application access from the existing to the new storage.

Before starting these steps, shut down the application and prepare the system to move to the new storage. In general, these steps include:

- ▶ Quiesce all I/O on the affected volumes.
- ▶ Unmount all file systems existing on the affected volumes.
- ▶ Remove the corresponding devices.
- ▶ Unmap the source volume(s) from the host(s).
- ▶ Verify that all changes have been propagated to the auxiliary site.
- ▶ Stop the Global Mirror relationship.
- ▶ Remove the Global Mirror relationships.

- ▶ Map the target volume(s) on the auxiliary system to the host(s).
- ▶ Reconfigure devices, file systems etc. on the host(s).
- ▶ Verify integrity of data on the new volume(s).
- ▶ Resume IO.

Unmapping the volume from the old storage system

After you have unconfigured and removed the volume at the host level, you can safely unmap it from the old storage system. In order to do that, perform the following steps:

1. In the **Volumes** view, right-click the volume and select **Remove Private Mappings** (Figure 7-26).

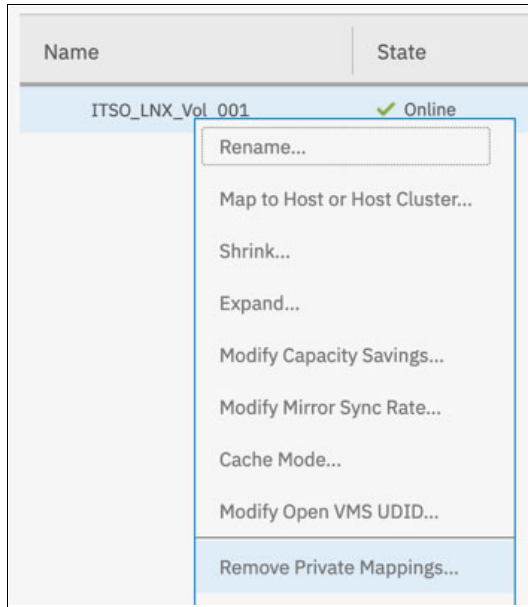


Figure 7-26 Remove Private mappings menu

2. In the **Remove Private Mappings** page, confirm your action - Figure 7-27.

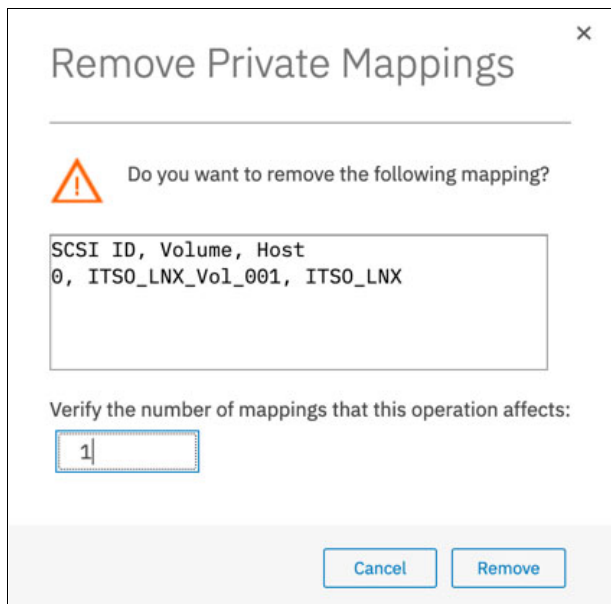


Figure 7-27 Confirm mapping removal

Stopping the Remote Copy relationships

To stop the Global Mirror relationship using the IBM Spectrum Virtualize GUI, perform the following steps:

1. In the Master system, go the **Copy Services** → **Remote Copy** view and right-click the consistency group containing the relationships for the volumes you want to switch access to, then select **Stop**, as shown in Figure 7-28.

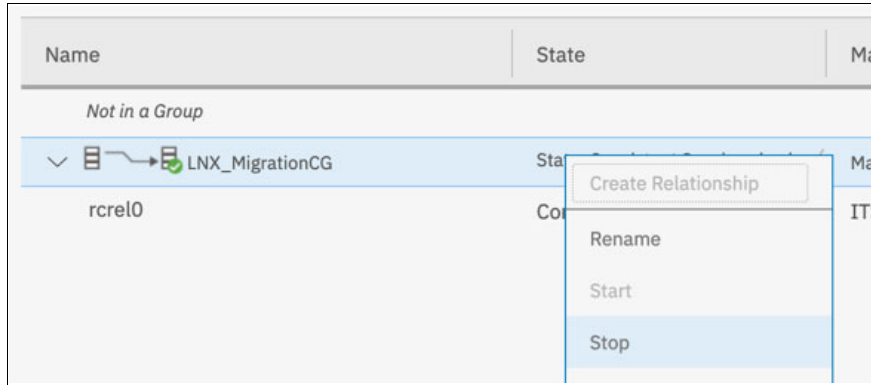


Figure 7-28 Stop consistency group menu

2. Make sure you check the **Allow secondary read/write access** option to allow the hosts to access the target volumes, then click **Stop Consistency Group** (Figure 7-29).

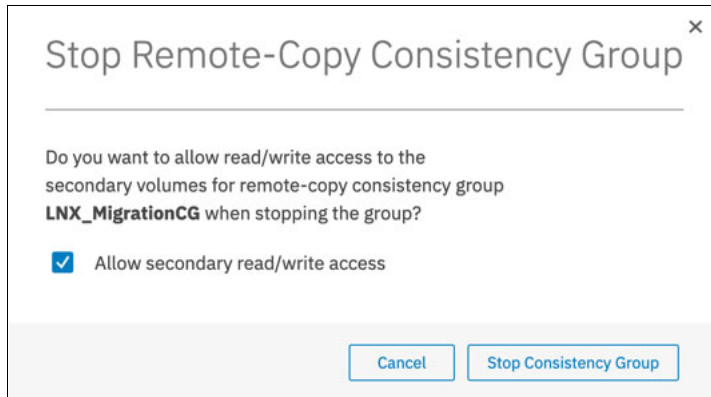


Figure 7-29 Stop Remote-Copy Consistency Group page

3. The consistency group will move to the **Idling** state. this means that the relationships still exist, but no data is being copied.

Name	State	Master Volume	Auxiliary Volume
Not in a Group			
LNX_MigrationCG	State: Idling	Master System: sbssvc	Auxiliary System: ITSO9100
rcrel0	Idling	ITSO_LNX_Vol_001	ITSO_LNX_Vol_001_GM_Target

Figure 7-30 Consistency group in Idling state

Deleting the relationship

Before you decouple the Master and Auxiliary volumes by deleting the Global Mirror relationship(s), you might want to make sure that all data has been correctly replicated to the Auxiliary site. This verification can be done with the `lsrcrelationship` command in the IBM Spectrum Virtualize CLI.

In CLI shell, issue `lsrcrelationship <relationship id or name>` command against the SAN Volume Controller system.

The relationship details are listed (see Example 7-2). The `sync` attribute has a value of `in_sync` when the contents are synchronized (identical) between volumes.

If write operations take place on either the master or auxiliary volume after a consistent (stopped) or idling state occurs, they will no longer be synchronized, and the `sync` attribute will have a value of `out_of_sync`.

Example 7-2 Verifying the relationship sync state

```
IBM_2145:sbssvc:superuser>lsrcrelationship rcrel0
id 0
name rcrel0
master_cluster_id 00000203200046A0
master_cluster_name sbssvc
master_vdisk_id 0
master_vdisk_name ITS0_LNX_Vol_001
aux_cluster_id 0000020420017E6C
aux_cluster_name ITS09100
aux_vdisk_id 24
aux_vdisk_name ITS0_GM_LNX_Vol_001_GM_Target
primary
consistency_group_id 0
consistency_group_name LNX_MigrateCG
state idling
bg_copy_priority 50
progress
freeze_time
status
sync in_sync
copy_type global
cycling_mode none
cycle_period_seconds 300
master_change_vdisk_id
master_change_vdisk_name
aux_change_vdisk_id
aux_change_vdisk_name
```

When you have verified the sync state of the relationships for the volumes you want to bring online on the new storage system, proceed with their deletion by removing the relationship, using the following steps:

1. Right-click the relationship you want to remove and select **Delete**, as shown in Figure 7-31 on page 117.

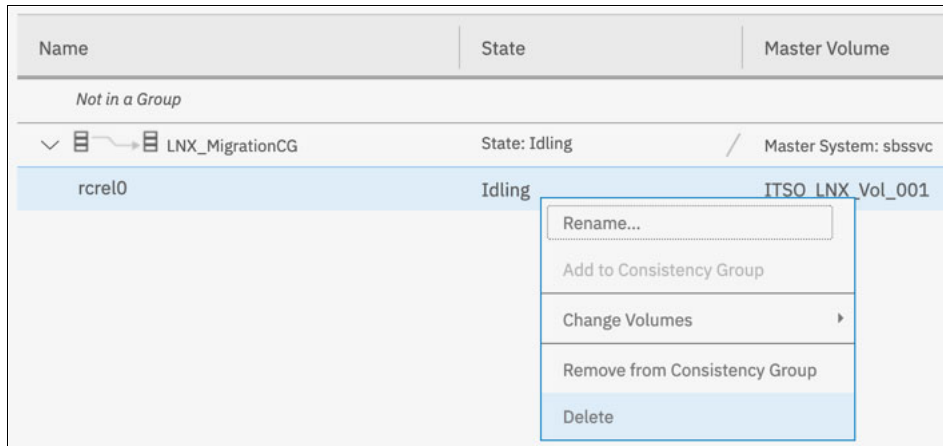


Figure 7-31 Delete Global Mirror relationship menu

2. Confirm the deletion in the **Delete Relationship** page shown in Figure 7-32.

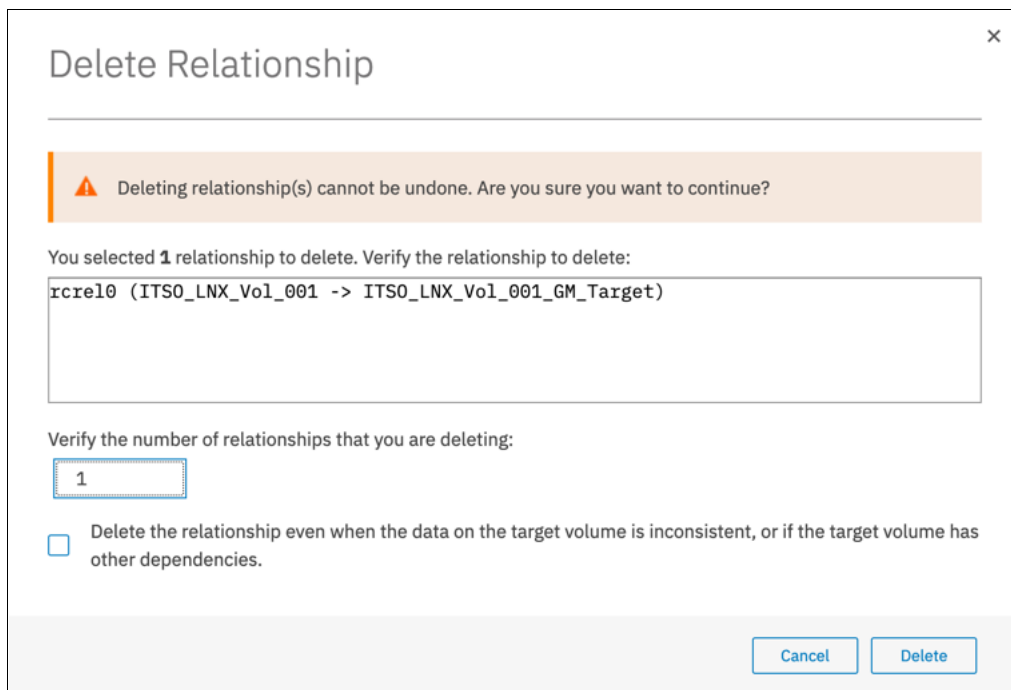


Figure 7-32 Confirm the relationship deletion

Map the volume on the new storage system and resume host IO

At this point the volume on the Auxiliary system is completely decoupled from the original volume on the Master system.

After you connect the host(s) to new storage system and define the corresponding host or host cluster object, you can present this volume to resume application IO.

Proceed as follows:

1. From the **Volumes** view, right-click the volume and select **Map to Host or Host Cluster...** (Figure 7-33).

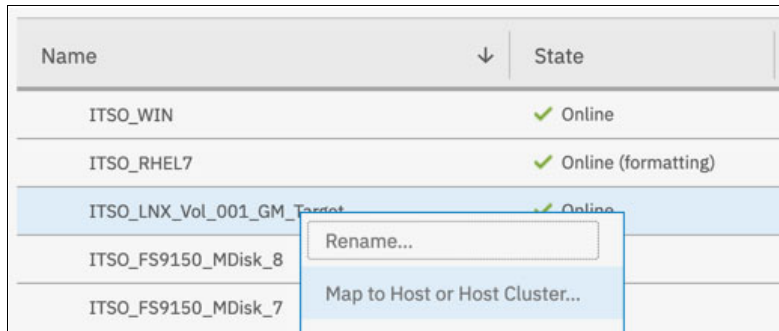


Figure 7-33 Map to Host or Host Cluster menu

2. Select the host or host cluster, then click **Next** (Figure 7-34).

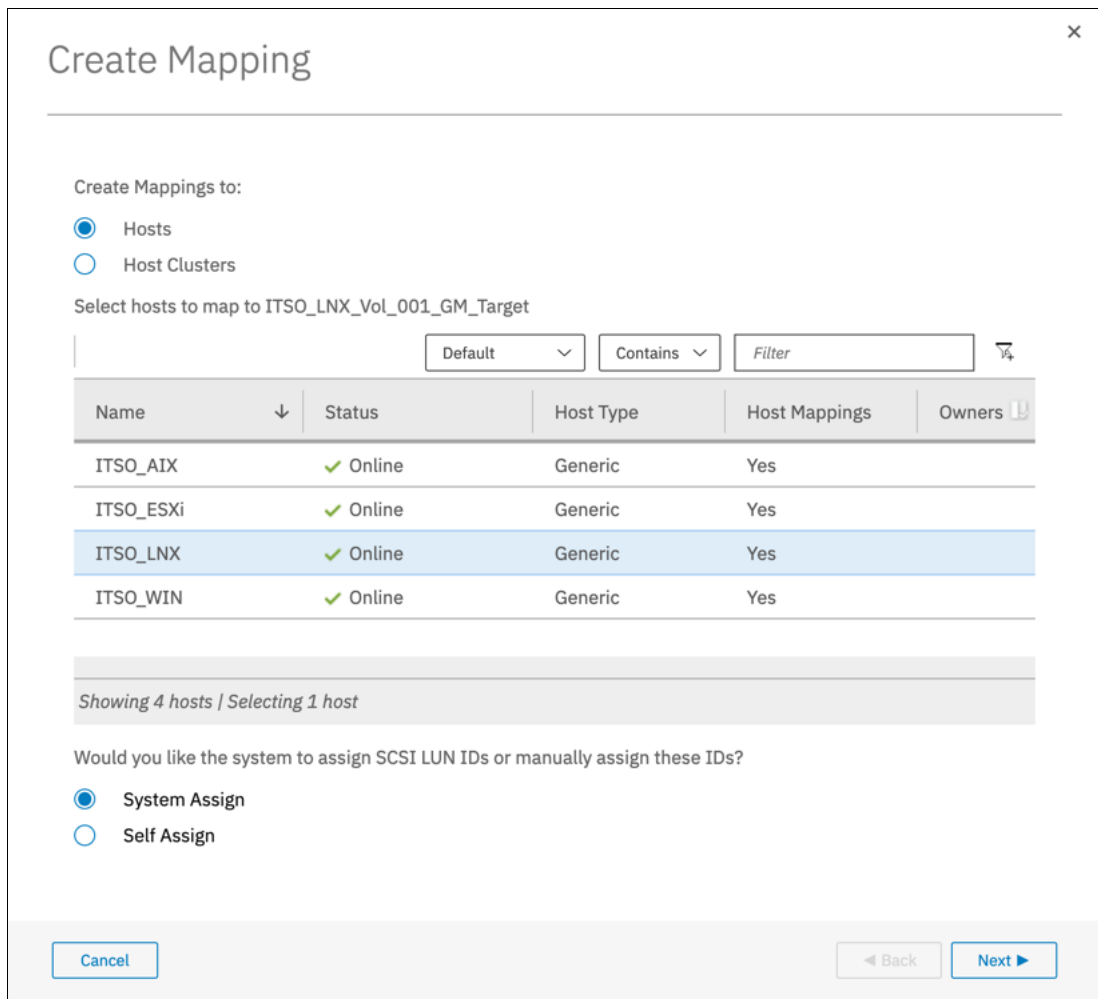


Figure 7-34 Create volume mapping page

3. Confirm the mapping. See Figure 7-35.

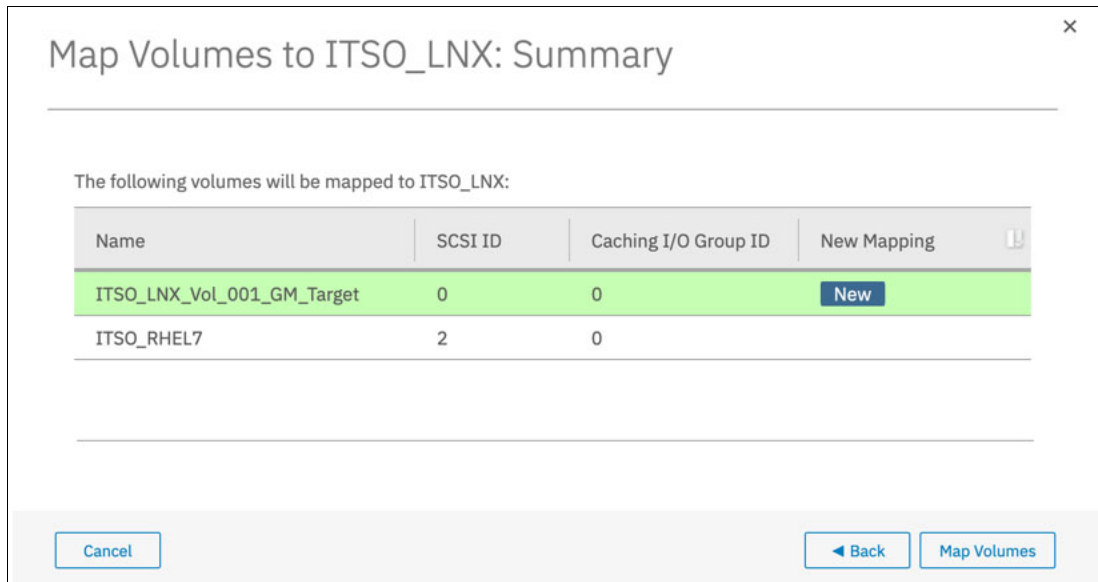


Figure 7-35 Mapping confirmation page

After you have mapped the volume(s), perform a rescan on your host, then reconfigure the device(s), file system(s), and so on, and verify the migrated data. Example 7-3 shows the commands used for the volume in our scenario.

Example 7-3 Verify migrated data

```
# Verify device details
[root@comp-1-31 ~]# multipath -ll
mpathd (360050768108005f9b0000000000001f) dm-5 IBM ,2145
size=50G features='1 queue_if_no_path' hwhandler='0' wp=rw
|-+- policy='service-time 0' prio=50 status=active
|  |- 5:0:1:1 sdd 8:48 active ready running
|  `-- 6:0:6:1 sdp 8:240 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
   |- 5:0:2:1 sdg 8:96 active ready running
   `-- 6:0:5:1 sdm 8:192 active ready running

# Mount the file system
[root@comp-1-31 ~]# mount /dev/mapper/mpathd1 /itso

# Verify file system details
[root@comp-1-31 ~]# df -h
Filesystem                Size  Used Avail Use% Mounted on
...
/dev/mapper/mpathd1        50G   53M   47G   1% /itso

#Verify file system data
[root@comp-1-31 ~]# ls /itso/
bert.txt  chris.txt  duncan.txt  francesco.txt  harry.txt  lost+found  markus.txt
matt.txt  roger.txt  steve.txt
```




IBM Spectrum Virtualize in the Cloud

This chapter describes the migration capabilities to the IBM Spectrum Virtualize product that is implemented in a cloud environment. This product is referred to as IBM Spectrum Virtualize for Public Cloud.

This chapter includes a brief overview of the technology that is behind the product and introduces the drivers and business values of using IBM Spectrum Virtualize in the context of public cloud services. It also shows a use case for migrating an on-premises solution to an IBM Spectrum Virtualize for Public Cloud implementation.

8.1 IBM Spectrum Virtualize for Public Cloud

Companies recognize the value of delivering services by using the cloud, and many of them are already using public clouds to some degree. The role of the cloud is maturing and it is more often being considered as a platform for innovation and business value.

Nevertheless, one of the challenges is how to integrate those public cloud capabilities with the existing on premises, back-end systems. Organizations want to retain flexibility without introducing new complexity or requiring significant new capital investment.

IBM Spectrum Virtualize for Public Cloud can support clients in their IT architectural transformation and migration towards the cloud service model, enabling hybrid cloud strategies or, for a multicloud workload, providing the benefits of familiar and sophisticated storage functions on public cloud data centers, enhancing the existing cloud offering.

In the specific context of single-site deployment, IBM Spectrum Virtualize for Public Cloud supports extra features that enhance the public cloud block-storage offering. At the storage level, IBM Spectrum Virtualize for Public Cloud resolves some limitations due to the standardized model of public cloud providers: a maximum number of LUNs per host, a maximum volume size, and poor granularity in the choice of tiers for storage snapshots.

Designed for Software Designed Storage (SDS) environments, IBM Spectrum Virtualize for Public Cloud represents a solution for public cloud implementations, and includes technologies that both complement and enhance public cloud offering capabilities.

Table 8-1 elaborates on the benefits of an IBM Spectrum Virtualize for Public Cloud single site deployment.

Table 8-1 Benefits of IBM Spectrum Virtualize for Public Cloud single site deployment.

Feature	Benefits
Single point of control for cloud storage resources.	<ul style="list-style-type: none"> ▶ Designed to increase management efficiency and to help to support application availability.
Pools the capacity of multiple storage volumes.	<ul style="list-style-type: none"> ▶ Helps to overcome volume size limitations. ▶ Helps to manage storage as a resource to meet business requirements, and not just as a set of independent volumes. ▶ Helps administrator to better deploy storage as required beyond traditional “islands”. ▶ Can help to increase the use of storage assets. ▶ Insulate applications from maintenance or changes to a storage volume offering.
Manages tiered storage.	<ul style="list-style-type: none"> ▶ Helps to balance performance needs against infrastructures costs in a tiered storage environment. ▶ Automated policy-driven control to put data in the right place at the right time automatically among different storage tiers and classes.

Feature	Benefits
Easy-to-use IBM FlashSystem family management interface.	<ul style="list-style-type: none"> ▶ Has a single interface for storage configuration, management, and service tasks regardless of the configuration that is available from the public cloud portal. ▶ Same single interface as for on-premises systems and same support for as many as 440 different vendor systems. ▶ Helps administrators use storage assets and volumes more efficiently. ▶ Has IBM Spectrum Control and IBM Spectrum Protect for extra capabilities to manage capacity and performance.
Dynamic data migration.	<ul style="list-style-type: none"> ▶ Migrates data among volumes and LUNs without taking applications that use that data offline. ▶ Manages and scales storage capacity without disrupting applications.
Data Reduction	<ul style="list-style-type: none"> ▶ Starting with IBM Spectrum Virtualize v 8.3.1, Data Reduction Pools are supported on Cloud.
Advanced network-based copy services.	<ul style="list-style-type: none"> ▶ Copy data across multiple storage systems with IBM FlashCopy. ▶ Copy data across metropolitan and global distances as needed to create high-availability storage solutions between multiple data centers.
Thin provisioning and snapshot replication.	<ul style="list-style-type: none"> ▶ Reduces volume requirements by using storage only when data changes. ▶ Improves storage administrator productivity through automated on-demand storage provisioning. ▶ Snapshots are available on lower tier storage volumes.
IBM Spectrum Protect Snapshot application-aware snapshots.	<ul style="list-style-type: none"> ▶ Perform near-instant and application-aware snapshot backups, with minimal performance impact for IBM Db2, Oracle, SAP, VMware, Microsoft SQL Server, and Microsoft Exchange. ▶ Provide advanced and granular restoration of Microsoft Exchange data.
Third-party native integration	<ul style="list-style-type: none"> ▶ Integration with VMware vRealize.

8.2 Workload relocation into the public cloud

In this section, a use case for IBM Spectrum Virtualize for Public Cloud is illustrated where an entire workload segment is migrated from a client's enterprise into the cloud. The suggested process for relocating a workload into the cloud is by using IBM Spectrum Virtualize Remote Copy; however, there are also other mechanisms that can accomplish this task.

8.2.1 Business justification

All the drivers that motivate businesses to use virtualization technologies make deploying services into the cloud even more compelling because the cost of idle resources is further absorbed by the cloud provider. However, certain limitations in regulatory or process controls may prevent a business from moving all workloads and application services into the cloud.

An ideal case with regards to a hybrid cloud solution is the relocation of a specific segment of the environment that is well suited, such as development. Another might be a specific application group that does not require either the regulatory isolation or low response time integration with on-premises applications.

Although performance might be a factor, it should not be assumed that cloud deployments automatically create a diminished performance. Depending on the location of the cloud service data center and the intended audience for the migrated service, the performance can conceivably be superior to on-premises pre-migration.

In summary, moving a workload into the cloud might provide similar functions with better economies due to scaling of physical resources in the cloud provider. Moreover, the cost of services in the cloud is structured, measurable, and predictable.

8.2.2 Data migration

The suggested method for data migration to the cloud is to use IBM Spectrum Virtualize Remote Copy services. For instance, Global Mirror or Global Mirror with Change Volumes, allows for an asynchronous copy of data to be sent to the cloud and synced from your on-premises copy. When you are ready for the cloud instance to become production, you are able to stop your host I/O and applications, ensure that your latest data is synced in the cloud, and change your host to now point to your cloud data and restart your applications.

Other methods like host-side migrations which are discussed in later chapters may be available for cloud migrations. Host-based migrations, when relocating to the cloud, are often impractical because each host you want to migrate must access both source and target environments simultaneously with iSCSI being the only connection method currently available for IBM Spectrum Virtualize in the Public Cloud.

8.2.3 Host provisioning

In addition to the replication of data, it may be necessary for compute nodes and networking to be provisioned within the cloud provider upon which to run the relocated workload. Currently, in the AWS Cloud the EC2 compute nodes are available with storage that is provisioned to the EC2 compute instance by using an iSCSI connection. In the IBM Cloud™, bare-metal and virtual servers are available. Within the bare-metal options, Intel processor-based servers and Power based servers with OpenPOWER provide high performance on Linux-based platforms. As IBM Spectrum Virtualize in the Public Cloud matures and expands into other Cloud Service Providers, other platforms might become available.

8.2.4 Implementation considerations

The following list describes implementation considerations for the workload relocation into the public cloud use case:

- ▶ Naming conventions: This is an important consideration in the manageability of a standard on-premises IBM Spectrum Virtualize environment, but given the multiple layers of virtualization in a cloud implementation, maintaining a consistent and meaningful naming convention for all objects (managed disks (MDisks), volumes, FlashCopy mappings, Remote Copy relationships, hosts, and host clusters) is necessary.

- ▶ **Monitoring integration:** Integration into IBM Storage Insights, IBM Spectrum Control or some other performance monitoring framework is useful for maintaining metrics for reporting or troubleshooting. IBM Storage Insights and IBM Spectrum Control are well suited for managing IBM Spectrum Virtualize environments.
- ▶ **Planning and scheduling:** Regardless of the method that is chosen, gather as much information ahead of time as possible: File system information, application custodians, full impact analysis of related systems, and so on.
- ▶ **Be sure to ensure a solid backout:** If inter-related systems or other circumstances require rolling back the application servers to on-premises, plan the migration to ensure as little difficulty as possible in the roll-back, which might mean keeping zoning in the library (even if it is not in the active configuration), and not deleting source volumes for a certain period.

8.3 Migration using Global Mirror with Change Volumes

This section describes how to configure replication from an on-premises solution that could be any IBM Spectrum Virtualize system to an IBM Spectrum Virtualize for Public Cloud on AWS solution. It is possible to manage all IBM Spectrum Virtualize, including for Public Cloud, using web-based browsers for simplified GUI management and monitoring.

Our example uses a SAN Volume Controller in the on-premises data center referred to as “SVC” and a two-node IBM Spectrum Virtualize for Public Cloud on AWS referred to as “AWS Cloud”, as shown in Figure 8-1.

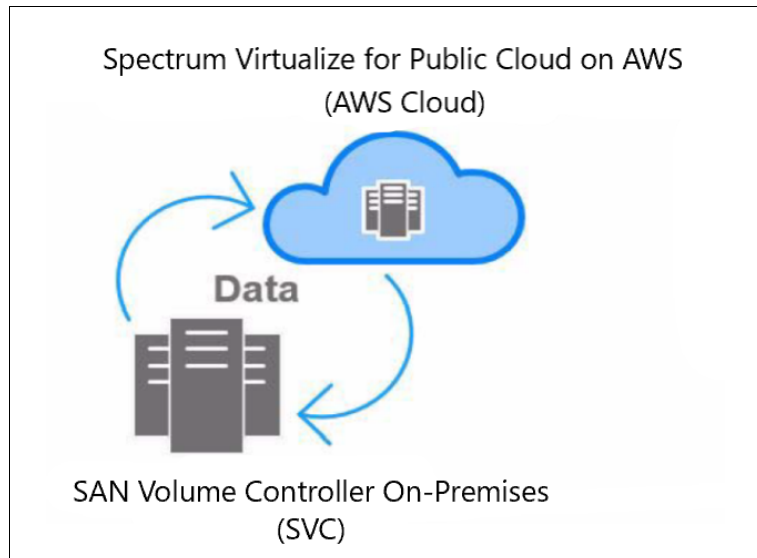


Figure 8-1 Use case migrating SAN Volume Controller data to the cloud.

This scenario uses IBM Spectrum Virtualize Global Mirror with Change Volumes (GMCV) to replicate the data from the SAN Volume Controller to AWS Cloud.

Because IBM Spectrum Virtual on Public Cloud is available only through iSCSI connections, this implementation starts with the assumption that the IP connectivity between the on-premises SAN Volume Controller and AWS Cloud is established through a Multiprotocol Label Switching (MPLS) or VPN connection. Because there are multiple ways to implement the IP connectivity, this section does not consider that specific configuration. For more information, contact your organizations’s network technical specialist.

To configure the GMCV, complete the following steps:

1. On the AWS Cloud, configure your IBM Spectrum Virtualize Private IP ports in the **Settings** → **Network** → **Ethernet Ports** menu. Select the wanted port, right-click and select **Modify Remote Copy** so that the port is enabled for remote copy, as shown in Figure 8-2.

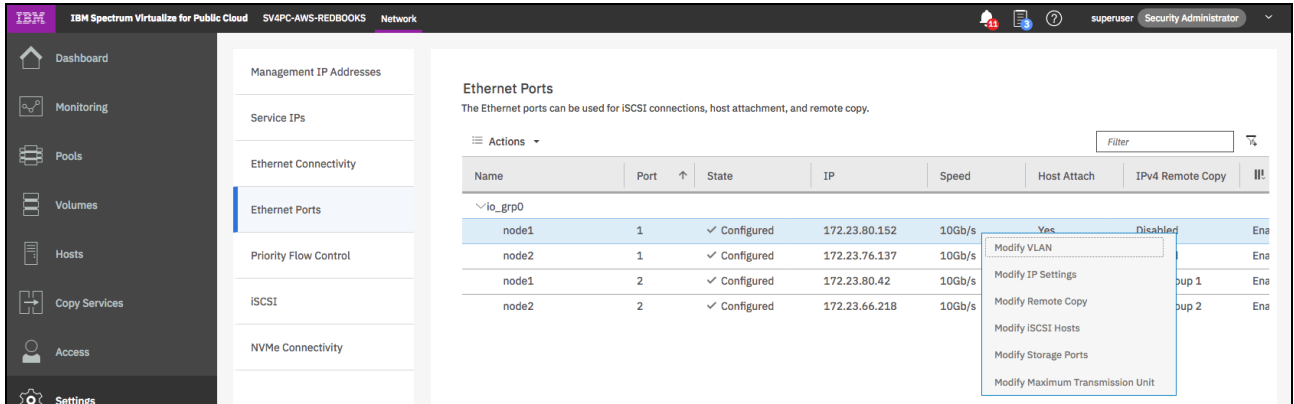


Figure 8-2 Remote copy IP port example

- a. You are redirected to choose which copy group to use, as shown in Figure 8-3.

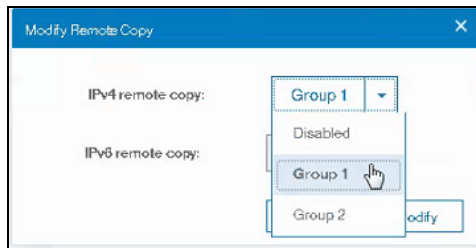


Figure 8-3 Group 1 configuration example

- b. Repeat steps 1 on page 126 and a for all of the IP ports that you want to configure, and you end up with a similar configuration as shown in Figure 8-4.

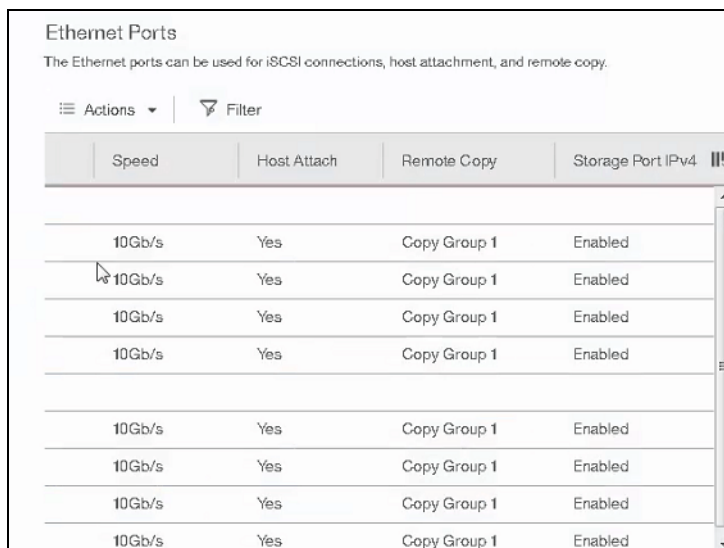


Figure 8-4 IBM Spectrum Virtualize configuration complete

- c. Run the same configuration for the on-premises SAN Volume Controller, as shown in Figure 8-5.

Note: It is important to understand what versions of IBM Spectrum Virtualize software are supported. For supported and interoperability versions, see [IBM Spectrum Virtualize Family of Products Inter-System Metro Mirror and Global Mirror Compatibility Cross Reference](#).

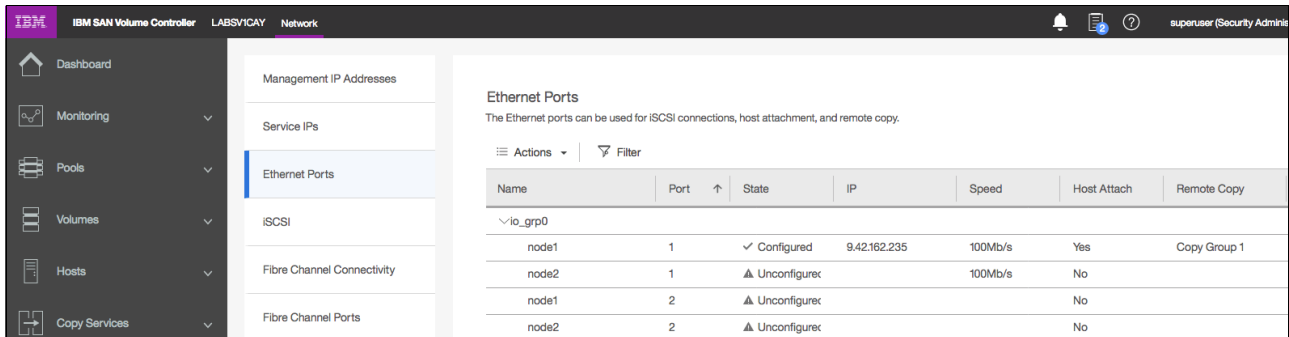


Figure 8-5 On-premises SAN Volume Controller configuration completion example

2. Create a cluster partnership between the on-premises data center and IBM Spectrum Virtualize for Public Cloud on AWS from the on-premises GUI by going to **Copy Services** → **Partnerships** and clicking **Create Partnership**, as shown in Figure 8-6.

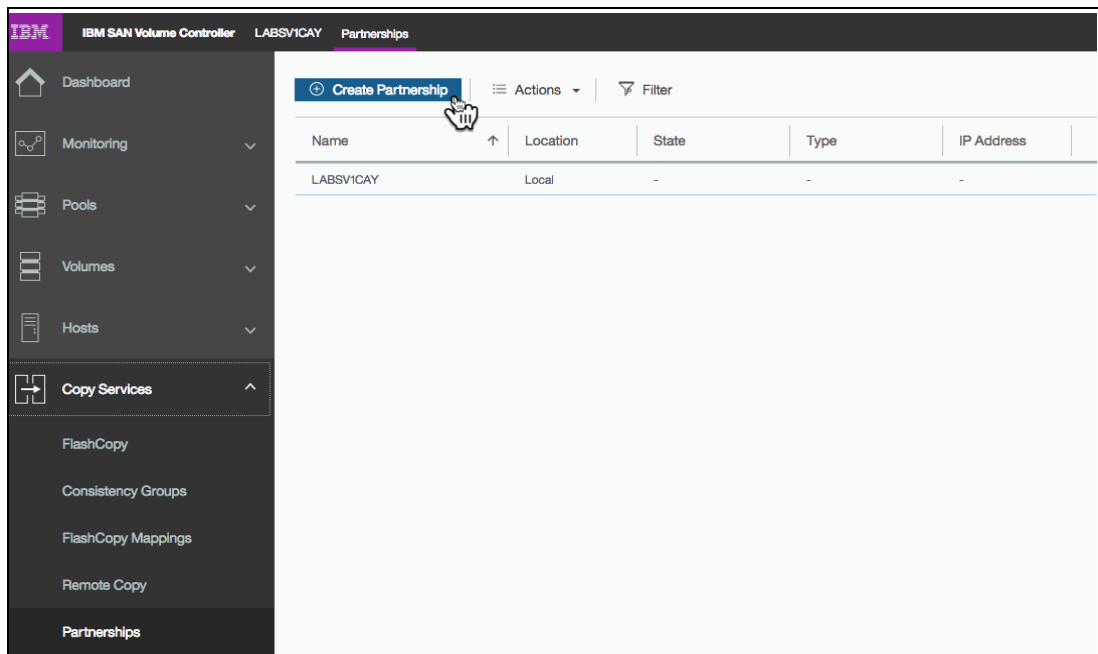


Figure 8-6 Create Partnership setup example

3. Complete the partnership creation from on-premises SAN Volume Controller, as shown in Figure 8-7.

This system is in the replication layer. Partnerships can be created with other systems in the replication layer.

Type: Fibre Channel IP

Partner system IP address:

Link bandwidth: Mbps

Background copy rate: %

Partner system's CHAP secret:

Figure 8-7 Inserting an IP address example

4. As you can see from the SAN Volume Controller GUI in Figure 8-8, the partnership is partially configured.

Name	Location	State	Type	IP Address
LABSVICAY	Local	-	-	-
IAMBACK	Remote	▲ Partially Configur...	IPv4	9.42.162.65

Figure 8-8 Partnership partially configured

5. Next you must complete the partnership configuration in the AWS Cloud GUI by going to **Copy Services** → **Partnerships** and clicking **Create Partnership**, as shown in Figure 8-9 and Figure 8-10 on page 129.

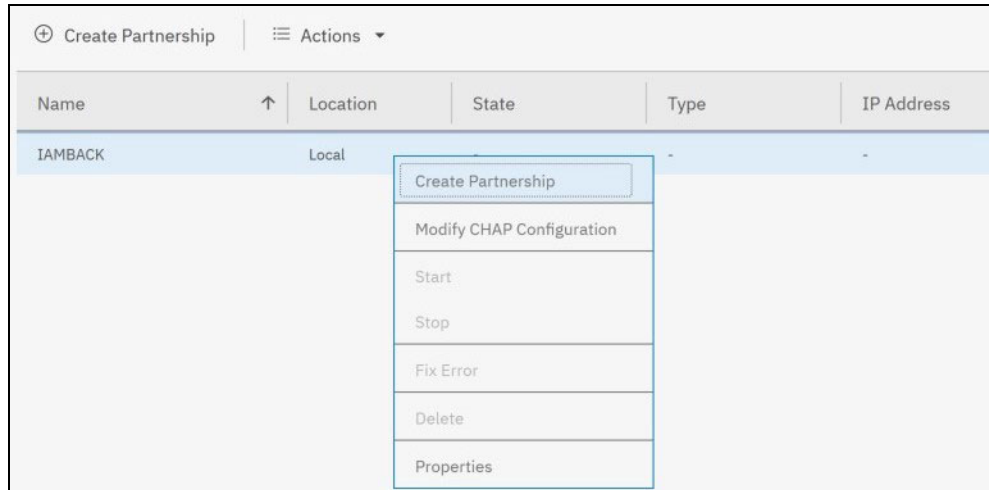


Figure 8-9 Create partnership example from AWS Cloud

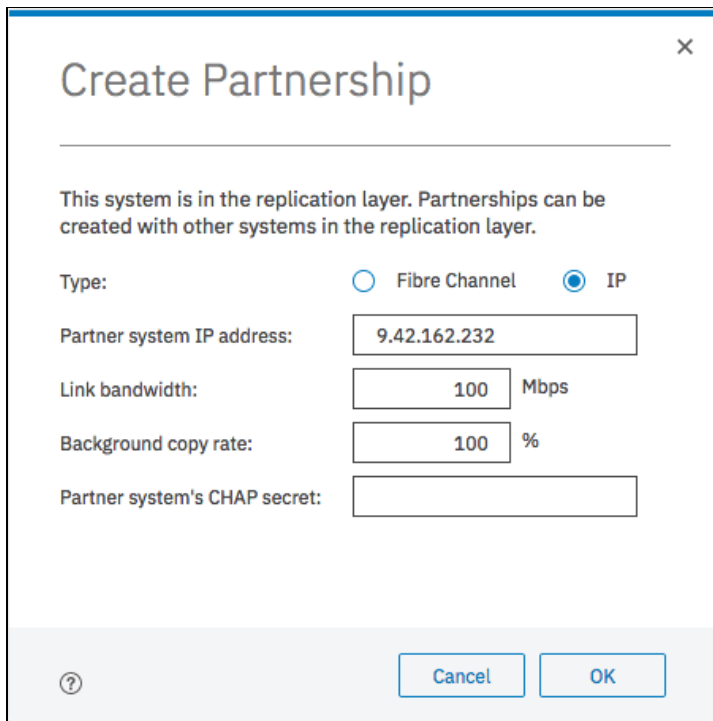


Figure 8-10 Partnership example

6. Now, your partnership is fully configured, as shown in Figure 8-11.

Name	Location	State	Type	IP Address
LABSV1CAY	Remote	✓ Fully Configured	IPv4	9.42.162.232
IAMBACK	Local	-	-	-

Figure 8-11 Fully configured example

Note: The connection might take a few seconds to synchronize, but double-clicking **Partnership** reveals the confirmed status of the partnership quicker.

- In our example, we have existing volumes on the SAN Volume Controller, a 100 GiB volume named SVCPC_AWS with a 100 GiB Change Volume (CV) named SVCPC_AWS_CV that must be replicated to the AWS Cloud instance that is defined.

Create your AWS volumes if they haven't been created and ensure they are large enough to handle the data you want to migrate. The on-premises volumes are thin-provisioned, but this is not a specific requirement. It is just a choice. The CV can be thin-provisioned or fully provisioned, regardless of whether the master or auxiliary volume is thin-provisioned or space-efficient.

The CV must store only the changes that accumulated during the cycle period, so it should use real capacity when possible, as shown in Figure 8-12.

Name	State	Pool	Protocol Type	UID	Capacity
SVCPC_AWS	Online	Pool1		600507607293A2FAA8000000000000...	100.00 GiB
SVCPC_AWS_CV	Online	Pool1		600507607293A2FAA8000000000000...	100.00 GiB

Figure 8-12 Volumes example

- Create a volume remote copy relationship for a GMCV from the on-premises SAN Volume Controller GUI in **Copy Services** → **Remote Copy** followed by **Actions** → **Create Relationship**, as shown in Figure 8-13.

Name	State	Master Volume	Auxiliary Volume
> Not in a Group			

Figure 8-13 Creating a Remote Copy relationship

- Select the type of relationship. Here we select **Global Mirror with Change Volumes**, as shown in Figure 8-14 on page 131.

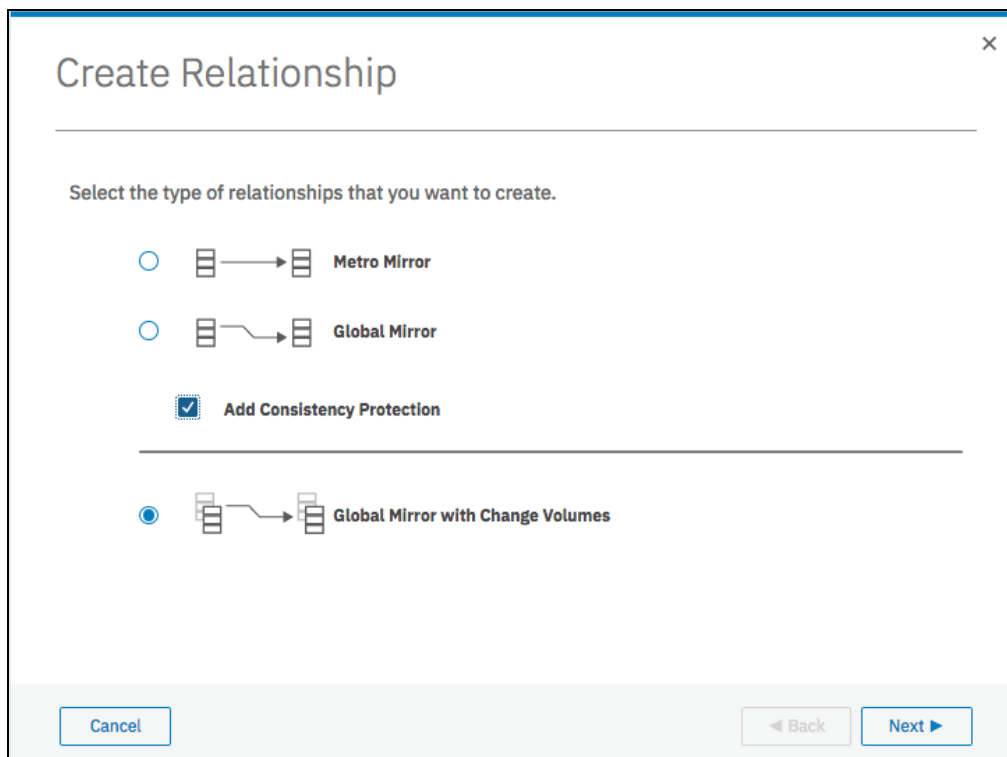


Figure 8-14 Global Mirroring with Change Volumes example

10. Next select where the AWS Cloud (auxiliary) volumes are located. Select the **On another system** option and select the system name, as shown in Figure 8-15.

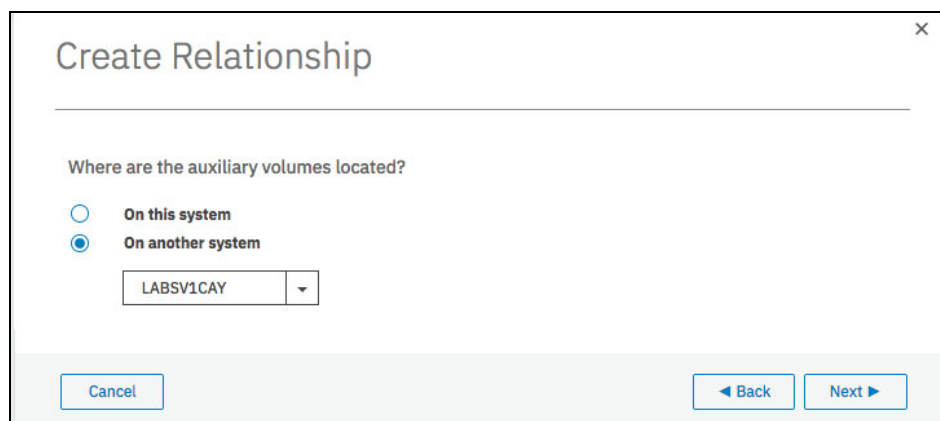


Figure 8-15 Create relationship to remote system

11. Next select the Volumes to be used in the relationship on both the SAN Volume Controller (Master) and the AWS Cloud (Auxiliary), as seen in Figure 8-16 on page 132. Note volumes on your auxiliary system should already be created with sufficient size to migrate data to. Below the auxiliary volume, you need to click **Add** to add the relationship.

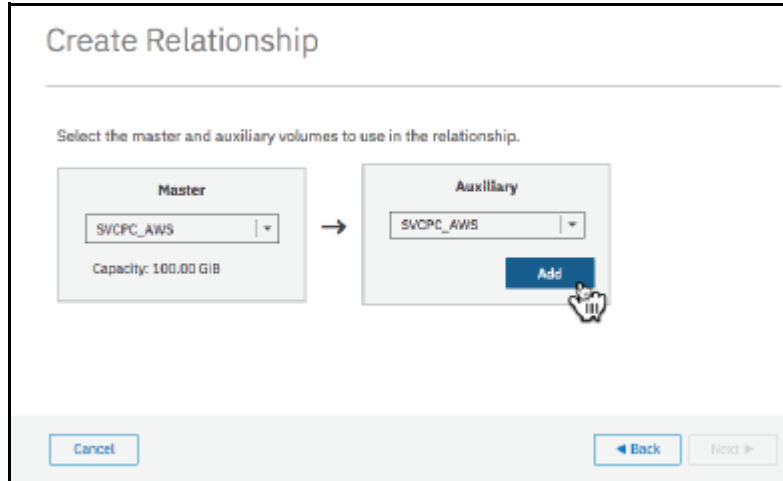


Figure 8-16 Master and auxiliary volumes example

12. In our example, we select **No, do not add volumes**, as shown in Figure 8-17, to add a Master Change Volume. They will be added later, click **Finish**.

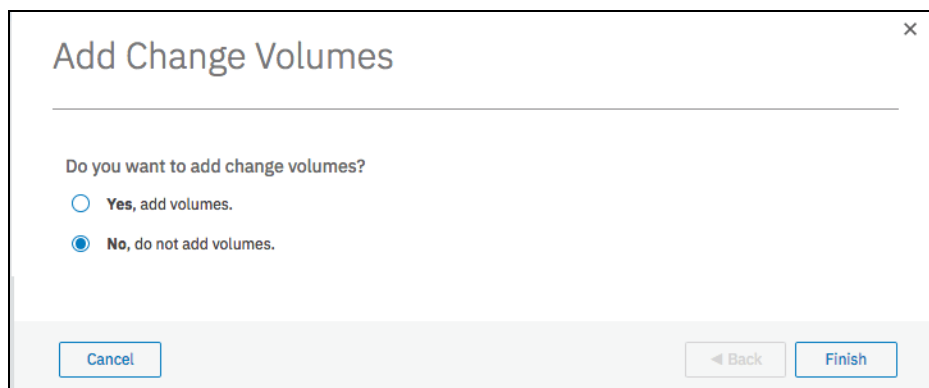


Figure 8-17 Do not add master change volume example

13. You will be back on the Create Relationship screen where you can now click the **Next** button where you will be prompted if the volumes are synchronized. We select **No, the volumes are not synchronized**, as shown in Figure 8-18 and select **Next**.

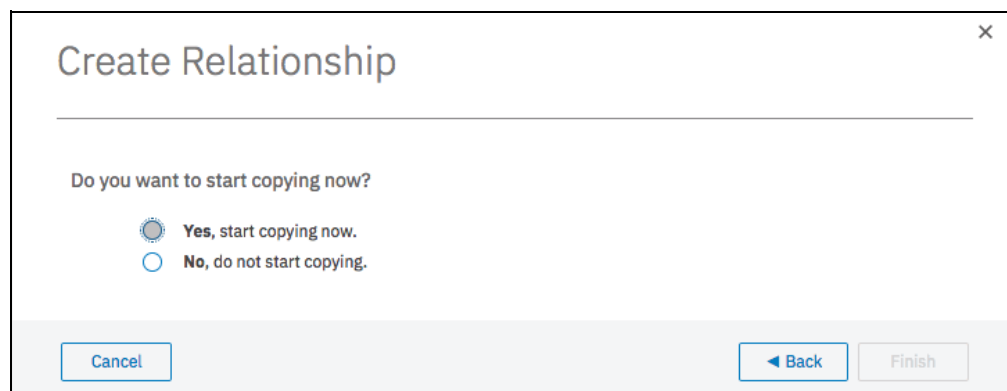


Figure 8-18 No volumes are not synchronized

14. You will be prompted for one more question if you want to start copying data. We select **No, do not start copying** (see Figure 8-19) as we will start the relationship copying later and click the **Finish** button.

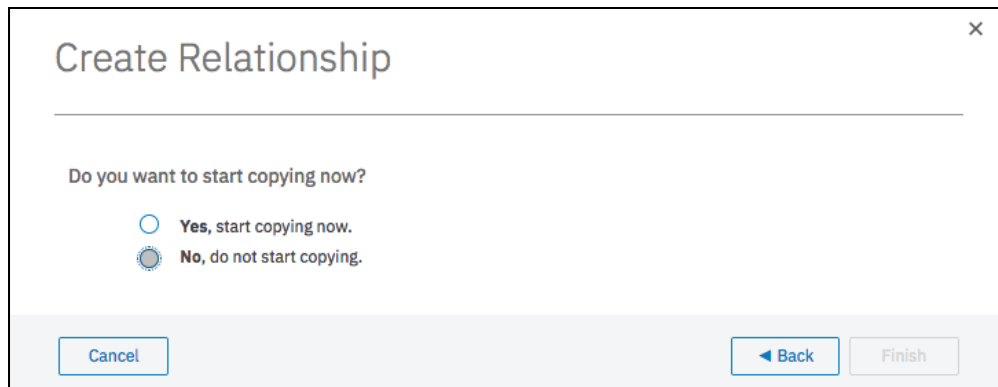


Figure 8-19 Do not start relationship example

15. You will now see the relationship under the Not in a Group drop-down. You can right-click the relationship and click **Edit Relationship** to set the cycling mode and period. Edit the relationship and set the **Cycling Mode** to **Multiple** and **Cycling Period** to **300** (see Figure 8-20) and then, click **OK**.

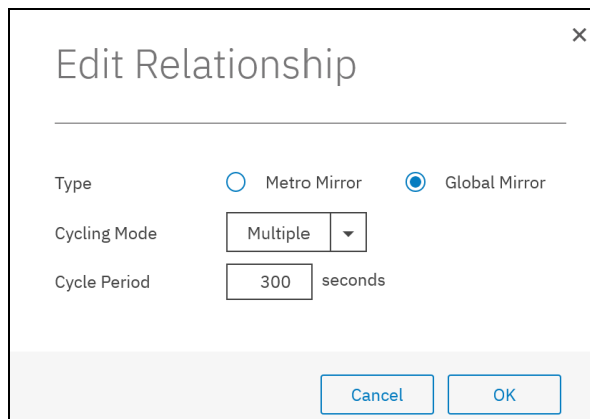


Figure 8-20 Edit relationship to set Cycling mode and period

16. Add the CV volumes to your relationship on both the SAN Volume Controller and the AWS Cloud GUI by right-clicking the relationship, then select **Change Volumes** → **Add Existing** as shown in Figure 8-21. Select the CV in the drop-down menu, as shown in Figure 8-22 and click **Add** and confirm the cycling settings are still accurate and click **OK** as shown in Figure 8-23. Do this on both sites.

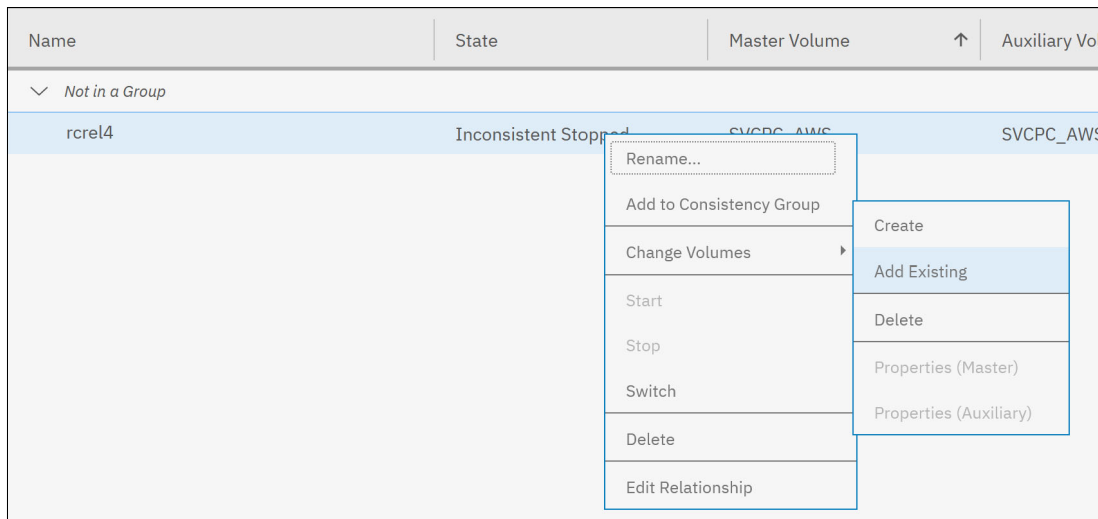


Figure 8-21 Adding a change volume from the on-premises site

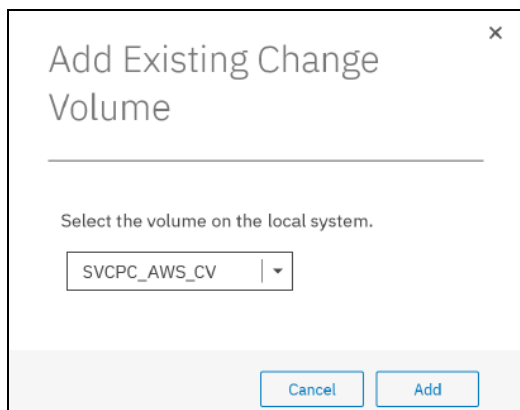


Figure 8-22 Selecting the change volume from the on-premises site

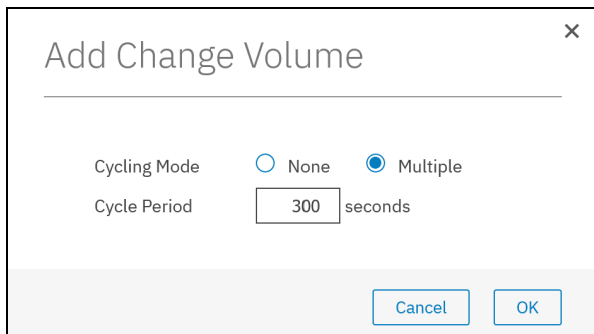


Figure 8-23 Confirm Cycle settings

17. Now, start your relationship from the on-premises SAN Volume Controller site by right-clicking the relationship and selecting **Start**, as shown in Figure 8-24.

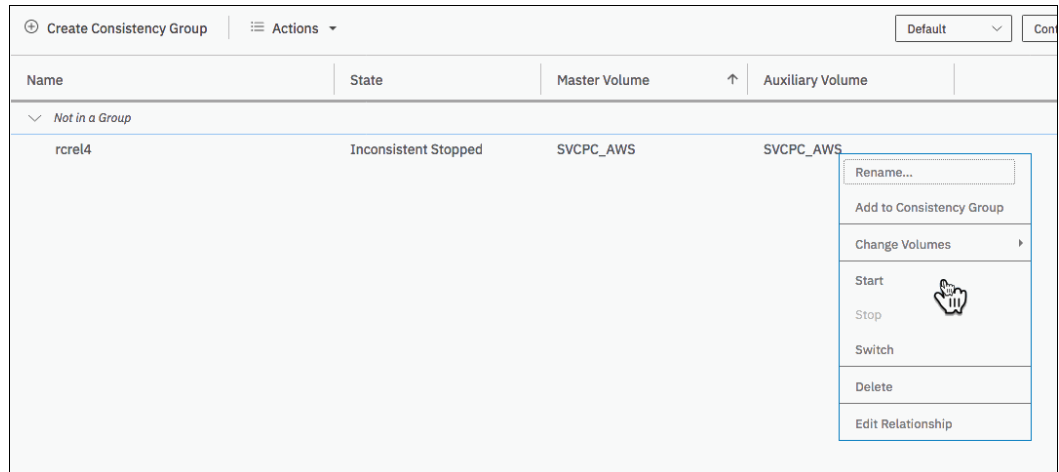


Figure 8-24 Starting the relationship

18. Now, you can create a GM consistency group by clicking **Create Consistency Group**, enter in the name, in our case **SVCPC_AWS_RB** and click **Next** as show in Figure 8-25.



Figure 8-25 Create Consistency Group

19. You will be prompted to enter where the auxiliary volumes are located, select the AWS Cloud system as seen in Figure 8-15 on page 131 and click next, and then select **No, create an empty consistency group**, as you will add the existing relationship in the next step and click **Finish** as seen in Figure 8-26.

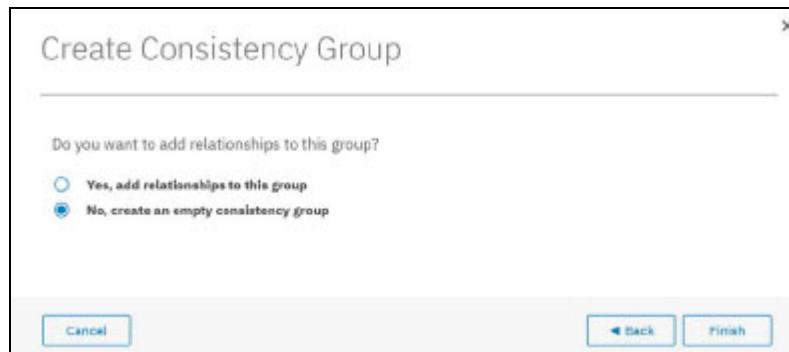


Figure 8-26 Create Empty Consistency Group

20. Now you can add your relationship to the Consistency Group you created, by right-clicking the relationship selecting **Add to Consistency Group** as shown in Figure 8-27 and selecting the Consistency Group and clicking **Add to Consistency Group** as seen in Figure 8-28.

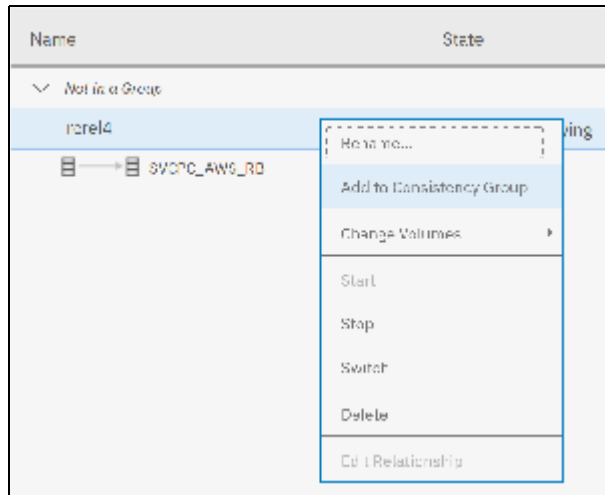


Figure 8-27 Adding a consistency group

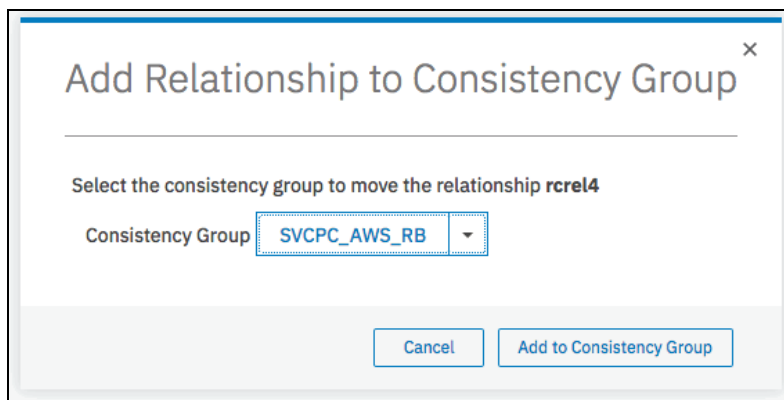


Figure 8-28 Add Relationship to Consistency Group

21. Now you can see the status of your consistency group, as shown in Figure 8-29 and Figure 8-30 on page 137.

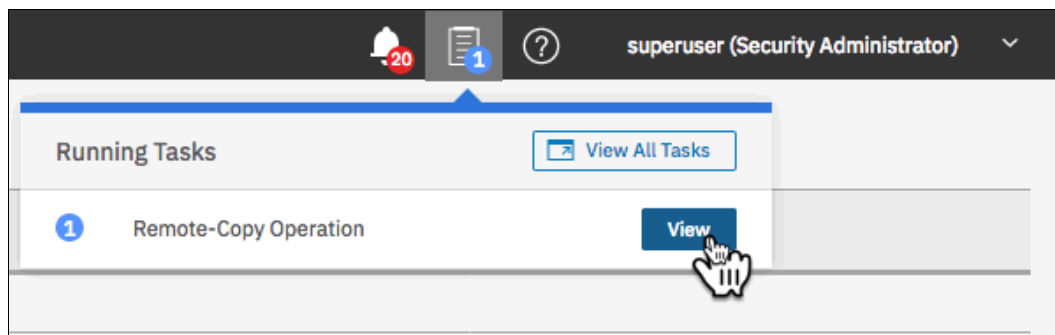


Figure 8-29 Consistency group status




Figure 8-30 Copying status

In our example, we show the status by using the AWS Cloud GUI.

When the copy approaches completion, the CV algorithm starts to prepare a freeze time in accordance with the cycling windows. When your copy reaches 100%, a FlashCopy is taken from the auxiliary volume to the auxiliary-CV to be used in case of real disaster or DR test.

At 100%, the status is *consistent copying*.

8.3.1 Stopping the Remote Copy and switching to the new storage system

At any point in time after the initial synchronization of the data between the migration source system and the migration target system, you can follow the following steps to switch your application workload to the new IBM Spectrum Virtualize instance in the cloud.

This requires the application to be shut down because you have to switch application access from existing to cloud storage.

Before starting these steps, shut down the application and prepare the system to move to the new storage. In general, these steps include:

- ▶ Quiesce all I/O on the affected volumes;
- ▶ Unmount all file systems insisting on the affected volumes;
- ▶ Remove the corresponding devices;
- ▶ Unmap the source volume(s) from the host(s);
- ▶ Verify that all changes have been propagated to the auxiliary site;
- ▶ Stop the Global Mirror relationship;
- ▶ Remove the Global Mirror relationships;
- ▶ Map the target volume(s) on the auxiliary system to the host(s);
- ▶ Reconfigure devices, file systems etc. on the host(s);
- ▶ Verify integrity of data on the new volume(s);
- ▶ Resume I/O.

Unmapping the volume from the old storage system

After you have unconfigured and removed the volume at the host level, you can safely unmap it from the old storage system. In order to do that, perform the following steps:

1. In the **Volumes** → **Volumes** view, right-click the volume and select **Remove Private Mappings** (Figure 8-31 on page 138).

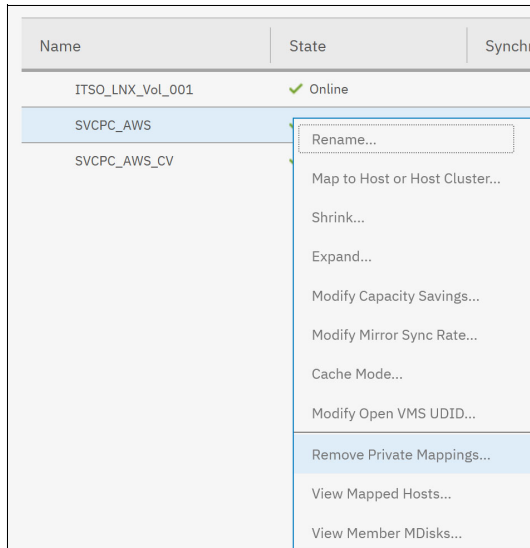


Figure 8-31 Remove Private mappings menu

2. In the **Remove Private Mappings** page, confirm your action to unmap.

Stopping the Remote Copy relationships

Before stopping your consistency group, be sure that there has not been any new I/O to your volumes after the time stamp of your freeze time for the consistency group. You can see the consistency group freeze time, as an additional column, if not shown in your default view of the GUI, as seen in Figure 8-32. You can also see it via the CLI using the `lsrcrelationship` command, as seen in Example 8-1.

Name	State	Master Volume	Auxiliary Volume	Freeze Time
<i>Not in a Group</i>				
<div style="display: flex; align-items: center;"> ▼ 📁 → SVCPC_AWS_RB </div>	State: Consistent Copying	Master System: IAMBACK	Auxiliary System: LABSV1CAY	Freeze Time: Feb 28, 2020, 11:04:45 /
rcrel4	Consistent Copying	SVCPC_AWS	SVCPC_AWS	Feb 28, 2020, 11:04:45 AM

Figure 8-32 Confirm that Freeze Time is after last I/O to your volumes

Example 8-1 Verifying the relationship sync state

```
IBM>lsrcrelationship rcrel4
...
progress 100
freeze_time 2020/02/28/11/04/45
status online
sync in-sync
...
```

To stop the Global Mirror relationship using the IBM Spectrum Virtualize GUI, perform the following steps:

1. In the Master system's **Copy Services** → **Remote Copy** view, right-click the consistency group containing the relationships for the volumes you want to switch access to and select **Stop**, as shown in Figure 8-33 on page 139.

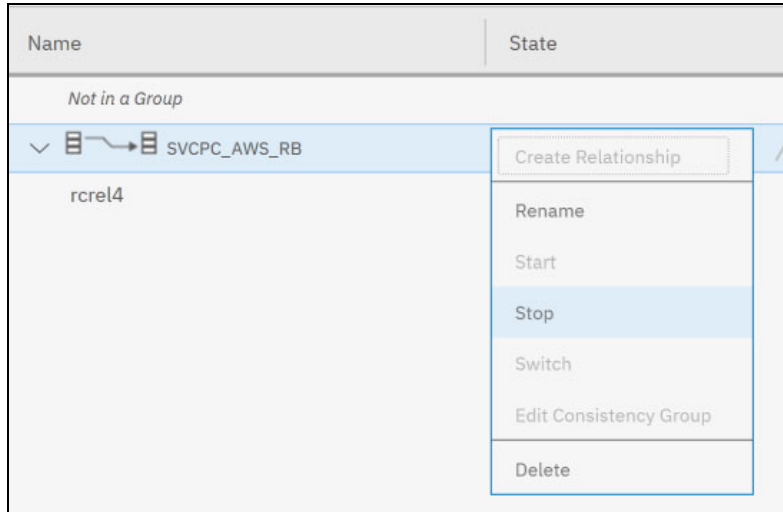


Figure 8-33 Stop consistency group menu

2. Make sure you check the **Allow secondary read/write access** option to allow the host(s) to access the target volume(s), then click **Stop Consistency Group** (Figure 8-34).

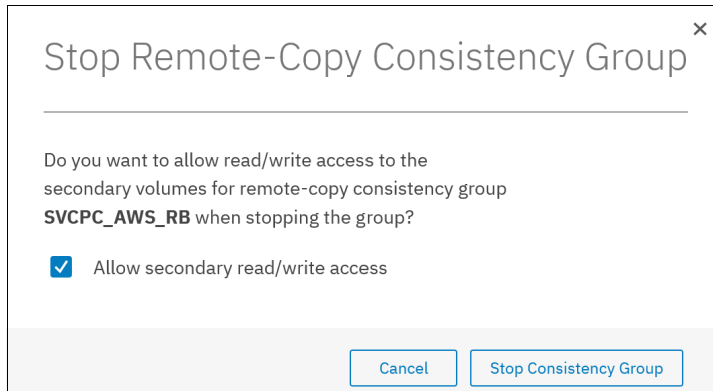


Figure 8-34 Stop Remote-Copy Consistency Group page

3. The consistency group will move to the **Idling** state (see Figure 8-35), which means that the relationships still exist, but no data is being copied.

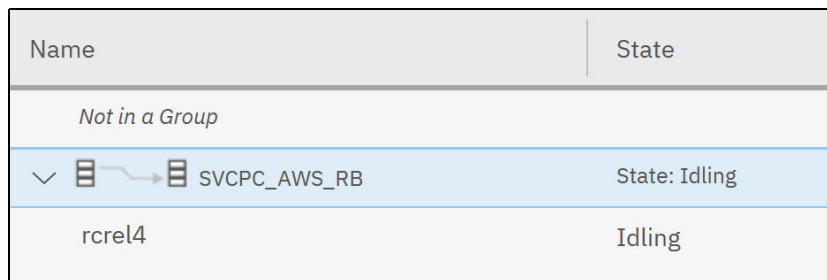


Figure 8-35 Consistency group in Idling state

Deleting the relationship

When you have verified that the sync was complete and stopped the relationships for the volumes you want to bring online on the new storage system, proceed with the relationship removal using the following steps:

1. Right-click the relationship you want to remove and select **Delete**, as shown in Figure 8-36.

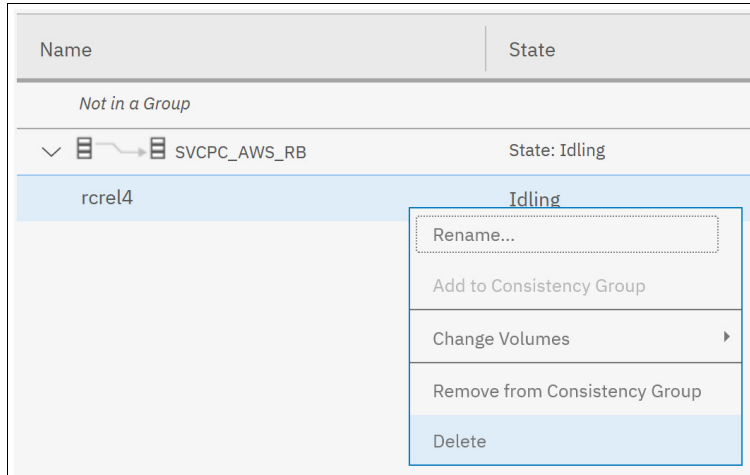


Figure 8-36 Delete Global Mirror with Change Volume relationship menu

2. Complete the number of relationships and confirm the deletion in the **Delete Relationship** page (Figure 8-37).

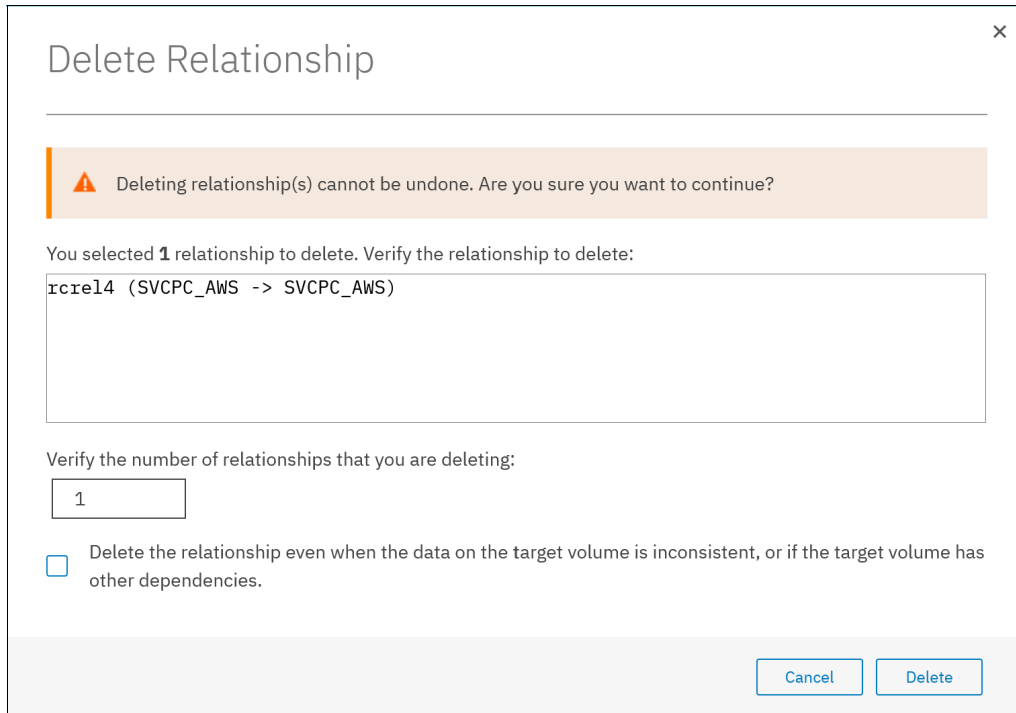


Figure 8-37 Confirm the relationship deletion

Mapping the volume on the new storage system and resume host I/O

At this point the volume(s) are in the cloud. In our use case we were not moving hosts to the cloud but that may be a preference. Whichever route you choose, you can now connect the host to the volumes by mapping the volumes to resume host I/O.

What we have described is just an example of how to configure a GMCV relationship from an on-premises solution to migrate to an IBM Spectrum Virtualize for Public Cloud on AWS solution. The previous steps were completed by using the GUI, but they can also be done by using the CLI.

For more information around hybrid multicloud implementations, see the following publications:

- ▶ *Implementing IBM Spectrum Virtualize for Public Cloud Version 8.3*, REDP-5466
- ▶ *IBM Spectrum Virtualize for Public Cloud on AWS Implementation Guide*, REDP-5534



Host-based migrations in Microsoft Windows

This chapter addresses migrating data from an existing storage source volume to a FlashSystem 9100 target volume using Microsoft Windows Logical Disk Manager (LDM) mirroring and Microsoft Storage Spaces.

This chapter contains the following sections:

- ▶ Logical Disk Manager and Microsoft Storage Spaces
- ▶ Preparation and planning
- ▶ Migrating data using Windows Logical Disk Manager
- ▶ Migration with Microsoft Storage Spaces

9.1 Logical Disk Manager and Microsoft Storage Spaces

The Microsoft Windows platform supports two volume management applications, the Logical Disk Manager (LDM) initially introduced with Microsoft Windows NT, and Storage Spaces made available with Microsoft Windows 10.

Both volume management applications create a layer of virtualization within the operating system and provide the following basic functions:

- ▶ Extend logical volumes across several physical disks
- ▶ Stripe data across several physical disks to improve performance
- ▶ Mirror data for high availability and migration
- ▶ Create Software RAID Volumes

Volumes provided by a storage system are presented to the Logical Disk Manager and Storage Spaces as physical disks.

Logical Disk Manager

The normal process for a migration in LDM is to set up a mirror of the data on the existing volumes in the current storage system to new volumes in the replacement storage system. After the volumes are synchronized, they can be split to cut-over to the new storage system. The volume management software provides commands that can be used to automate the process.

Logical Disk Manager mirroring is an effective way to get data from one storage volume to another. The process is straightforward. You have a volume on an existing storage system (data source) with data that you want to relocate or migrate to a volume on a new storage system (data target).

The process for LDM mirror involves the following steps:

1. Configure the destination volume on the new storage system. This volume will be presented as a physical disk to LDM
2. Have the operating system recognize that new physical disk
3. Use the Windows LDM to set the new physical disk online
4. Establish a mirror between the source and destination volumes
5. Synchronize the mirror
6. Split the mirror

The data is now on the destination volume, on the FlashSystem 9100.

Storage Spaces

The normal process for a migration using Storage Spaces is to add the new physical disk to the existing storage pool and retire the original physical disk, then start a repair of the virtual Disk. You wait until the Virtual Disk is repaired and then remove the original physical disk.

The process for Storage Spaces involves these steps:

1. Configure the destination Volume, in this case on the FlashSystem 9100.
2. Have the operating system add the new Physical Disk to the same Pool as the source Virtual Disk belongs to.
3. Use the windows PowerShell to retire the source Physical Disk.
4. Use the windows PowerShell to repair the source Virtual Disk.

5. Wait until the repair process completes.
6. Use the Storage Spaces GUI or PowerShell to remove the source Physical Disk from the Pool.

The data is now on the destination volume.

Considerations

The application using the source logical volume or file system does not need to be stopped during the LDM mirror migration process or the Storage Spaces Physical Disk replacement process. Either process is nondisruptive.

The drawback in using LDM or Storage Spaces as the migration method is that both require substantial system administration, intervention, and attention. Production data is manipulated while production is running, and so migration requires host cycles.

Important: If you are planning to use LDM or Storage Spaces functions to migrate data, be careful with limitations. One essential limitation is that neither of them can be used for migrations with Microsoft Failover Clustering. Testing is also essential, if you have limited prior experience with LDM or Storage Spaces.

9.2 Preparation and planning

As highlighted in 1.4, “Data migration techniques” on page 16, any migration project must start with a preparation and planning phase. Create a diagram of the overall environment and create a list of all the components in your environment. The diagram for our use case is shown in Figure 9-1 on page 146, depicting a migration from FlashSystem A9000R to FlashSystem 9100.

Check the compatibility of the following components in your environment:

► Prepare the host servers

Verify that hosts with application volumes that need to be migrated can connect to the existing storage (source) and the target FlashSystem 9100 for the appropriate communication protocol in use in your environment.

Check for compatibility of HBAs and protocol in use. See the [IBM System Storage Interoperation Center \(SSIC\)](#).

- Check for the proper multipath IO feature. The Microsoft Device Specific Module (MSDSM) should now be used.

Note: The IBM Subsystem Device Driver Device Specific Module (SDDDSM) will be out of support as of June 2020. See the [IBM Support page](#)

Therefore, all Windows installations should now use the Microsoft Device Specific Module (MSDSM) which is included with the standard Windows operating system.

See the following for instructions:

- For [installing MSDSM](#)
- For [migrating from SDDDSM to MSDSM](#)

- Verify the connectivity to the host, for the protocol in use:
 - if using FC: Zone the host HBAs to the FS9100. A best practice is to use is four paths.
 - if using iSCSI or FCoE: check Ethernet connectivity.
 - if using SAS: check connectivity.
- ▶ **Prepare the existing storage system**
 - Verify that the existing storage system is in optimal condition
 - Check that all connections, ports, arrays, and volumes have a good status
 - Check firmware levels, and upgrade if necessary
 - Discover and document the logical and physical layout of the storage systems and volume mappings to the hosts.

9.3 Migrating data using Windows Logical Disk Manager

Dynamic disks were first introduced with the Windows 2000 server and provide features that basic disks do not. One of these features is the ability to create fault-tolerant volumes. This section addresses how to create a mirror using the Logical Disk Manager(LDM).

9.3.1 Windows migration use case

The environment for our sample use case is shown in Figure 9-1, and consists of migrating volumes attached to a Windows Server 2019 host, from FlashSystem A9000R to FlashSystem 9100.

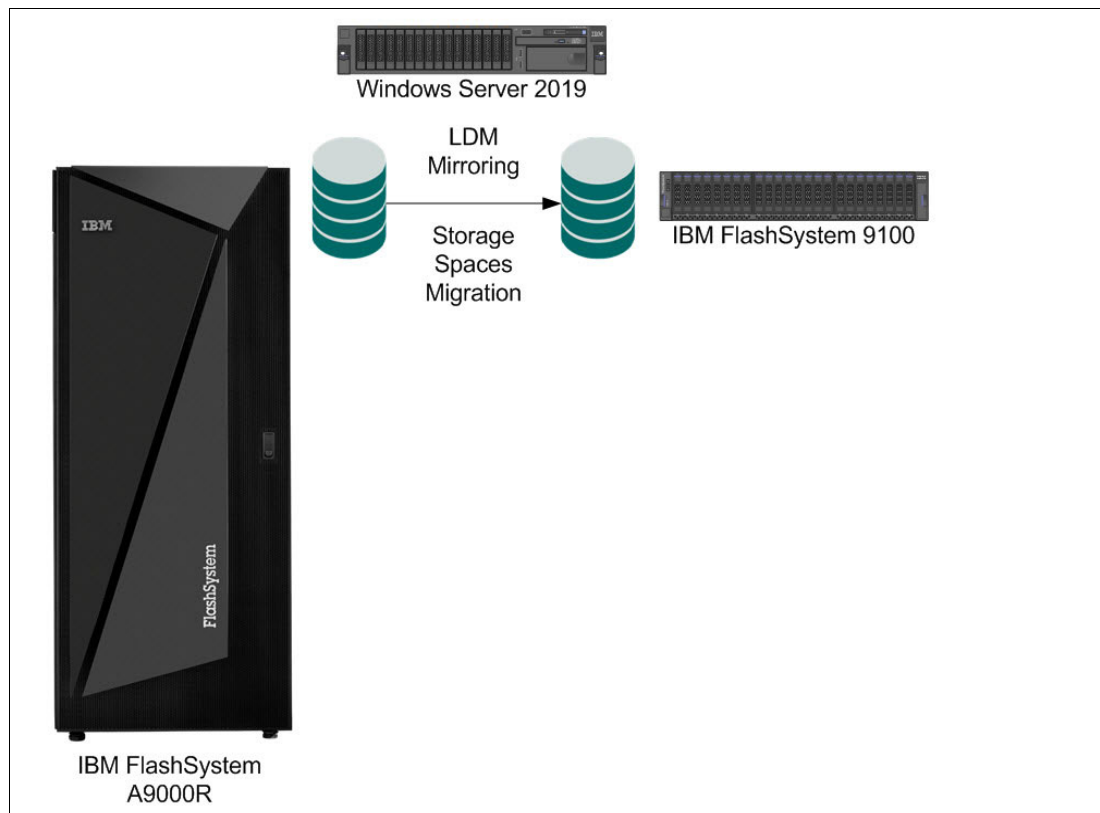
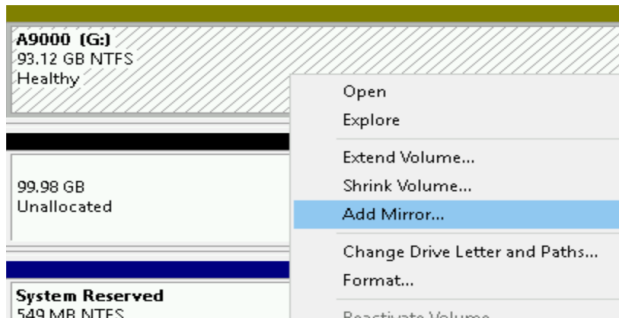


Figure 9-1 Microsoft Windows use case environment

9.3.2 Migration using Windows LDM mirroring

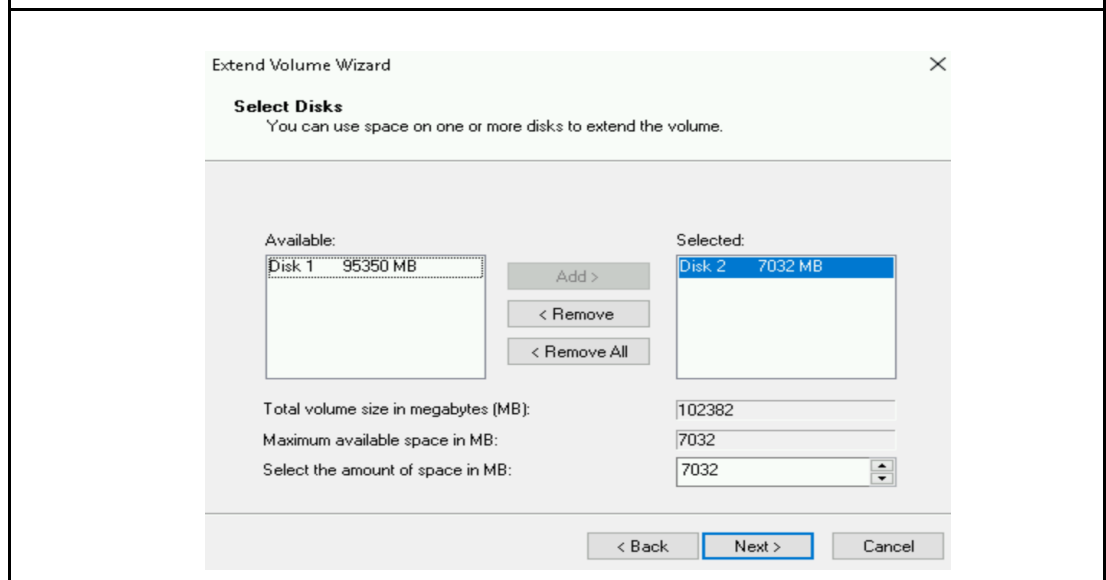

Table 9-1 shows the migration steps when using Logical Disk Manager (LDM) mirroring.

Table 9-1 Migration procedure using Windows LDM

Process	Commands	Explanation
Step 1. Identify the source (A9000R) volumes on the host server to be migrated.	<ol style="list-style-type: none"> 1. Right-click This PC. 2. Click Manage. 3. Click Disk Management. 	This identification is done on the Windows host server.
Note: Disk manager can also be opened with diskmgmt.msc		
Step 2. Assign the FS9100 volumes to the Windows host server and discover the new volumes.	<ol style="list-style-type: none"> 1. Map new volume in FS9100 GUI to win host 2. Click This PC. 3. Click Manage. 4. Click Disk Management. 5. Click Action from the toolbar. 6. Click Rescan disks. 	<p>Step 1 is done on the FS9100</p> <p>Steps 2-6 are performed on the Windows host server.</p>
Step 3. Bring the FS9100 target disks online.	<ol style="list-style-type: none"> 1. Right-click a new disk 2. Click Online. 	This step is done on the Windows host server.
Step 4. Identify the FS9100 LUN IDs.	<ol style="list-style-type: none"> 1. Right-click the Online disk and choose Properties 	There you can see IBM 2145 (SAN Volume Controller machine type), which is your FS9100
Step 5. Initiate the mirroring process.	<ol style="list-style-type: none"> 1. Right-click the source volume (x:). 2. Click Add Mirror. 3. Select the chosen disk. 4. Click Add Mirror. 	The synchronization of the source and target volumes is automatic.
Note: To be able to add mirror disk needs to be Dynamic, but the add mirror process automatically takes care of the conversion.		
		
Step 6. Verify that the volumes are copied and synced.	Visual check	Look for a disk that is functioning correctly.

Process	Commands	Explanation
<p>Note 1: Before disks are fully synchronized a warning sign is shown. Note 2: Because it is a mirror it also has the same disk label, A9000 Vol.</p>		
<p>Step 7. Remove the mirror.</p>	<ol style="list-style-type: none"> 1. Right-click the source disk. 2. Click Remove Mirror. 3. Select the source disk to remove. 4. Click Remove Mirror. 5. Click Yes. 6. Right-click the selected volume. 7. Click Remove Mirror. 8. Click Yes. 	<p>Make sure that you are removing the source disk from the mirror, not the FS9100 target disk.</p>
<p>Step 8. Verify that the mirror is removed.</p>	<p>Visual check</p>	<p>Check that the disk is now removed.</p>
<p>Step 9. Change the name of the new target Volume to FS9100.</p>	<ol style="list-style-type: none"> 1. Right-click the volume. 2. Click Properties. 3. Enter the volume name under the General tab. 	<p>Change the name of the new target disk to something meaningful.</p>

Process	Commands	Explanation
Step 10. Repeat steps 5-9.	N/A	Repeat the mirroring and renaming process for the remaining source and target disks.
Step 11. If needed extend the FS9100 disks to their full capacity.	<ol style="list-style-type: none"> 1. Right-click the volume. 2. Click Extend Volume. 3. Click the target disk. 4. Click Next. 5. Click Finish. 	In the example, the target disks are larger than the source Volumes, so you must expand them to use the remaining capacity.



Step 12. Delete the source volumes from the Windows host server.	<ol style="list-style-type: none"> 1. Unassign the volume in the A9000R from the Windows host server. 2. Delete the zone in the fabric. 3. Disconnect the A9000R storage device from the fabric or host server. 	Follow the procedures for each component listed in the right column.
---	--	--

Process	Commands	Explanation
Step 13. Rescan the Windows host server.	<ol style="list-style-type: none"> 1. Right-click This PC. 2. Click Manage. 3. Click Disk Management. 4. Click Action from the toolbar. 5. Click Rescan disks. 	If a volume is displayed as missing, right-click the volumes and click Delete .
Step 14. Verify that the device definitions are removed.	Visual check	Check that source volumes are gone.

9.4 Migration with Microsoft Storage Spaces

The Microsoft Storage Spaces (MSS) was introduced with Microsoft Windows 10 and can also be used for migration from an existing storage volume to a new storage volume.

For a description of Storage Spaces and its basic functionality, see the following [Storage Space Overview](#) document from Microsoft.

In our test case we explore the movement of data from FlashSystem A9000R volume to a FlashSystem 9100 volume.

Open **Server Manager** → **File and Storage Services** → **Volumes** → **Storage Pools**, as seen in Figure 9-2. In our test environment, we have a Windows Server 2019 host with a Physical Disk from FlashSystem A9000R in a Storage Spaces Storage Pool. We also created a corresponding Virtual Disk, named `ITSO_Test_vDisk` in the same Storage Spaces Storage Pool, as seen in Figure 9-2.

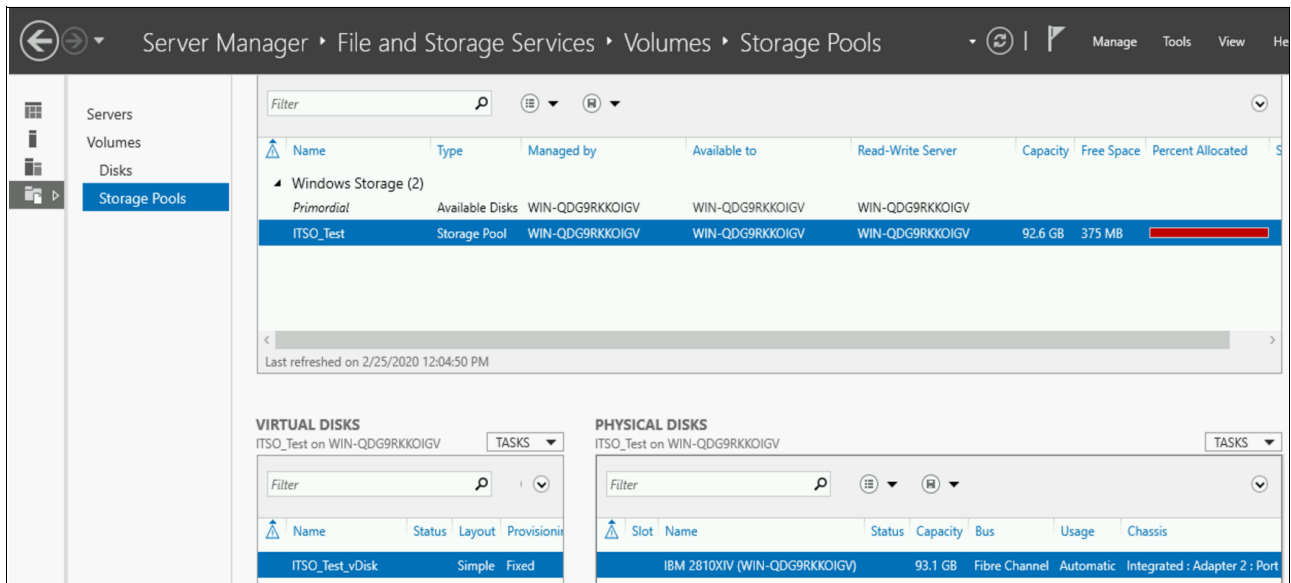


Figure 9-2 MSS GUI view

The data that resides on the Virtual Disk is shown in Figure 9-3.




Figure 9-3 Data on the FlashSystem A9000R disk

After the new storage disk from FlashSystem 9100 has been mapped to the Windows host it is time to add the Physical Disk to the Storage Pool. Click the **TASKS** pull-down in the PHYSICAL DISKS pane and select **Add Physical Disk**, as shown in Figure 9-4.

If no new disk can be seen, a rescan is needed. To rescan, click **TASKS** and **Rescan** above the Storage Pool section.

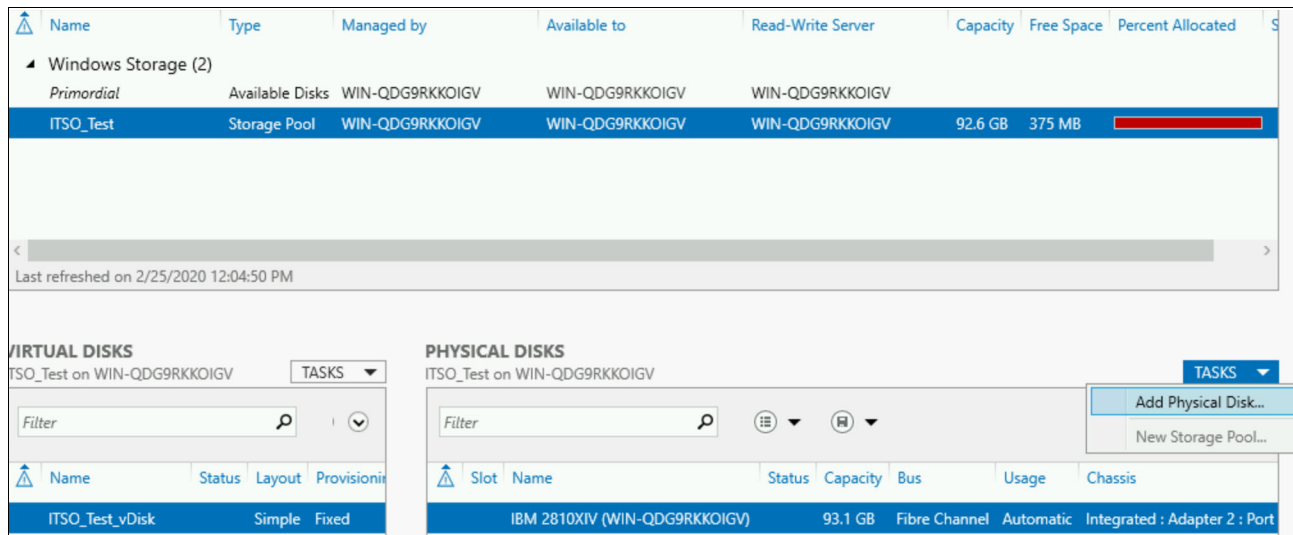


Figure 9-4 Add Physical Disk to the MS Storage Spaces Pool

Select the FlashSystem 9100 (FS9100) disk, which is recognized as IBM 2145 (2145 is the original machine type of the SAN Volume Controller), as shown in Figure 9-5.

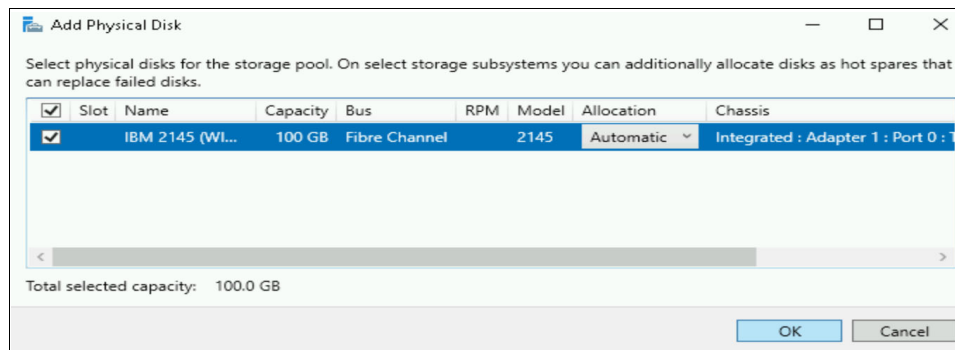


Figure 9-5 Selecting FlashSystem 9100 disk

Now, the MSS Pool has two Physical Disks attached, as shown in Figure 9-6.

The screenshot shows a window titled 'PHYSICAL DISKS' with the subtitle 'ITSO_Test on WIN-QDG9RKKOIGV'. It features a search filter, a list view icon, and a refresh icon. Below is a table with columns: Slot, Name, Status, Capacity, Bus, Usage, and Chassis. Two rows are visible, both with a blue background:

Slot	Name	Status	Capacity	Bus	Usage	Chassis
	IBM 2145 (WIN-QDG9RKKOIGV)		100 GB	Fibre Channel	Automatic	Integrated : Adapter 1 : Port
	IBM 2810XIV (WIN-QDG9RKKOIGV)		93.1 GB	Fibre Channel	Automatic	Integrated : Adapter 2 : Port

Figure 9-6 Two physical disks in MS Storage Spaces Pool

The next step is to move the data from the A9000R (IBM 2810XIV) Physical Disk to the new Physical Disk in the FS9100 (IBM 2145).

To start the movement of data from the A9000R disk to the FS9100 Disk, you need to retire the Physical Disk from A9000R and then repair the Virtual Disk.

The only way to retire a Physical disk from a Storage Pool in MSS is from the PowerShell. Open PowerShell as administrator and issue the commands as shown in Example 9-1.

Example 9-1 Retire the Physical Disk in PowerShell

```
# Show the PhysicalDisks
PS C:\Users\Administrator> Get-PhysicalDisk
DeviceId FriendlyName      SerialNumber      MediaType  CanPool
OperationalStatus HealthStatus Usage              Size
-----
-----
-----
2        IBM 2145                020420017e6cXX00 Unspecified True
OK        Healthy                Auto-Select      100 GB
1        IBM 2810XIV            7805C2300000000039213 Unspecified False
OK        Healthy                Auto-Select      93.13 GB
0        IBM ServerRAID M5210 005962020d5ee13e22c01a050db00506 Unspecified False
OK        Healthy                Auto-Select      557.86 GB

# Set the PhysicalDisk with DeviceId 1 to Retired
PS C:\Users\Administrator> Get-PhysicalDisk | where { $_.DeviceId -eq 1 } |
Set-PhysicalDisk -Usage Retired

# Show the PhysicalDisks
PS C:\Users\Administrator> Get-PhysicalDisk

DeviceId FriendlyName      SerialNumber      MediaType  CanPool
OperationalStatus HealthStatus Usage              Size
-----
-----
-----
2        IBM 2145                020420017e6cXX00 Unspecified False
OK        Healthy                Auto-Select      100 GB
1        IBM 2810XIV            7805C2300000000039213 Unspecified False
OK        Healthy                Retired          93.13 GB
0        IBM ServerRAID M5210 005962020d5ee13e22c01a050db00506 Unspecified False
OK        Healthy                Auto-Select      557.86 GB
```

Note: MSS assigns a FriendlyName which is dependent on the storage system it comes from. It can happen that it assigns the same FriendlyName to several Physical Disks. That makes it impossible to use the -FriendlyName parameter to Retire a Physical Disk.

However, if you only have one FriendlyName of the same kind, you can also use the -FriendlyName parameter with the **Set-PhysicalDisk** command, for example.

```
Set-PhysicalDisk -FriendlyName "IBM 2810XIV" -Usage Retired
```

In the GUI we can see that the Physical Disk from A9000R called "IBM 2810XIV" is in Retired status, as shown in Figure 9-7.

Slot	Name	Status	Capacity	Bus	Usage	Chassis
	IBM 2145 (WIN-QDG9RKKOIGV)		100 GB	Fibre Channel	Automatic	Integrated : Adapter 1 : Port
	IBM 2810XIV (WIN-QDG9RKKOIGV)	Retired	93.1 GB	Fibre Channel	Retired	Integrated : Adapter 2 : Port

Figure 9-7 Physical disk is retired

To start the movement of data to the empty Physical Disk, you need to repair the **Virtual Disk**. Go back to PowerShell again and repair the **Virtual Disk**.

As shown in Example 9-2, first list the Virtual Disks and then use the friendly name to start repair. You can follow the progress of the Repair on the top of the PowerShell screen.

Example 9-2 Repair the VirtualDisk

```
# Show the PhysicalDisks
PS C:\Users\Administrator> Get-VirtualDisk

FriendlyName      ResiliencySettingName FaultDomainRedundancy OperationalStatus
HealthStatus      Size FootprintOnPool StorageEfficiency
-----
-----
-----
ITS0_Test_vDISK Simple          0          OK
Healthy          92 GB          92 GB          100.00%
```

```
# Repair the VirtualDisk
PS C:\Users\Administrator> Repair-VirtualDisk -FriendlyName ITS0_Test_vDisk

# Monitor the Repair
#----> This will appear on top of screen ---->

Repair-VirtualDisk -FriendlyName ITS0_Test_vDisk
1/2 completed
[oooooooooooooooooooooooooooooooooooooooooooooooooooooooooooooooooooooooooooo
]
00:00:00 remaining.
Repairing virtual disk
Running
```

[00000000000000000000000000000000
]

After the repair process has finished and all data has been moved to the new Physical Disk, it is possible to remove the physical disk. In our test case, we remove the physical disk from the A9000R (2810XIV), as seen in Figure 9-8.

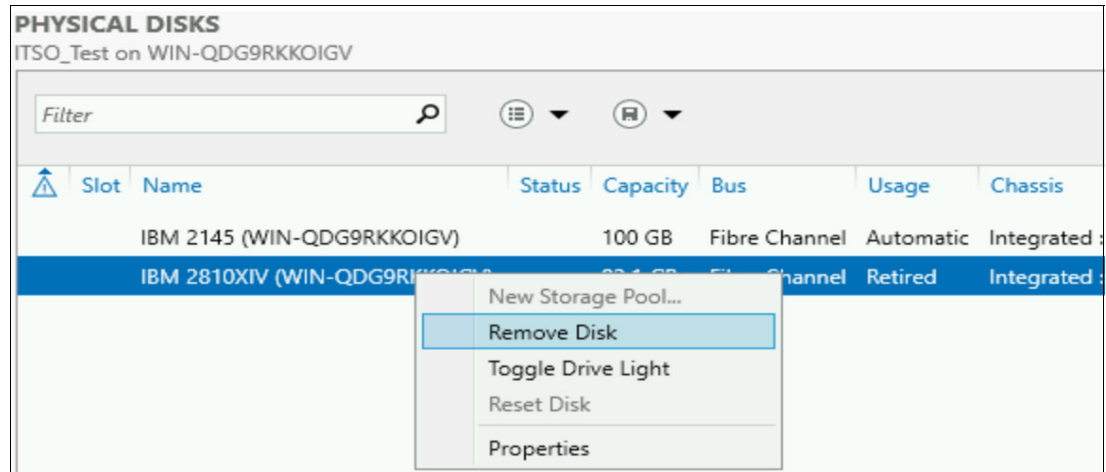


Figure 9-8 Remove Physical disk from Pool

Afterwards you can see that there only one Physical Disk left in the Pool, which is the FlashSystem 9100 disk (IBM 2145), as shown in Figure 9-9.

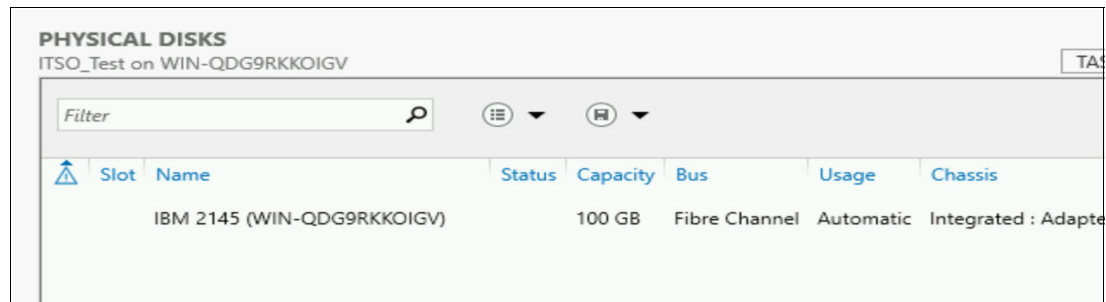


Figure 9-9 FS9100 Disk is the only member of the Pool

To make it more user friendly, it is a good idea to rename the disk label to some meaningful name, such as FS9100_Disk, as shown in Figure 9-10 on page 155, and then verify that data is readable.

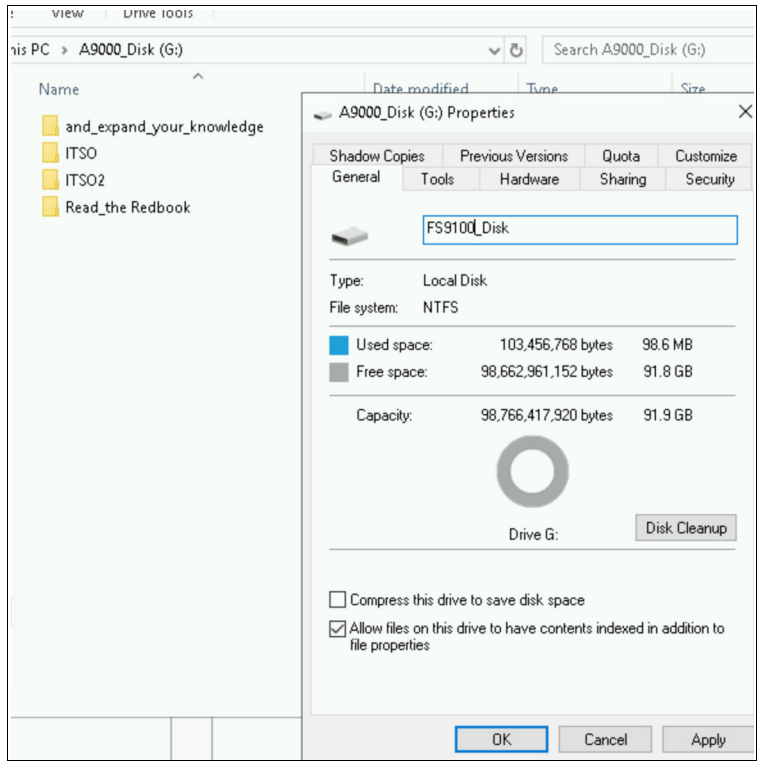


Figure 9-10 Rename of disk and verify data



Host-based migrations in AIX

This chapter addresses migrating data from a disk storage source device to a Spectrum Virtualize target device using the IBM AIX Logical Volume Manager (LVM)-based mirroring (split mirroring).

This chapter contains the following sections:

- ▶ Concepts of Logical Volume Manager
- ▶ Preparation and planning
- ▶ Data migration using AIX LVM mirroring
- ▶ PowerHA Storage Migration Considerations

10.1 Concepts of Logical Volume Manager

Some form of logical volume management is available for most open systems. For AIX and Linux, it is called Logical Volume Manager (LVM). For Solaris and HP-UX, it can be known as LVM or VxVM. The LVM creates a layer of virtualization within the operating system.

Every volume management software provides the following basic functions:

- ▶ Extend logical volumes across several physical disks
- ▶ Stripe data across several physical disks to improve performance
- ▶ Mirror data for high availability and migration

The LUNs provided by the storage system to AIX are presented to the LVM as physical SCSI disks.

The normal process for a migration is to set up a mirror of the data on the old disks to the new disks. You wait until the disks are synchronized, and then split them at the cut-over time. Some volume management software provides commands that automate this process.

Disk mirroring was first mentioned in a 1978 patent awarded to Norman Ken Ouchi of the IBM Corporation. Logical Volume mirroring came into general use in UNIX operating systems around the mid 1990s. At that time, it was used primarily for data protection. However, the “split mirror” technique of moving data from one disk to another became a handy tool for the system administrator.

A lot has changed over the years, but LVM mirroring is still an effective way to get data from one disk to another. The process is straightforward. You have one or more logical volumes in a volume group on one or more disks (data source) with data that you want to relocate or migrate to disks in another storage system (data destination). The process involves these steps:

1. Configure the destination LUNs on the replacement storage system
2. Have the operating system recognize the LUNs
3. Extend the AIX volume group with new LUNs
4. Mirror the volume group with the included logical volumes
5. Synchronize the mirror.
6. Split the mirror
7. The data is now on the destination storage system.

The application using the source logical volume or file system does not need to be stopped during the logical volume mirror migration process. The process is nondisruptive if the operating system allows you to add and remove devices dynamically. However, typically you want the application to use the data on the destination logical volume after the migration. Therefore, the application must be quiesced and reconfigured to use the destination logical volume.

A major advantage of using the volume management software for data migration is that it also allows for various types of consolidation. This consolidation is possible thanks to the virtualization nature of volume management software.

The major disadvantage of the volume management software mirroring method is that it requires substantial system administration, intervention, and attention. Production data is manipulated while production is running, and so migration requires host cycles.

Important: If you are planning to use volume management software functions to migrate data, be careful with limitations. These limitations include the total number of physical disks in the same volume group. In addition, if you are consolidating volumes with different sizes, check the procedures to see whether consolidation is possible and ensure enough free space is available on the new storage system. Take into consideration any changes to compression or deduplication you might have been getting on the old storage system.

10.2 Preparation and planning

As highlighted in 1.4, “Data migration techniques” on page 16, any migration project must start with a preparation and planning phase.

Create a diagram of the overall environment and create a list of all the components in your environment. Check the compatibility of the following components in your environment, to ensure they are supported System Storage Interoperation Center (SSIC) at: <http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>.

► Prepare the host servers

Verify that both the host and host HBAs, CNAs or Ethernet adapter can connect to the existing storage system (source) and the new storage system (destination).

- Check vendor and model type.
- Check the firmware level and upgrade if necessary. Supported firmware levels can be available in SSIC.
- Determine the number and size of each source volume (LUN) on each host.
- Determine how the LUN is spread across multiple disks for performance considerations.
- Determine whether the existing multipath is running AIXPCM. Since IBM Spectrum Virtualize version 7.6, IBM recommends the use of native AIXPCM for AIX systems instead of SDDPCM. SDDPCM is not supported with Spectrum Virtualize 8.3.0 or later. If you are deploying a new AIX system, use AIXPCM. If you are migrating an existing AIX host which runs SDDPCM to the Spectrum Virtualize system, our recommendation is to migrate from SDDPCM to AIXPCM whenever it is possible. For more information on how to migrate SDDPCM to AIXPCM, see the following links:
 - The recommended Multi-path Driver to use on IBM AIX and VIOS:
<https://www.ibm.com/support/pages/node/697363>
 - How To Migrate SDDPCM to AIXPCM:
<https://www.ibm.com/support/pages/node/698075>

► Prepare the IBM Spectrum Virtualize system and volumes.

- Make sure that the storage system is running the recommended code level or a higher code level. More details are available at:
<https://www.ibm.com/support/pages/node/690527>
- Check if enough usable capacity is available to do the migration. Check if the source storage system is using data reduction features and what capacity is used. You might also need Data Reduction Pools and data reduction features such as thin-provisioning, compression and deduplication. If your storage system is not showing the Data Reduction ratio or not using Data Reduction it might be necessary to run the [Data Reduction Tool \(DRET\)](#).

- Create the needed volumes.
- Map the volumes to the AIX host.
- ▶ Prepare the SAN or Ethernet connections.
 - Depending on connection protocols in use in your environment, check FC, iSCSI, iSER, FCoE, or SAS compatibility for connection to both the existing storage system and the new Spectrum Virtualize system. See the Interoperability matrix available at: <https://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>
 - Verify that you have enough free ports to connect the new Spectrum Virtualize system to the fabric. Add more if needed.
 - Verify the cabling.
 - Order the correct number of cables.
 - Connect the storage system to the fabric or direct to the host.
 - Provide cabling diagrams to the cabling vendor or the person responsible for the cabling.
 - Verify the zoning.
 - Zone the host HBAs to the Spectrum Virtualize system. The optimum number of paths from each host to one IBM Spectrum Virtualize volume is four paths, even though the maximum supported is eight paths.

For more details see:

IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines, SG24-7521-06.
- ▶ Prepare the existing storage system being replaced.
 - Verify that the existing storage system is in optimal condition.
 - Check that all connections, ports, arrays, and LUNs are functioning properly.
 - Check firmware levels (upgrade if necessary).
 - Discover and document the logical and physical layout of the arrays and LUN mappings to the hosts.

10.3 Data migration using AIX LVM mirroring

This section outlines the steps for using the AIX Logical Volume Manager (LVM) mirroring to migrate data from a FlashSystem 900 to a FlashSystem 9100 with the `mirrorvg` command.

Note: You can use the `mk1vcopy` commands on each logical volume if you do not want to migrate the entire volume group.

You can connect the new storage system in parallel with the existing storage system. Assign the LUNs from the new storage system to the same compatible host server adapters that the current LUNs are assigned to. In this way, you can mirror LUNs or volumes on the existing storage system to the new storage system using software (LVM) techniques.

After the volumes are mirrored and synced, break the mirror from the existing storage system, remove the old volume and storage definitions and remove the old storage system.

Prepare the new storage system with the correct Array to LUN spread and LUN size considerations as it applies to the current setup of the existing storage system. For more information about FlashSystem 9200 or 9100 preparation, see *IBM FlashSystem 9200 and 9100 Best Practices and Performance Guidelines*, SG24-8448.

10.3.1 IBM FlashSystem 900 to IBM FlashSystem 9100 migration

The use case presented is about migrating from IBM FlashSystem 900 to IBM FlashSystem 9100, using AIX LVM.

The physical test environment consists of the following components:

- ▶ IBM Power System Model 8284-22A LPAR running AIX level 7.2, TL 04, SP 01.
 - AIXPCM (Native MPIO multipathing)
 - Virtual Fibre Channel Adapter (NPIV)
- ▶ IBM FlashSystem 900 running 1.5.2.6 code.
- ▶ IBM FlashSystem 9100 running 8.3.1 code.
- ▶ IBM Switch 2498-F48, Fabric OS firmware version 7.4.2c

10.3.2 High-level migration plan using the AIX LVM mirrorvg

The high-level steps for a migration based on the `mirrorvg` command are shown in Table 10-1. The table illustrates a FlashSystem 900 to FlashSystem 9100 migration example. You can apply the same process to other open system storage platforms.

Table 10-1 Example migration procedure using mirrorvg

Process	Commands	Explanation
Step 1. Identify the FlashSystem 900 source disks.	<code>lsdev -Cc disk</code> <code>lspv</code>	Gather information about the AIX host, such as the number of FlashSystem 900 LUNs and sizes.
Step 2. Create the FlashSystem 9100 Volumes and assign to the host.		In FlashSystem 9100 GUI, create Volumes and map to host definition.
Step 3. Discover the 9100 LUNs.	<code>cfgmgr</code>	On the AIX host, discover the newly assigned LUNs.
Step 4. Identify the sizes of the FS9100 target LUNs.	<code>getconf DISK_SIZE /dev/hdisk#</code>	Where # is the number of the new hdisk from the FS9100.
Step 5. Move the FS9100 LUNs into the appropriate VGs.	<code>extendvg vg-name hdisk#</code>	Where <code>vg_name</code> is the name of the VG and # is the number of the hdisk.
Step 6. Verify that the FS9100 LUNs are added to the VG.	<code>lsvg -p vg_name</code>	Where <code>vg_name</code> is the name of the VG.
Step 7. Identify the logical volumes (LVs) to migrate.	<code>lsvg -l vg_name</code>	Where <code>vg_name</code> is the name of the VG.
Step 8. Mirror the VG data from the FS900 source LUNs to the FS9100 target LUNs.	<code>mirrorvg vg_name</code>	Where <code>vg_name</code> is the name of the VG.

Process	Commands	Explanation
Step 9. Verify that the LUNs are copied.	<code>lsvg -p vg_name</code> <code>lspv -l hdisk#</code>	Where <code>vg_name</code> is the name of the VG and <code>#</code> is the number of the hdisk.
Step 10. Unmirror VG to remove old FS900 HDisks.	<code>unmirrorvg vg_name hdisk#</code>	Where <code>vg_name</code> is the name of the VG and <code>#</code> is the number of the hdisk on the FS900.
Step 11. Remove the FS900 source LUNs from the VGs and verify that the source FS900 LUNs are removed from the VG.	<code>reducevg vg_name hdisk#</code> <code>lsvg -p vg_name</code>	Where <code>vg_name</code> is the name of the VG and <code>#</code> is the number of the hdisk.
Step 12. Delete the device definitions from the host ODM.	<code>rmdev -dl hdisk#</code>	Where <code>#</code> is the number of the appropriate hdisk.
Step 13. Verify that the device definitions are removed	<code>lsdev -Cc disk</code>	Check that there are no defined disks remaining from the FS900 source storage.
Step 14. Remove the source zone definition from the switch fabric.	n/a	Removes data path between AIX host and disk subsystem
Step 15. In the FS900 GUI, unassign the LUNs from the host server.	8.	This sequence is done from the FS900 GUI.

10.3.3 Migration back-out considerations

The data migration process described removes the VG copy. After the VG copy is removed, you can no longer use that copy. The only accessible copy of the VG is on the new set of LUNs. If, for whatever reason, you decide to go back to the old set of LUNs, you need to repeat the migration process in the reverse direction. Depending on the amount of data, this process might take a significant amount of time. The downside of this method is that at the end of the migration, you have only one physical copy of the data. During the planning phase, before the migration starts, you need to decide whether the risk associated with having only one copy of data is acceptable.

The alternative approach allows you to preserve an extra copy of the data. The `splitvg` command allows you to split a copy from the original VG and creates a VG from it. To maintain consistency of data, however, this command requires the application to be stopped and the file system to be unmounted. If you decide that having an extra copy is more important than doing the migration non-disruptively, use `splitvg`.

The `splitvg` command allows you to set which mirror to split with the `-c` option. If you need to go back to the original mirrored set of LUNs you can use the `joinvg` command again to join the volume group again. This process does not require any data to be copied.

10.3.4 Data Migration with LVM mirroring example

To migrate from the FlashSystem 900 (source) to FlashSystem 9100 (target) we need to create the zoning to the virtual worldwide port names (WWPNs) of FlashSystem 9100 on the switch, create a host and a volume on the FlashSystem 9100 and map it to the AIX host.

Run `cfgmgr` on the AIX so that HBA ports log in to FlashSystem 9100 ports.

Open the FlashSystem 9100 GUI by entering `https://<cluster-ip-address>` in a browser and sign in. Select **Hosts** → **Hosts** and click **Add Host**, the Add Host wizard opens, complete name and select the worldwide port names (WWPNs); then click **Add** as illustrated in Figure 10-1.

Add Host

i Because NPIV is enabled on this system, host traffic is only allowed over the storage system's virtual ports. Ensure that SAN zoning allows connectivity between virtual ports and the host.

Required Fields

Name:

Host connections:

Host port (WWPN):

Host port (WWPN):

Optional Fields

Host type:

I/O groups:

Host cluster:

Figure 10-1 Add AIX host on FlashSystem 9100

To create and map a volume to the AIX host, select **Volumes** → **Volumes** from the menu, click **Create Volumes**, select Pool and choose quantity and size, click **Create and Map** as shown in Figure 10-2.

The screenshot shows the 'Create Volumes' interface with the following details:

- Navigation:** 'Basic' is selected, with 'Mirrored' and 'Custom' as alternatives.
- Instructions:** 'Create a preset volume with all the basic features.'
- Pool:** A dropdown menu is set to 'Pool0'.
- Capacity Details:** A progress bar indicates 'Total 66.54 TiB'.
- Volume Details:**
 - Quantity:** A spinner control is set to '1'.
 - Capacity:** A dropdown menu is set to '100 GiB'.
 - Name:** A text input field contains 'ITSO_AIX_Data'.
 - Capacity savings:** A dropdown menu is set to 'None', and the 'Deduplicated' checkbox is unchecked.
- Actions:** A link labeled 'Define another volume' is present.
- I/O group:** A dropdown menu is set to 'Automatic'.
- Summary:** A small table shows:

Summary
1 volume
Volume name: ITSO_AIX_Data
- Footer:** Includes a 'Need Help' link and three buttons: 'Cancel', 'Create and Map', and 'Create'.

Figure 10-2 Create AIX volume on FlashSystem 9100

In the Create Mapping window, select the host *ITSO_AIX*. Click **Next** to continue as depicted in Figure 10-3.

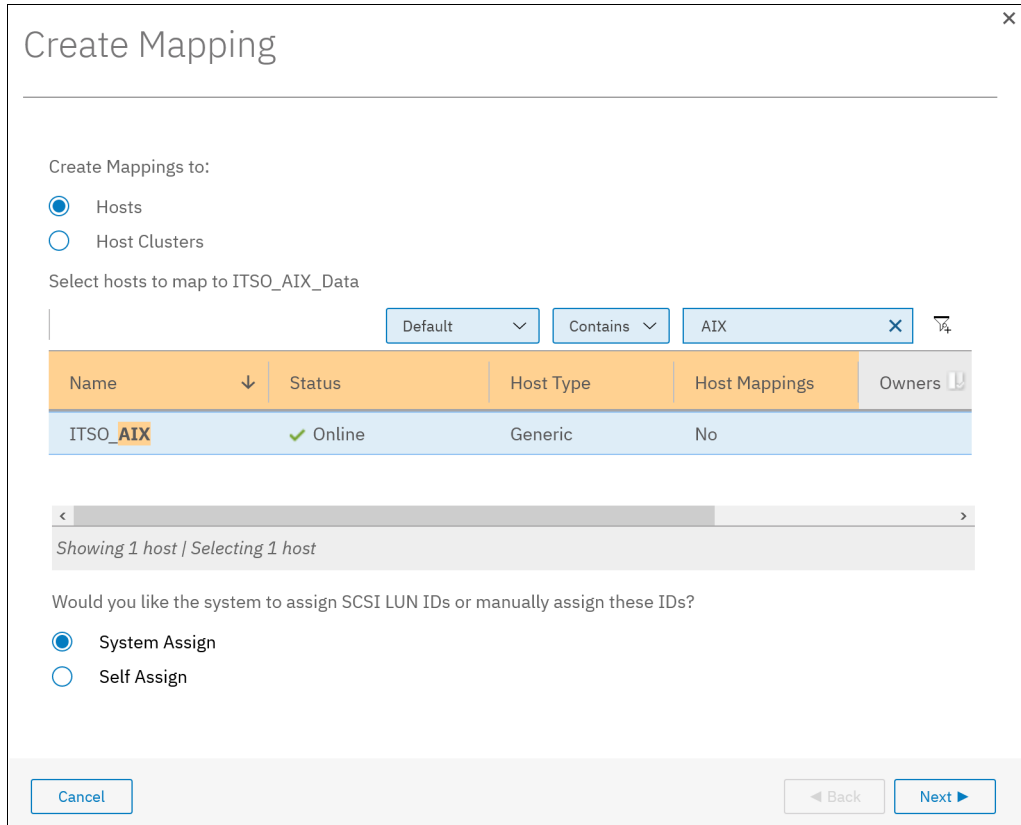


Figure 10-3 Map AIX host and volume on 9200

In the Summary window, click **Map Volumes** to confirm your host and volume mapping, as shown in Figure 10-4.

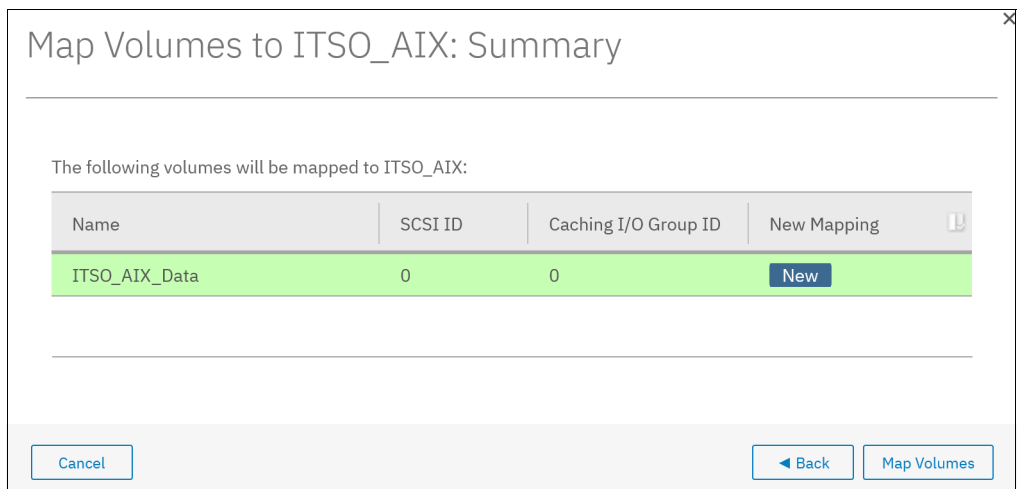


Figure 10-4 Summary window for host and volume mapping

You can now see the new volume on the FlashSystem 9100 formatting in the background, as illustrated in Figure 10-5.

Name	State	Synchronized	Pool	Protocol Type	UID	Host Mappings	Capacity
ITSO_AIX_Data	✓ Online (formatting)		Pool0	SCSI	60050768108005F9B0000000000000...	Yes	100.00 GiB

Figure 10-5 New Volume formatting in the background

Back on your AIX host, you can run the `cfgmgr` command again to refresh the configuration for AIX to see the FlashSystem 9100 disk. Show the output of all present disks with `lspv` where you will see an HDisk with no mappings, as shown in Example 10-1 for `hdisk2`.

Example 10-1 The `lspv` and `cfgmgr` commands to show new `hdisk2` FlashSystem 9100 disk on AIX

```
# lspv
hdisk0          00fa00f0ff5158ab          rootvg          active
hdisk1          00fa00f0793833b6          datavg          active

# cfgmgr

# lspv
hdisk0          00fa00f0ff5158ab          rootvg          active
hdisk1          00fa00f0793833b6          datavg          active
hdisk2          none                       None            active
```

Extend your volume group, with the `extendvg` command, to include this new HDisk from the FlashSystem 9100. Notice the before and after volume group listing information, now showing the additional physical volume (Total PVs) as highlighted in Example 10-2.

Example 10-2 Add new FlashSystem 9100 HDisk to existing volume group

```
# lsvg datavg
VOLUME GROUP:      datavg          VG IDENTIFIER:
00fa00f000004c0000000170793833e4
VG STATE:          active          PP SIZE:          128 megabyte(s)
VG PERMISSION:    read/write      TOTAL PPs:        799 (102272
megabytes)
MAX LVs:          256          FREE PPs:         6(768 megabytes)
LVs:              2          USED PPs:         793(101504 megabytes)
OPEN LVs:         2          QUORUM:           2 (Enabled)
TOTAL PVs:        1          VG DESCRIPTORS:  2
STALE PVs:        0          STALE PPs:        0
ACTIVE PVs:       1          AUTO ON:          yes
MAX PPs per VG:  32512
MAX PPs per PV:  1016          MAX PVs:          32
LTG size (Dynamic): 512 kilobyte(s)
HOT SPARE:        no          AUTO SYNC:        no
PV RESTRICTION:  none          BB POLICY:        relocatable
DISK BLOCK SIZE: 512          INFINITE RETRY:  no
FS SYNC OPTION:  no          CRITICAL VG:     no
CRITICAL PVs:    no
```

```
# lsvg -l datavg
datavg:
LV NAME          TYPE          LPs          PPs          PVs          LV STATE          MOUNT POINT
loglv00          jfs2log      1            1            1            open/syncd        N/A
fslv00           jfs2         792          792          1            open/syncd        /data
```



```
# extendvg datavg hdisk2
0516-1254 extendvg: Changing the PVID in the ODM.
```

```
# lsvg datavg
VOLUME GROUP:      datavg                VG IDENTIFIER:
00fa00f000004c0000000170793833e4
VG STATE:          active                 PP SIZE:        128 megabyte(s)
VG PERMISSION:    read/write            TOTAL PPs:      1598 (204544
megabytes)
MAX LVs:          256                   FREE PPs:       805(103040 megabytes)
LVs:              2                     USED PPs:       793(101504 megabytes)
OPEN LVs:         2                     QUORUM:         2 (Enabled)
TOTAL PVs:      2                   VG DESCRIPTORS: 3
STALE PVs:        0                     STALE PPs:      0
ACTIVE PVs:       2                     AUTO ON:        yes
MAX PPs per VG:   32512
MAX PPs per PV:   1016                   MAX PVs:        32
LTG size (Dynamic): 512 kilobyte(s)     AUTO SYNC:      no
HOT SPARE:        no                     BB POLICY:      relocatable
PV RESTRICTION:   none                   INFINITE RETRY: no
DISK BLOCK SIZE: 512                     CRITICAL VG:    no
FS SYNC OPTION:   no                     CRITICAL PVs:   no
```

```
# lsvg -l datavg
datavg:
LV NAME          TYPE      LPs    PPs    PVs  LV STATE  MOUNT POINT
loglv00          jfs2log  1      1      1    open/syncd  N/A
fs1v00           jfs2     792    792    1    open/syncd  /data
```

Mirror the volume group to start the migration process so data will exist on both your original source storage system and the FlashSystem 9100 storage system. You can use the **mirrorvg** command as shown in Example 10-3 or specify the **-S** option to run the sync in the background and return to the command line immediately.

Example 10-3 AIX mirrorvg command

```
# mirrorvg datavg
0516-1804 chvg: The quorum change takes effect immediately.
```

After the mirror sync has completed, you will be able to see the command output of your volume group as illustrated in Example 10-4 showing two physical volumes (PVs) and the logical volumes (LVs) in a open/syncd state.

Example 10-4 AIX volume group showing synced with 2 copies

```
# lsvg -l datavg
datavg:
LV NAME          TYPE      LPs    PPs    PVs  LV STATE  MOUNT POINT
loglv00          jfs2log  1      2      2    open/syncd  N/A
fs1v00           jfs2     792    1584   2    open/syncd  /data
```

You can now break the volume group mirror, with the **unmirrorvg** command as shown in Example 10-5 on page 168, specifying to remove the hdisk from the old storage system, keeping the FlashSystem 9100 volume only.

You can then remove this HDisk from the volume group using the **reducevg** command and remove the HDisk from the configuration with the **rmdev** command.

Example 10-5 Remove old storage mirror

```
# unmirrorvg datavg hdisk1
0516-1804 chvg: The quorum change takes effect immediately.

# lspv
hdisk0          00fa00f0ff5158ab          rootvg          active
hdisk1          00fa00f0793833b6          datavg          active
hdisk2          00fa00f07994adaa          datavg          active

# lsvg -l datavg
datavg:
LV NAME          TYPE          LPs          PPs          PVs          LV STATE          MOUNT POINT
loglv00          jfs2log       1            1            1            open/syncd        N/A
fslv00           jfs2          792          792          1            open/syncd        /data

# reducevg datavg hdisk1

# lspv
hdisk0          00fa00f0ff5158ab          rootvg          active
hdisk1          00fa00f0793833b6          None
hdisk2          00fa00f07994adaa          datavg          active

# rmdev -d1 hdisk1
hdisk1 deleted

# lspv
hdisk0          00fa00f0ff5158ab          rootvg          active
hdisk2          00fa00f07994adaa          datavg          active
```

10.4 PowerHA storage migration considerations

Naturally, customers would like to do storage migration in a PowerHA cluster without disruption to their operations. This is often possible, but not always. Explore the following sections discussing a few use cases for storage migrations with PowerHA environments. For additional information, see *IBM PowerHA SystemMirror V7.2.3 for IBM AIX and V7.22 for Linux*, SG24-8434.

This section discusses the following use cases:

- ▶ The new storage can be brought online along with the existing storage.
- ▶ The new storage cannot be brought online along with the existing storage.
- ▶ Migrating the cluster repository disk.

10.4.1 The new storage can be brought online along with the existing storage

This is the simplest use case. Correct use of systems facilities allows you to do nondisruptive storage migration. The basic idea is to add the new disks to an existing volume group (VG), create a new mirror of the VG on the new disks, and then remove the old disks from the VG. There are three cases to consider:

► **rootvg**

A new disk or disks can be added to rootvg with the **extendvg** command. A new mirror of rootvg can be made with the **mirrorvg** command. After completion, the existing disks can be evacuated with the **rmlvcopy** command, and those disks removed from the VG with the **reducevg** command. For rootvg's that contain only a single disk, **replacepv** can be a more convenient way of performing this operation.

There are two important steps that must be performed for rootvg after you complete this operation:

- a. Run **savebase** to update the boot image.
- b. Run **bootlist** to indicate that the new disks are to be used for the next boot operation.

The administrator must check the man pages for the **mirrorvg**, **rmlvcopy**, and **reducevg** to check the documented restrictions. These restrictions are not expected to apply to the common run of rootvg configurations, but a quick check in advance can avoid a problem later.

► **Shared VGs managed by PowerHA**

For this type of VG, it is important to use Cluster Single Point of Control (C-SPOC) operations to change the VG to maintain a consistent image, cluster wide. That is, use the C-SPOC SMIT panels to add new disks to the VG, create new mirrors of the logical volumes, remove existing mirrors, and remove existing disks. C-SPOC takes care of any needed **savebase** operations. Notice that when the new mirror is created, Logical Volume Manager (LVM) must copy all of the existing partitions to the new mirror. This process can be lengthy. For shared VGs with a single disk or a few disks, it is more convenient to use the C-SPOC disk replacement operation.

► **The repository disk, see 10.4.3, "Migrating the cluster repository disk" on page 170**

Although the process of migrating the repository to a new disk is conceptually the same as the two prior cases, LVM commands must never be used on caavg_private in an attempt to accomplish this. Such attempts can easily produce the following situation:

- LVM thinks caavg_private is on one disk.
- Cluster Aware AIX (CAA) thinks the repository is on another disk.

This can happen because CAA keeps information in the boot record, which is not part of the LVM image. Such situations require significant expertise to unravel in a nondisruptive manner. Most customers, when faced with the torturous series of steps required, prefer to take an outage, and do a scrub and rebuild.

The process of migrating the repository consists of first using the SMIT panels (or **clmgr** commands) to add a backup repository, then doing a *replace repository* operation. This makes the existing repository a backup. The last step is to remove the original repository disk. Because repository replacement updates the PowerHA HACMPsircol Object Data Manager (ODM), a verify and sync is necessary (depending on the software level). The way to check is to run **odmget HACMPcluster | grep handle** to see the handle. If the handle is 0, you know that a verify and sync are required.

If any backup repositories are defined, be sure to remove them and add the new disks.

10.4.2 The new storage cannot be brought online along with the existing storage

This use case can arise if the customer is actually changing the attachment mechanism for the disks, say from MultiPath I/O (MPIO) to VIOS. In this case, a restart is unavoidable. Here are some points to keep in mind to make the restart successful:

- ▶ rootvg

Such a replacement of the attachment mechanism can cause the disk names to change. Then, it might be necessary to use the Hardware Management Console (HMC) to modify the boot list so that it points to the new disk.

- ▶ Shared VGs managed by PowerHA

Because LVM finds disks by physical volume identifier (PVID), which is unaltered by the change of the attachment mechanism, this case requires no work.

- ▶ The repository disk, see 10.4.3, “Migrating the cluster repository disk” on page 170

Because CAA finds disks by UUID, and a change of attachment mechanism almost certainly changes the UUID, more work might be required here.

Note: On recent versions of CAA, it also tracks the PVID. If the output for the `odmget -q 'name=cluster0' CuAt` command shows a PVID stanza, you are fortunate. CAA finds the repository by PVID if a UUID search fails

CAA comes up, using the bootstrap repository, even if it cannot find the repository disk. However, if the **caavg_private** VG is not online, it is possible to bring it online by running `clusterconf -r <hdisk name>`. If that does not work, and the change of attachment mechanism changed the disk name, the **chrepos** command can be tried to replace the repository disk under the old name with its new name.

In the worst case, there is always the brute force solution: shut down PowerHA and CAA on all nodes, scrub the repository disk, and do a verify and sync to re-create the CAA cluster.

The important case that has not been discussed is the use of raw disks. For these, none of the statements in this appendix apply. You must consult the manufacturer of the disk attachment mechanism (for example, Oracle)

10.4.3 Migrating the cluster repository disk

The following procedure is valid for clusters that are PowerHA SystemMirror v7.2.0 and later.

1. Verify the cluster level on any node in the cluster by executing the **hallevel -s** command as shown in Example 10-6.

Example 10-6 Verify cluster level

```
TST[root@aixdc79p:/] # hallevel -s
7.2.1 SP2
```

The repository disk is the only disk in the **caavg_private** volume group (VG) and requires special procedures. You do *not* use Logical Volume Manager (LVM) on it. It is recommended that you run a verification on the cluster. Any errors must be addressed and corrected before you replace the repository disk.

2. To run cluster verification, which is a non-intrusive procedure, execute the `clmgr verify cluster` command. Confirm that the command completes normally as shown in Example 10-7.

Example 10-7 Cluster verification

```
Completed 70 percent of the verification checks
Completed 80 percent of the verification checks
Completed 90 percent of the verification checks
Completed 100 percent of the verification checks
Verification has completed normally.
```

Note: If errors are detected, you must correct them before you replace the repository disk.

3. The new repository disk must be the same size as the existing repository disk. In Example 10-8, both disks are 5 GB. You verify this fact by running the `bootinfo -s hdisk#` command.

Example 10-8 Verify the disk sizes

```
TST[root@aixdc285:/] # lspv
hdisk0 00f811ec5f0952f5 rootvg active
hdisk1 00f811ec68f206bd volgrp1 active
hdisk2 00f811ec63beda3a vgams01 active
hdisk4 00f811ec32a8dc6f caavg_private active
hdisk5 00f811eca2b2d20f vgshr02 concurrent
hdisk3 00f9be345a22c61e volgrp1 active
hdisk6 00f9be345a232b5a rootvg active
hdisk7 none None
hdisk8 00f9be345a249a2f None
hdisk9 00f9be345a24490d vgams01 active

TST[root@aixdc285:/] # bootinfo -s hdisk4
5120
TST[root@aixdc285:/] # bootinfo -s hdisk7
5120
```

4. Configure the old disk with respect to the new disk as follows:
In this case, **hdisk4** is the existing repository disk and **hdisk7** is the new one to replace it. The two repository disks must be configured the same and then synchronized with all nodes in the cluster. However, it is recommended that you do configuration of one node first with the `cfgmgr` command (in this example, on **aixdc285**). Then, manually run the following command to put a physical volume identifier (PVID) on the disk as shown in Example 10-9.

Example 10-9 Add volume identifier

```
chdev -l hdisk7 -a pv=yes

TST[root@aixdc285:/] # lspv |grep hdisk7
hdisk7 00f9be345a23c5db None
```

Note: Make a note of the PVID to use later in this procedure. The PVIDs of the old and new repository disks must match

5. Verify that the *reserve_policy* on the new repository disk is set to *no_reserve* by running this command as shown in Example 10-10.

Example 10-10 Verify reserve policy

```
TST[root@aixdc285:/] # lsattr -El hdisk7|grep reserve
reserve_policy no_reserve Reserve Policy True
```

If you do not see this setting, run the following command in Example 10-11 to update the setting.

Example 10-11 Updating the reserve policy

```
chdev -l hdisk7 -a reserve_policy=no_reserve
```

6. On the second node — **aixdc284** in this example — run the **cfgmgr** command to verify that a disk is configured with the same PVID.

The HDisk name does not have to, and often does not, match across the systems, as in this example. Yet, you can see in Example 10-12 that the PVID on the new disk matches the PVID on the old disk.

Example 10-12 Confirm volume identifier

```
TST[root@aixdc284:/] # lspv |grep c5db
hdisk3 00f9be345a23c5db None
```

7. Verify that the *reserve_policy* on the repository disk is set to *no_reserve* as shown in Example 10-13.

Example 10-13 Verify reserve policy

```
TST[root@aixdc284:/] # lsattr -El hdisk3|grep reserve
reserve_policy no_reserve Reserve Policy True
```

8. To swap the repository disk, execute the following command:

```
clmgr modify cluster REPOSITORY=hdisk#
```

where *hdisk#* is the name of a specific HDisk.

You can run this command on either node in the cluster. Be sure to specify the correct HDisk. For example, from node **aixdc284** and **hdisk3**, you run this command as shown in Example 10-14. Check that you selected the correct disk. Then, confirm this operation by entering **y** and pressing **Enter**. This process is completed in 1 to 2 minutes.

Example 10-14 Swap repository disk

```
TST[root@aixdc284:/] # clmgr modify cluster REPOSITORY=hdisk3
***Warning: this operation will destroy any information currently stored on
"hdisk3". Are you sure you want to proceed? (y/n) y
```

9. After completion, verify on both nodes that the new disk shows up with **caavg_private** as shown in Example 10-15.

Example 10-15 Verify both nodes status

```
TST[root@aixdc284:/] # lspv |grep caa
hdisk3 00f9be345a23c5db caavg_private active
```

```
TST[root@aixdc285:/] # lspv |grep caa
```

```
hdisk7 00f9be345a23c5db caavg_private active
```

10. **Migrations or swaps only:** Delete the original repository disk if your goal is to migrate from the original repository disk and not retain a backup repository disk.

In contrast to the stated purpose of this section, the default behavior in a repository disk swap is that the previous repository disk automatically becomes a *backup* repository disk. For a migration or a swap, this behavior is *not* desired. You must manually remove the backup repository disk from the cluster, and then synchronize the cluster again. One approach is by using the following steps:

- a. Run the following command to verify that the original repository disk (**hdisk4**, in this example) is still defined as a repository candidate disk.

Example 10-16 Verify repository disk

```
TST[root@aixdc284:/] # clmgr query repository
hdisk3 (00f9be345a23c5db)
hdisk4 (00f811ec32a8dc6f)
```

- b. When you see that the original repository disk is still listed, you run the command shown in Example 10-17 to remove it.

Example 10-17 Remove repository disk

```
/usr/es/sbin/cluster/utilities/clmgr -f delete repository 'hdisk4'
```

- c. Example 10-18 shows how to verify success by repeating the **clmgr query repository** command. Notice that the query output no longer shows **hdisk4**.

Example 10-18 Verify repository disk

```
TST[root@aixdc284:/] # clmgr query repository
hdisk3 (00f9be345a23c5db)
```

- d. Run the command shown in Example 10-19 to complete the removal process by synchronizing the cluster. Upon successful cluster synchronization, the repository disk migration or swap process is complete.

Example 10-19 Synchronize the cluster

```
TST[root@aixdc284:/] # clmgr sync cluster
```



Host-based migrations in Linux

This chapter addresses migrating data online from one storage system to another using Logical Volume Manager (LVM)-based live migration. This chapter contains the following sections:

- ▶ Concepts of Logical Volume Manager
- ▶ Preparation and planning
- ▶ Data migration using the LVM `pvmove` command

11.1 Concepts of Logical Volume Manager

Logical Volume Manager (LVM) is a dynamic volume management solution for Linux. It was created in 1998, and has since become the de-facto standard for managing storage in Linux, partially because of Red Hat Enterprise Linux (and its derivatives) using it by default for all installations of recent versions.

LVM provides the following basic functions:

- ▶ Extend logical volumes across several physical volumes
- ▶ Stripe data across several physical volumes to improve performance
- ▶ Mirror data for high availability
- ▶ Live migration of extents between physical volumes

LVM uses three concepts to manage storage: Physical Volumes, Logical Volumes, and Volume Groups. Physical Volumes, or PVs, are any block device that LVM uses as its back-end storage, such as a hard disk or a volume on a storage array. PVs are grouped into pools called Volume Groups or VGs, which can then be subdivided into Logical Volumes, or LVs.

Most LVs are “Linear”, meaning they are normal volumes that are allocated linearly in free space on PVs. They can also be created as striped, mirrored, or striped with parity (RAID5/6). The specific configuration doesn’t matter for the purpose of migrating all of one PV to a different PV, but you should be aware of this to diagnose issues that might come up in slightly different situations (such as, trying to consolidate two PVs into one).

The **pvmove** command is used to facilitate data migration. It provides varying levels of control:

- ▶ Specify source only (“Evacuate this volume”)
- ▶ Specify source and destination (“Replace this volume with this one”)
- ▶ Specify source, destination, and LV (Like the previous example, but only move extents that belong to the named logical volume)
- ▶ Specify physical extents to move and where specifically to move them

The basic process looks as follows:

1. Configure the destination LUN
2. Make the operating system recognize the LUN
3. Format the LUN as a new PV
4. Add the new PV to the VG containing the existing storage system’s PV
5. Migrate all extents
6. Remove the existing PV from the VG
7. Optionally unformat the existing PV

The data is now on the destination LUN, and no new data will be written by LVM to the existing storage. At no point does I/O ever need to be paused or stopped to the volume.

A large advantage of using LVM for data migration is that it also allows for consolidation of LUNs. The process for consolidating LUNs is the same as migrating, except the destination must have enough space to contain all the data stored on both LUNs.

The primary concern with this method is performance. As with any migration, this operation is heavy on disk I/O, but the major difference is that it also uses host CPU cycles and network bandwidth. If the system using this storage is under heavy load, it might be worth considering whether the migration can be run during off-peak hours, or what can be done to reduce the load during the operation.

11.2 Preparation and planning

As highlighted in Chapter 1.4, “Data migration techniques” on page 16, any migration project must start with a preparation and planning phase. Create a diagram of the overall environment and create a list of all the components in your environment.

Check the compatibility of the following components in your environment:

- ▶ Prepare the host servers. Verify that both the host and host HBAs can connect to the existing storage (source) and the new storage (target).
 - Check vendor and model type (change or add HBAs, if necessary).
 - Check the HBA firmware level and upgrade if necessary. Supported HBA firmware levels are available at the System Storage Interoperation Center (SSIC):
<http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>
 - Determine the number and size of each source volume (LUN) on each host.
 - Determine how the LUN is spread across multiple disks for performance considerations.
- ▶ Prepare the new storage system.
 - Make sure that the recommended or later firmware is loaded.
 - Create volumes to meet the specifications of your workload and make sure enough capacity is available.
- ▶ Prepare the SAN fabric.
 - Check fabric compatibility for connection to both the existing existing storage system and the new storage system.
 - Verify that fabric firmware levels are supported with the new storage system (upgrade if necessary).
 - Verify that you have enough free ports to connect the new storage system to the fabric. Add more if needed.
 - Verify the cabling.
 - Connect the new storage system to the fabric.
 - Provide cabling diagrams to the cabling vendor or the person responsible for the cabling.
 - Verify the zoning.
 - Zone the host HBAs to the new storage system.
- ▶ Prepare the existing storage system.
 - Verify that the existing storage system is in optimal condition.
 - Check that all connections, ports, arrays, and LUNs are functioning properly.
 - Check firmware levels (upgrade if necessary).
 - Check that the system is in a data-protected state and not exposed to data loss in case of a disk failure, as migration can be a high-stress workload.
 - Discover and document the logical and physical layout of the storage systems and LUN mappings to the hosts.

11.3 Data migration using the LVM `pvmove` command

Linux's LVM is architecturally similar to other volume management facilities on other operating systems in terms of its storage virtualization and operating flexibility.

The Logical Volume Management on Linux is called LVM. The current version of LVM is called LVM2 and is compatible with earlier versions of Linux LVM1. LVM2 is required for certain clustering and snapshot features, and `pvmove`, which is the process used in this document for migration. LVM1 VGs can be converted to LVM2 using the `vgconvert` command.

LVM2 and the Linux multipath driver, DM-MPIO, require the device mapper kernel driver, a generic device mapper that provides the framework for the volume management.

In this scenario, the system already has LVM2 installed and the volumes to be migrated are part of an existing LVM2 VG.

LVM can move segments used by one or more LVs to the new target devices concurrently using the `pvmove` command. A segment is a contiguous group of extents (by default, 4 MiB chunks) on a given PV. `pvmove` works by moving segments to a new PV using a behind-the-scenes mirror. It can do this one of two ways: non-atomically or atomically.

Tip: You can view the physical extents on PVs connected to your server with the command `pvdisplay -v -m`.

The non-atomic mode creates a temporary LV and uses it to copy the data from the source PV to the target PV in segments. The original LV metadata is updated to point at the corresponding segment in the temporary mirror LV (as opposed to pointing at a physical segment on a PV as in normal operation) until the source and target segments are in sync. It then breaks the mirror and the LV uses the target location for that segment. This process is repeated for each segment to be moved. After all segments are mirrored, the temporary mirror LV is removed and the LV metadata is updated to use the new physical target segments.

The atomic mode works similarly, except it creates one mirror for all segments to be moved. The primary difference is that if the move is cancelled, in non-atomic mode your data is now split across both the old and new PVs, whereas in atomic mode, all data stays on the old PV.

Note: Atomic mode support was introduced in LVM version 2.02.107.

11.3.1 LVM migration scenario

General outline of the LVM migration

The following are the high-level steps used to migrate data with LVM.

1. Configure a target volume on the new storage system.
2. Configure the host to recognize the newly created volume.
3. Configure the Linux multipath driver, DM-MPIO (if not already configured).
4. Initialize the multipath device as an LVM PV.
5. Extend the existing LVM VG to include the new PV.
6. Use `pvmove` to migrate extents off the old PV onto the new PV.
7. Use `vgreduce` to remove the old PV from the VG.

8. Remove the old mpath device from the system and unmap it on the array.

Migration environment

The physical test environment is composed of the following components:

- ▶ Application Server: Red Hat Enterprise Linux 7; kernel 3.10.0-1062.12.1.el7.x86_64
- ▶ Fibre Channel HBA: Qlogic QLE2562 HBAs
- ▶ Source Storage System: IBM FlashSystem 900 (FS900)
- ▶ Source volume: One 50GiB volume
- ▶ Target Storage System: IBM FlashSystem 9100 (FS9100)
- ▶ Target Volume: One 50GiB volume
- ▶ LVM version 2.02.185(2)-RHEL7
- ▶ LVM configuration:
 - One volume group
 - One 40GiB logical volume on the volume group, formatted ext4, with 36GiB used

Important: Back up the data on your volumes before you start the migration process.

LVM configuration before migration

The current LVM configuration for the example consists of a single VG, called `migration_demo`, with one PV on the FS900, `/dev/mpathd.migration_demo` contains one LV, `migratevol` (40 GiB). Note the total number of extents and free extents available on each PV. The extents will be migrated from the old PVs to the new PVs in the VG.

Example 11-1 LVM Configuration before migration

```
[root@migratedemo ~]# vgdisplay -v migration_demo
--- Volume group ---
VG Name                migration_demo
System ID
Format                 lvm2
Metadata Areas         1
Metadata Sequence No  15
VG Access               read/write
VG Status               resizable
MAX LV                 0
Cur LV                 1
Open LV                 1
Max PV                 0
Cur PV                 1
Act PV                 1
VG Size                 <50.00 GiB
PE Size                 4.00 MiB
Total PE                12799
Alloc PE / Size        10240 / 40.00 GiB
Free PE / Size          2559 / <10.00 GiB
VG UUID                e6R4gy-tzQ7-vBcL-DTuu-GUvE-DgnV-o00Zb6

--- Logical volume ---
LV Path                 /dev/migration_demo/migratevol
LV Name                 migratevol
VG Name                 migration_demo
```

```

LV UUID                2vgnuK-j4Sg-gzks-STfN-CMdQ-d5Le-YVfuTm
LV Write Access        read/write
LV Creation host, time comp-1-31, 2020-02-26 12:05:58 -0700
LV Status              available
# open                 1
LV Size                40.00 GiB
Current LE             10240
Segments              1
Allocation             inherit
Read ahead sectors    auto
- currently set to    8192
Block device          253:2

```

--- Physical volumes ---

```

PV Name                /dev/mapper/mpathd
PV UUID                bOHNLI-xMuH-fIHP-G3YJ-F27L-1fvn-4MCXuX
PV Status              allocatable
Total PE / Free PE    12799 / 2559

```

Configure target volume on FS9100

Connect FS9100 to your SAN and zone each HBA port on your Linux host to the configured ports on the FS9100. For best results, use no more than one host initiator per zone. However, more than one storage port can be a member of the zone. Ideally, you would have multiple zones with one host initiator, and one or more storage targets (from the same array) in each zone for multiple logical paths. DM-MPIO handles the I/O and balancing the load on each logical path. In this example we created one 50GiB volume and mapped it to our RHEL7 host.

Configure the Linux host to recognize the new volume

Use procedures appropriate for your Linux distribution and Host Bus Adapter Vendor to add device definitions for the new storage devices from the host system. This example uses the **rescan-scsi-bus.sh** utility (provided by the package `sg3_utils`) to dynamically discover and activate the new FS9100 devices as shown in Example 11-2.

Example 11-2 Dynamic discovery of new FS9100 devices

```

[root@migratedemo ~]# rescan-scsi-bus.sh
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning for device 0 0 0 0 ...
OLD: Host: scsi0 Channel: 00 Id: 00 Lun: 00
      Vendor: TSSTcorp Model: CDDVDW TS-L633F Rev: IT03
      Type: CD-ROM ANSI SCSI revision: 05
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 3 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 4 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning for device 4 0 0 0 ...
OLD: Host: scsi4 Channel: 00 Id: 00 Lun: 00
      Vendor: PNY Model: USB 2.0 FD Rev: 8192
      Type: Direct-Access ANSI SCSI revision: -1
Scanning host 5 for all SCSI target IDs, all LUNs
Scanning for device 5 0 0 0 ...
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 00
      Vendor: IBM Model: 2810XIV-LUN-0 Rev: 0000
      Type: RAID ANSI SCSI revision: 05

```

```

Scanning for device 5 0 0 1 ...
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 01
      Vendor: IBM      Model: 2810XIV      Rev: 0000
      Type:   Direct-Access      ANSI SCSI revision: 05
Scanning for device 5 0 1 0 ...
OLD: Host: scsi5 Channel: 00 Id: 01 Lun: 00
      Vendor: IBM      Model: 2810XIV-LUN-0 Rev: 0000
      Type:   RAID      ANSI SCSI revision: 05
Scanning for device 5 0 1 1 ...
OLD: Host: scsi5 Channel: 00 Id: 01 Lun: 01
      Vendor: IBM      Model: 2810XIV      Rev: 0000
      Type:   Direct-Access      ANSI SCSI revision: 05
Scanning for device 5 0 2 0 ...
NEW: Host: scsi5 Channel: 00 Id: 02 Lun: 00
      Vendor: IBM      Model: 2145      Rev: 0000
      Type:   Direct-Access      ANSI SCSI revision: 06
Scanning for device 5 0 2 1 ...
OLD: Host: scsi5 Channel: 00 Id: 02 Lun: 01
      Vendor: IBM      Model: 2145      Rev: 0000
      Type:   Direct-Access      ANSI SCSI revision: 06
...
[snip]
...
Scanning host 7 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning for device 7 0 0 0 ...
NEW: Host: scsi7 Channel: 00 Id: 00 Lun: 00
      Vendor: IBM-207x Model: ST9300603SS Rev: B53E
      Type:   Direct-Access      ANSI SCSI revision: 05
Scanning for device 7 0 1 0 ...
NEW: Host: scsi7 Channel: 00 Id: 01 Lun: 00
      Vendor: IBM-207x Model: ST9300603SS Rev: B53E
      Type:   Direct-Access      ANSI SCSI revision: 05
Scanning for device 7 1 3 0 ...
OLD: Host: scsi7 Channel: 01 Id: 03 Lun: 00
      Vendor: LSILOGIC Model: Logical Volume Rev: 3000
      Type:   Direct-Access      ANSI SCSI revision: 02
7 new or changed device(s) found.
  [5:0:2:0]
  [5:0:3:0]
  [6:0:3:0]
  [6:0:4:0]
  [6:0:8:0]
  [7:0:0:0]
  [7:0:1:0]
0 remapped or resized device(s) found.
0 device(s) removed.

```

Install and Configure DM-MPIO

If DM-MPIO is not already installed and configured, follow the appropriate steps for your Linux distribution to do so. For RHEL 7, first ensure the `device-mapper-multipath` package is installed, then run the commands to enable multipath discovery as in Example 11-3.

Example 11-3 Enable and configure DM-MPIO

```
[root@migratedemo ~]# mpathconf --enable --with_multipathd y
```

```
[root@migratedemo ~]# systemctl enable multipathd
Created symlink from /etc/systemd/system/sysinit.target.wants/multipathd.service
to /usr/lib/systemd/system/multipathd.service.
[root@migratedemo ~]# pvs
```

PV	VG	Fmt	Attr	PSize	PFree
/dev/mapper/mpathd	migration_demo	lvm2	a--	<50.00g	<10.00g
/dev/sdo3	rhel_migratedemo	lvm2	a--	232.00g	0

Determine which mpath devices map to which volumes

There can be many mpath devices configured, and you need to be certain that you are referencing the correct ones. `dmsetup` provides a simple method to list the SCSI UUIDs of the multipath volumes, thus allowing you to verify which devices correspond to which volumes.

Example 11-4 List of device mapper devices and their UUIDs

```
[root@migratedemo ~]# dmsetup info -c
```

Name	Maj	Min	Stat	Open	Targ	Event	UUID
mpathd	253	5	L--w	1	1	0	mpath-360050761018f88110000000001000003
mpathc	253	7	L--w	0	1	0	mpath-360050768108005f9b000000000000024
mpathb	253	4	L--w	1	1	0	mpath-360050768108005f9b00000000000001f
mpatha	253	3	L--w	0	1	0	mpath-36001738cce051a5000000000002bb59
rhel_migratedemo-swap	253	1	L--w	2	1	0	
LVM-3L6PtoiGmL5Hxig9gvaQUhfhyWaiEBYnHnhg1TcMg1Ipphv0eF2B2DqccRrUjau							
rhel_migratedemo-root	253	0	L--w	1	1	0	
LVM-3L6PtoiGmL5Hxig9gvaQUhfhyWaiEBYa20qwJjMDwi5mCdQpxYSKuZxLI340J9y							
migration_demo-migratevol	253	2	L--w	1	1	0	
LVM-e6R4gytzQ7vBcLDTuuGUvEDgnVo00Zb62vgnuKj4SggzksSTfNCMdQd5LeYVFuTm							
mpathb1	253	6	L--w	0	1	0	part1-mpath-360050768108005f9b00000000000001f

Here, the original volume has the UUID 60050761018f88110000000001000003 and thus is device `mpathd`, while the new volume has the UUID 60050768108005F9B000000000000024 and thus is device `mpathc`.

Initialize the mpath devices for LVM

Initializing the mpath devices prepares the disk for use by the LVM as a physical volume (PV). The `pvcreate` command initializes the device as shown in Example 11-5.

Example 11-5 Initialize mpath partitions for LVM

```
[root@migratedemo ~]# pvcreate /dev/mapper/mpathc
Physical volume "/dev/mapper/mpathc" successfully created.
[root@migratedemo ~]# pvs
```

PV	VG	Fmt	Attr	PSize	PFree
/dev/mapper/mpathc		lvm2	---	50.00g	50.00g
/dev/mapper/mpathd	migration_demo	lvm2	a--	<50.00g	<10.00g
/dev/sdo3	rhel_migratedemo	lvm2	a--	232.00g	0

Extend the LVM VG to include the new mpath device

Tip: You can use the following commands to delete the partition table on the disk if one exists. Use these commands with caution.

```
# dd if=/dev/zero of=/dev/diskname bs=1k count=1
# blockdev --rereadpt /dev/diskname
```

The new mpath device needs to be added to the existing LVM VG, `migration_demo`. Add the devices with the `vgextend` command as shown in Example 11-6.

Example 11-6 Using vgextend to add the new mpath device to the VG

```
[root@migratedemo ~]# vgextend migration_demo /dev/mapper/mpathc
Volume group "migration_demo" successfully extended
```

Use **vgdisplay** to verify the addition of the new mpath device (Example 11-7). Notice the free Physical Extents (PE) available to the migration_demo volume group on each device.

Example 11-7 pvdisplay showing the new mpath device

```
[root@migratedemo ~]# vgdisplay -v migration_demo
--- Volume group ---
VG Name                migration_demo
System ID
Format                 1vm2
Metadata Areas        2
Metadata Sequence No  16
VG Access              read/write
VG Status              resizable
MAX LV                 0
Cur LV                1
Open LV                1
Max PV                 0
Cur PV                2
Act PV                 2
VG Size                99.99 GiB
PE Size                4.00 MiB
Total PE               25598
Alloc PE / Size       10240 / 40.00 GiB
Free PE / Size         15358 / 59.99 GiB
VG UUID                e6R4gy-tzQ7-vBcL-DTuu-GUvE-DgnV-o00Zb6

--- Logical volume ---
LV Path                /dev/migration_demo/migratevol
LV Name                migratevol
VG Name                migration_demo
LV UUID                2vgnuK-j4Sg-gzks-STfN-CMdQ-d5Le-YVfuTm
LV Write Access        read/write
LV Creation host, time comp-1-31, 2020-02-26 12:05:58 -0700
LV Status              available
# open                 1
LV Size                40.00 GiB
Current LE             10240
Segments               1
Allocation              inherit
Read ahead sectors     auto
- currently set to     8192
Block device           253:2

--- Physical volumes ---
PV Name                /dev/mapper/mpathd
PV UUID                b0HNL1-xMuH-fIHP-G3YJ-F27L-1fvn-4MCXuX
PV Status              allocatable
Total PE / Free PE    12799 / 2559

PV Name                /dev/mapper/mpathc
```

```
PV UUID          DYgPMm-FiCu-I7L4-7n7u-sUhT-58J6-8j9rLp
PV Status        allocatable
Total PE / Free PE 12799 / 12799
```

Use `pvmove` to migrate extents within the volume group

To use `pvmove` to migrate the extents within the VG, perform the following steps:

1. Start the migration of extents using the `pvmove` command. It migrates extents within a VG from a source device to a target device as shown in Example 11-8. Issue the command:

```
pvmove source target
```

Where `source` and `target` are your source and target devices in the volume group. If you want to use atomic mode, and your version of LVM supports it, pass the `--atomic` flag:

```
pvmove --atomic source target
```

Important: The `pvmove` command takes a long time to run because it performs a block-by-block copy.

Example 11-8 `pvmove` in progress

```
[root@migratedemo ~]# pvmove --atomic /dev/mapper/mpathd /dev/mapper/mpathc
/dev/mapper/mpathd: Moved: 0.05%
/dev/mapper/mpathd: Moved: 19.26%
/dev/mapper/mpathd: Moved: 39.41%
/dev/mapper/mpathd: Moved: 59.39%
/dev/mapper/mpathd: Moved: 79.56%
/dev/mapper/mpathd: Moved: 99.66%
/dev/mapper/mpathd: Moved: 100.00%
[root@migratedemo ~]#
```

If `pvmove` needs to be stopped during the migration, pressing `Ctrl+C` then issuing the `pvmove --abort` command will abort. If not using atomic mode, the extents are spread between the old device and the new device. The migration can be continued by issuing the `pvmove` command again. After pressing `Ctrl+C`, the migration will continue in the background. If you do not want to cancel, you can let the migration continue or run `pvmove` without any arguments to monitor the status of any currently running migrations.

2. After `pvmove` is complete, verify the migration with the `vgdisplay` command as shown in Example 11-9. Notice that all the extents are free on `mpathd` and are moved to `mpathc`.

Example 11-9 `vgdisplay` showing the extents migrated from device `mpathd` to `mpathc`

```
[root@migratedemo ~]# vgdisplay -v migration_demo
--- Volume group ---
VG Name          migration_demo
System ID
Format           lvm2
Metadata Areas   2
Metadata Sequence No 34
VG Access        read/write
VG Status        resizable
MAX LV           0
Cur LV          1
Open LV          1
Max PV           0
Cur PV          2
```

```

Act PV                2
VG Size               99.99 GiB
PE Size               4.00 MiB
Total PE              25598
Alloc PE / Size      10240 / 40.00 GiB
Free PE / Size       15358 / 59.99 GiB
VG UUID               e6R4gy-tzQ7-vBcL-DTuu-GUvE-DgnV-o00Zb6

--- Logical volume ---
LV Path                /dev/migration_demo/migratevol
LV Name                migratevol
VG Name                migration_demo
LV UUID                2vgnuK-j4Sg-gzks-STfN-CMdQ-d5Le-YVFuTm
LV Write Access        read/write
LV Creation host, time comp-1-31, 2020-02-26 12:05:58 -0700
LV Status              available
# open                 1
LV Size                40.00 GiB
Current LE             10240
Segments               1
Allocation              inherit
Read ahead sectors     auto
- currently set to    0
Block device           253:2

--- Physical volumes ---
PV Name                /dev/mapper/mpathd
PV UUID                b0HNL1-xMuH-fIHP-G3YJ-F27L-1fvn-4MCXuX
PV Status              allocatable
Total PE / Free PE    12799 / 12799

PV Name                /dev/mapper/mpathc
PV UUID                DYgPMm-FiCu-I7L4-7n7u-sUHT-58J6-8j9rLp
PV Status              allocatable
Total PE / Free PE    12799 / 2559

```

Removing the old source device from LVM

After all the physical extents are migrated, the physical volume can be removed from the volume group using the **vgreduce** command, and then wipe the LVM labels with the **pvremove** command (Example 11-10).

Example 11-10 The vgreduce command to remove a physical volume from a volume group

```

[root@migratedemo ~]# vgreduce migration_demo /dev/mapper/mpathd
Removed "/dev/mapper/mpathd" from volume group "migration_demo"
[root@migratedemo ~]# pvremove /dev/mapper/mpathd
Labels on physical volume "/dev/mapper/mpathd" successfully wiped.

```

Removing the old source device from the system

1. Remove the old multipath volume and verify that it is no longer listed (Example 11-11).

Example 11-11 Remove multipath device

```

[root@migratedemo ~]# multipath -f /dev/mapper/mpathd
[root@migratedemo ~]# dmsetup info -c

```

Name	Maj	Min	Stat	Open	Targ	Event	UUID
mpathc	253	7	L--w	1	1	0	
mpath-360050768108005f9b00000000000024							
mpathb	253	4	L--w	1	1	0	
mpath-360050768108005f9b0000000000001f							
mpatha	253	3	L--w	0	1	0	
mpath-36001738ccce051a5000000000002bb59							
rhel_migratedemo-swap	253	1	L--w	2	1	0	
LVM-3L6PtoiGmL5Hxig9gvaQUhfhyWaiEBYnHnhg1TcMg1Ipqhv0eF2B2DqccRrUjau							
rhel_migratedemo-root	253	0	L--w	1	1	0	
LVM-3L6PtoiGmL5Hxig9gvaQUhfhyWaiEBYa20qwJjMDwi5mCdQpxYSKuZxLI340J9y							
migration_demo-migratevol	253	2	L--w	1	1	0	
LVM-e6R4gytzQ7vBcLDTuuGUvEDgnVo00Zb62vgnuKj4SggzksSTfNCMdQd5LeYVFuTm							
mpathb1	253	6	L--w	0	1	0	
part1-mpath-360050768108005f9b0000000000001f							

2. Change SAN zoning to remove access to old source volumes.
3. Use procedures appropriate for your operating system and Host Bus Adapter Vendor to remove device definitions for the old storage controller devices from the system.

This example uses the **rescan-scsi-bus.sh** utility (provided by the package `sg3_utils`) to dynamically remove the unmapped FS900 devices as shown in Example 11-12.

Example 11-12 Qlogic dynamic reconfiguration tool to remove old storage devices

```
[root@migratedemo ~]# rescan-scsi-bus.sh -r
Syncing file systems
Scanning SCSI subsystem for new devices and remove devices that have disappeared
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 0 0 0 ...
OLD: Host: scsi0 Channel: 00 Id: 00 Lun: 00
      Vendor: TSSTcorp Model: CDDVDW TS-L633F Rev: IT03
      Type: CD-ROM ANSI SCSI revision: 05
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 3 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 4 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 4 0 0 ...
...
[snip]
...
sg17 changed: LU not available (PQual 1)
REM: Host: scsi6 Channel: 00 Id: 06 Lun: 00
DEL: Vendor: IBM Model: FlashSystem-9840 Rev: 1611
      Type: Direct-Access ANSI SCSI revision: 05
Host: scsi6 Channel: 00 Id: 06 Lun: 00
      Vendor: IBM Model: FlashSystem-9840 Rev: 1611
      Type: Direct-Access ANSI SCSI revision: 05
  Scanning for device 6 0 8 0 ...
OLD: Host: scsi6 Channel: 00 Id: 08 Lun: 00
      Vendor: IBM Model: 2145 Rev: 0000
      Type: Direct-Access ANSI SCSI revision: 06
  Scanning for device 6 0 8 1 ...
OLD: Host: scsi6 Channel: 00 Id: 08 Lun: 01
      Vendor: IBM Model: 2145 Rev: 0000
      Type: Direct-Access ANSI SCSI revision: 06
```

```
Scanning host 7 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 7 0 0 0 ...
NEW: Host: scsi7 Channel: 00 Id: 00 Lun: 00
      Vendor: IBM-207x Model: ST9300603SS      Rev: B53E
      Type:   Direct-Access                    ANSI SCSI revision: 05
  Scanning for device 7 0 1 0 ...
NEW: Host: scsi7 Channel: 00 Id: 01 Lun: 00
      Vendor: IBM-207x Model: ST9300603SS      Rev: B53E
      Type:   Direct-Access                    ANSI SCSI revision: 05
  Scanning for device 7 1 3 0 ...
OLD: Host: scsi7 Channel: 01 Id: 03 Lun: 00
      Vendor: LSILOGIC Model: Logical Volume   Rev: 3000
      Type:   Direct-Access                    ANSI SCSI revision: 02
2 new or changed device(s) found.
  [7:0:0:0]
  [7:0:1:0]
0 remapped or resized device(s) found.
4 device(s) removed.
  [5:0:4:0]
  [5:0:6:0]
  [6:0:2:0]
  [6:0:6:0]
```



Host-based migrations in VMware Storage vMotion

This chapter addresses migrating data from a block source storage system to a Spectrum Virtualize target device using VMware Storage vMotion.

This chapter contains the following sections:

- ▶ Preparation and planning
- ▶ Data migration using VMware Storage vMotion

12.1 Preparation and planning

As highlighted in Chapter 1.4, “Data migration techniques” on page 16, any migration project must start with a preparation and planning phase. Create a diagram of the overall environment and create a list of all the components in your environment.

Check the compatibility of the following components in your environment:

- ▶ Prepare the VMware ESXi hosts. Verify that both the host and host HBAs, CNAs, or network adapter can connect to the existing storage (source) and the Spectrum Virtualize system (target). See the Interoperability matrix available at:

<https://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>

- Check vendor and model type.
- Check the adapter firmware level and upgrade if necessary.
- Determine the number and size of each source volume (LUN) on each host.
- ▶ Prepare the Spectrum Virtualize system and volumes.
 - Make sure the storage system is running the recommended code level or a higher code level. More details are available at:
<https://www.ibm.com/support/pages/node/690527>
 - Check if enough usable capacity is available to do the migration. Check if the source storage system is using data reduction features and what capacity is used. You might need also Data Reduction Pools and data reduction features such as thin-provisioning, compression and deduplication. If your storage system is not showing the Data Reduction ratio or not using Data Reduction it might be necessary to run the Data Reduction Estimator Tool (DRET). More details about DRET are available at:
https://www.ibm.com/support/knowledgecenter/STSLR9_8.3.1/com.ibm.fs9200_831.doc/dret_overview.html
 - Create the needed volumes.
 - Map the volumes to the ESXi host.
- ▶ Prepare the SAN or Ethernet connections.
 - Depending on connection protocols in use in your environment, check FC, iSCSI, iSER, FCoE, or SAS compatibility for connection to both the existing storage system and the Spectrum Virtualize system. See the Interoperability matrix available at:
<https://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>
 - Verify that you have enough free ports to connect the new Spectrum Virtualize system to the fabric. Add more if needed.
 - Verify the cabling.
 - Order the correct number of cables.
 - Connect the storage system to the fabric or direct to the host.
 - Provide cabling diagrams to the cabling vendor or the person responsible for the cabling.
 - Verify the zoning.
 - Zone the host HBAs to the Spectrum Virtualize system. The optimum number of paths from each host to one IBM Spectrum Virtualize volume is four paths, even though the maximum supported is eight paths.

For more details see:

IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines, SG24-7521-06.

- ▶ Prepare the existing storage system being replaced.
 - Verify that the existing storage system is in optimal condition.
 - Check that all connections, ports, arrays, and LUNs are functioning properly.
 - Check firmware levels (upgrade if necessary).
 - Discover and document the logical and physical layout of the arrays and LUN mappings to the hosts.

12.2 Data migration using VMware Storage vMotion

ESXi host uses a special file system, called Virtual Machine File System (VMFS), to store virtual machines.

VMFS is a clustered file system that can be accessed concurrently by up to 32 ESXi hosts. VMFS maximum volume size is 64 TB, and the maximum file size on VMFS is 62 TB.

More details about the vSphere 6.7 configuration maximums can be found at the following web page:

<https://configmax.vmware.com/guest?vmwareproduct=vSphere&release=vSphere%206.7&categories=2-0>

On a VMFS, ESXi stores all the files of which a VM consists:

- ▶ Virtual machine configuration file (.vmx)
- ▶ Additional virtual machine configuration file (.vmxf)
- ▶ Virtual machine BIOS file (.nvram)
- ▶ Virtual disk characteristics (.vmdk)
- ▶ Virtual machine data disk (-flat.vmdk)
- ▶ Virtual machine snapshots (.vmsd)
- ▶ Virtual machine snapshot data file (.vmsn)
- ▶ Virtual machine swap file (.vswp)
- ▶ Virtual machine suspend file (.vmss)
- ▶ Virtual machine log files (.log)
- ▶ Old virtual machine log files (-#.log), where # is a number starting with 1

VMs can use two different concepts for their virtual drives: virtual disks and raw device mappings (RDM). Both concepts are shown in Figure 12-1 on page 192.

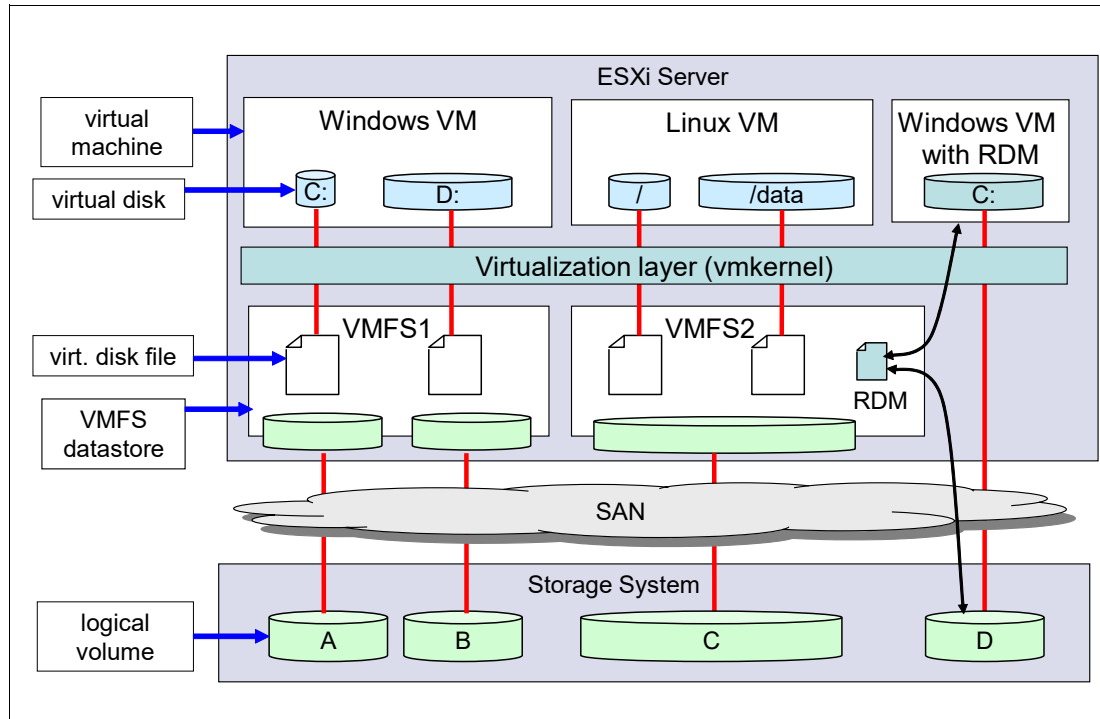


Figure 12-1 Storage concepts of ESXi and VMFS

VMware storage vMotion is a feature of vSphere and comes with VMware vSphere Standard, VMware vSphere Enterprise Plus, and VMware vSphere Platinum. Storage vMotion enables the ability to migrate live virtual machine's (VM's) files from one VMFS, Network File System (NFS), Virtual Volume (VVOL), or Virtual SAN (vSAN) data store to another VMFS, NFS, VVOL, or vSAN data store without downtime for the guest.

For more details about VMware datastores, see:

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.storage.doc/GUID-D5AB2BAD-C69A-4B8D-B468-25D86B8D39CE.html>

Storage vMotion copies the virtual machine metadata from the source to the destination data store. Making use of changed block tracking to maintain data integrity, storage vMotion copies the virtual machine virtual disks as well to the destination data store. After the first copy, storage vMotion determines what regions of the virtual disk were written during the first copy and starts a second copy of those regions that were changed by running applications. It will continue to do so until it has copied most of the data over to the other data store. Physical Mode Raw Device Mapping (RDM) data cannot be migrated to a VMDK disk with Storage vMotion, more details about VM migrations with RDMs can be found at the following web page:

<https://kb.vmware.com/s/article/1005241>

The switchover process is fast enough to be unnoticed by the guest.

Note that with VMware vSphere Storage APIs for Array Integration (VAAI), the data mover will try to offload the copy operation to the storage system, so that the copy between the source and destination data store will not use the resources of the ESXi host and will be handled by the storage system.

VAAI will be used for migrations (Storage vMotions) within one storage system, but not for migrations from one storage system to another storage system, which is the case for migrations for our example. To check if a storage system supports VAAI, see the VMware Storage/SAN Compatibility Guide at:

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

Figure 12-2 illustrates Storage vMotion of a virtual machine from IBM FlashSystem A9000R data store to an IBM FlashSystem 9100 data store.

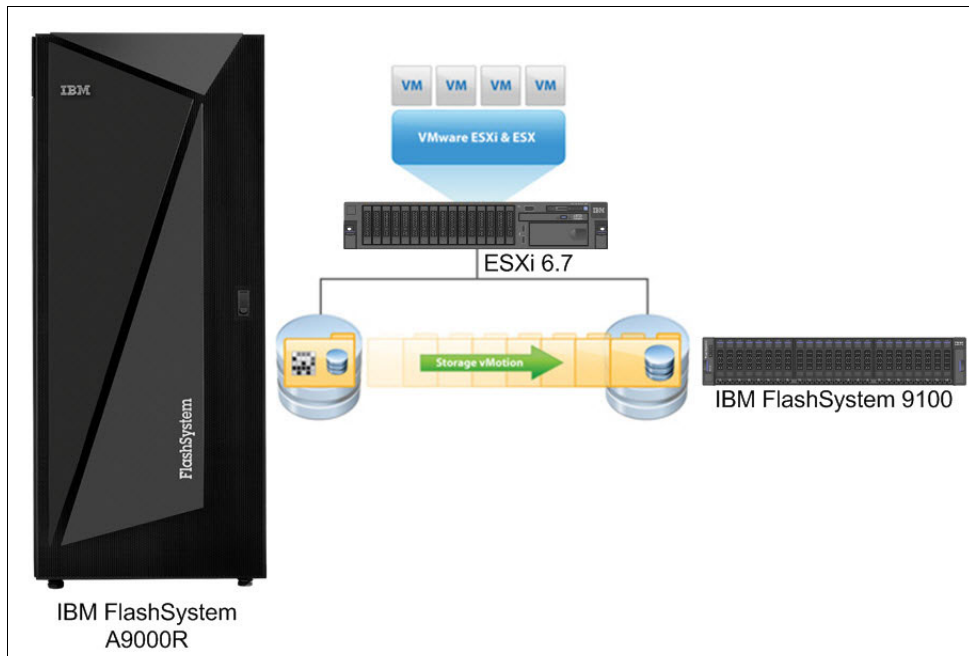


Figure 12-2 VMware Storage vMotion

The necessary steps for the migration process are documented in the following sections.

Open the vSphere Client (HTML5) on the vCenter Server `https://<vCenter IP address>/ui` in your browser and go to **Menu** → **Hosts and Clusters** as depicted in Figure 12-3.

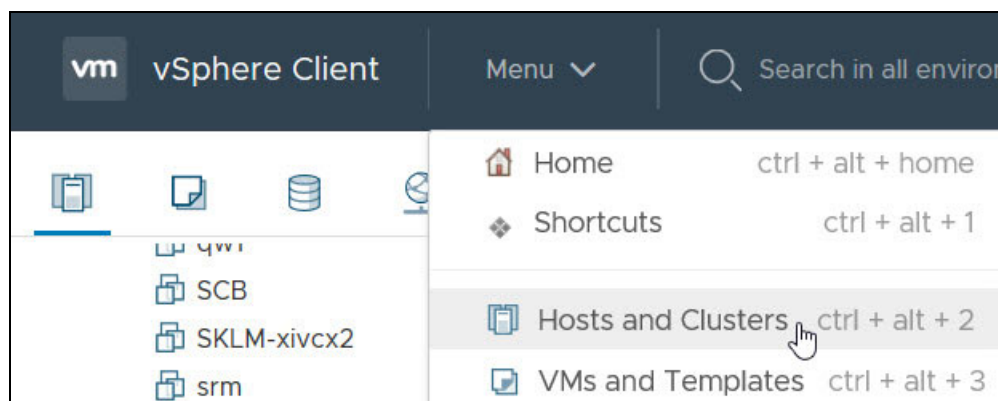


Figure 12-3 vSphere Client Hosts and Clusters

In our environment, we just have two virtual machines (VMs) running as shown in Figure 12-4 on page 194, one is the VMware vCenter Server Appliance (VCSA) and the other one is a Windows 2016 Guest VM running on a FlashSystem A9000R data store.

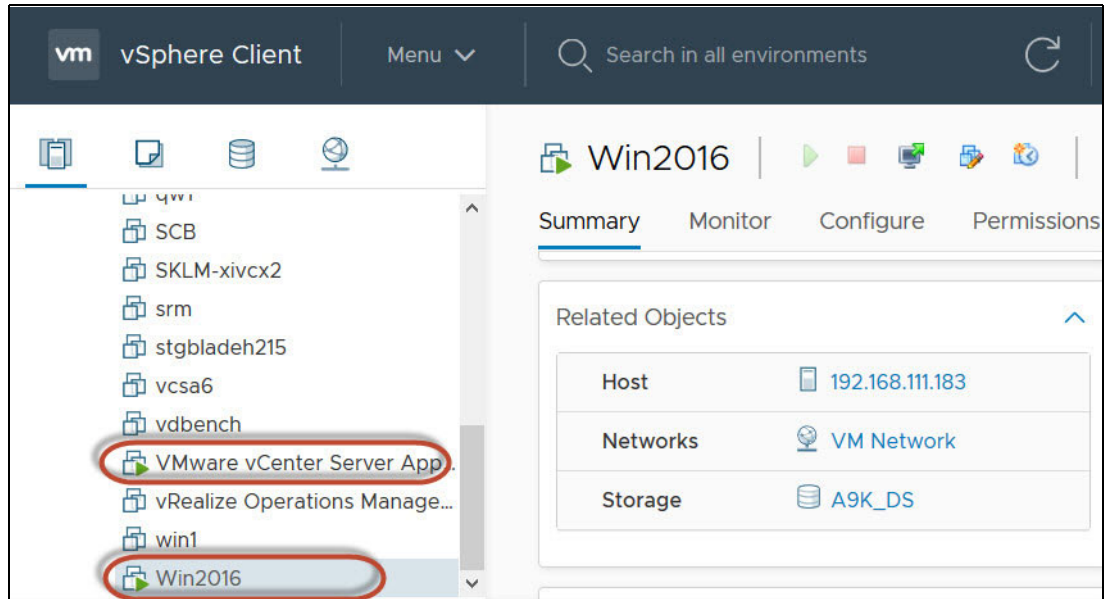


Figure 12-4 Windows 2016 VM

To migrate the VM from the FlashSystem A9000R (source) to FlashSystem 9100 (target) we need to create a host and a volume on the FlashSystem 9100 and map it to the VMware ESXi host and then create a data store on the volume.

Open the FlashSystem 9100 GUI by entering `https://<cluster-ip-address>` in a browser and signing in as illustrated in Figure 12-5.

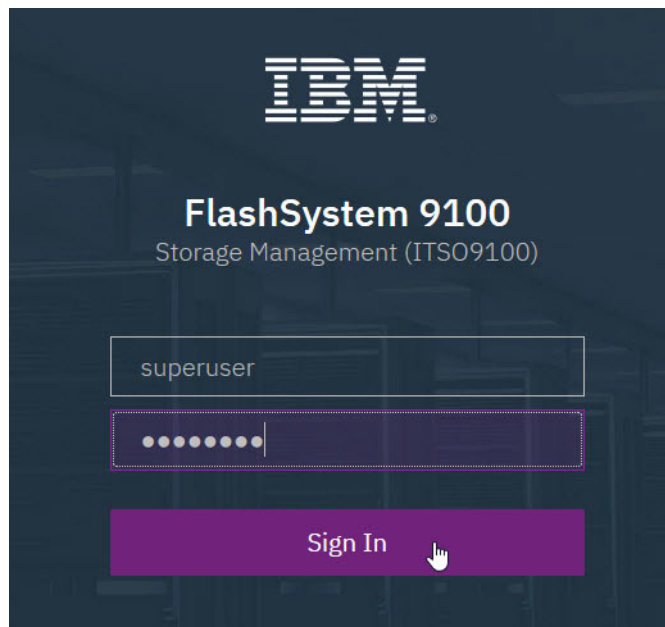


Figure 12-5 FlashSystem 9100 GUI Sign In

Select **Hosts** → **Hosts** in the menu as depicted in Figure 12-6.

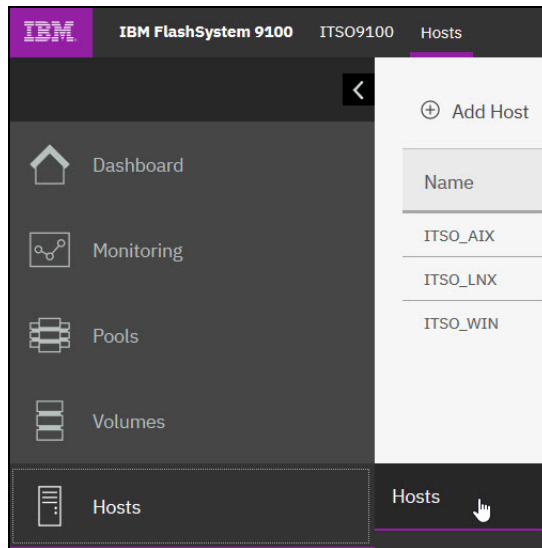


Figure 12-6 Hosts menu item

Click **Add Host**, as shown in Figure 12-7 and choose a name for the host, add the worldwide port names (WWPNs) of the HBA ports of the host and click **Add** as illustrated in Figure 12-8.




Figure 12-7 Add host

A screenshot of the 'Add Host' configuration form. The form has a title 'Add Host' and a warning message: 'Because NPIV is enabled on this system, host traffic is only allowed over the storage system's virtual ports. Ensure that SAN zoning allows connectivity between virtual ports and the host.' Below the warning, there are two sections: 'Required Fields' and 'Optional Fields'. The 'Required Fields' section includes: 'Name' (text input with 'ITSO_ESXi'), 'Host connections' (dropdown menu with 'Fibre Channel (SCSI)'), 'Host port (WWPN)' (text input with '10000090FAA30C8D'), and 'Host port (WWPN)' (text input with '10000090FAA30C8E'). The 'Optional Fields' section includes: 'Host type' (dropdown menu with 'Generic'). At the bottom right, there are 'Cancel' and 'Add' buttons. A mouse cursor is pointing at the 'Add' button.

Figure 12-8 Add ESXi host

The host is listed in the Hosts pane, as depicted in Figure 12-9.

Name	Status	Host Type	# of Ports	Host Mapping
ITSO_AIX	✓ Online	Generic	2	Yes
ITSO_ESXi	✓ Online	Generic	2	No
ITSO_LNX	✓ Online	Generic	2	Yes
ITSO_WIN	✓ Online	Generic	2	Yes

Figure 12-9 Hosts list

To create and map a volume to the ESXi host select **Volumes** → **Volumes** from the menu shown in Figure 12-10.

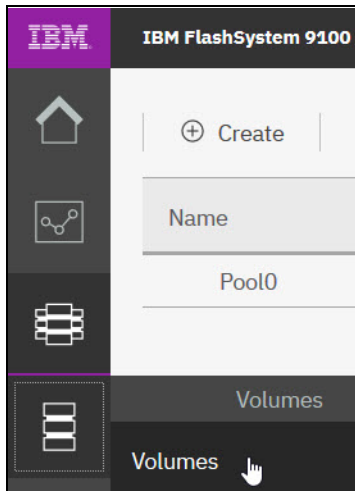


Figure 12-10 Volumes menu item

Click **Create Volumes** as illustrated in Figure 12-11.




Figure 12-11 Create Volumes

Choose the pool, the name, and the data reduction features for the volume, as size select an equal or bigger size as all the files of the VMs (which need to be migrated) need.

Click **Create and Map** as shown in Figure 12-12.

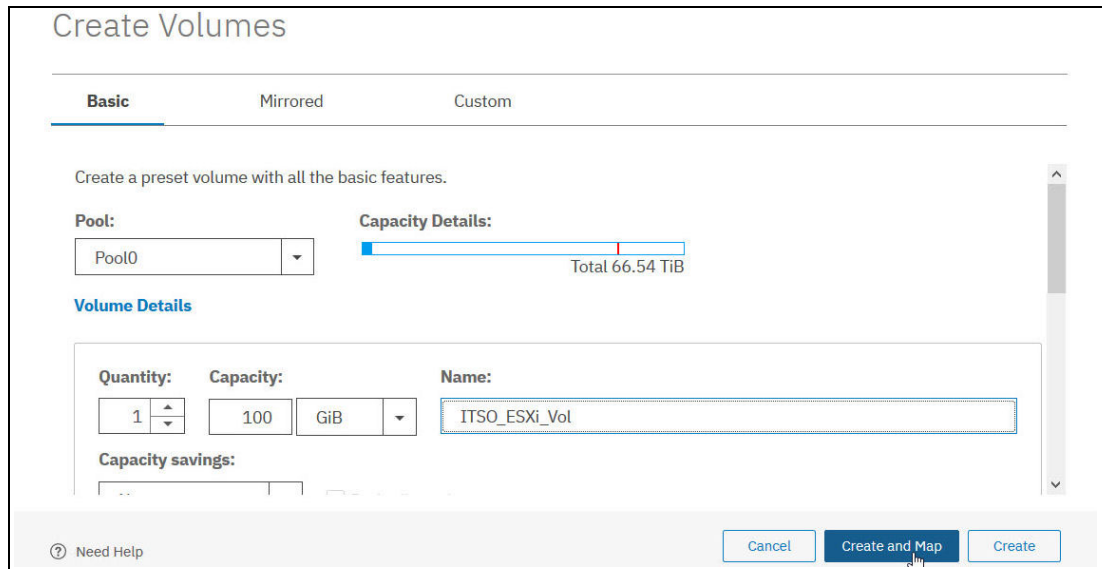


Figure 12-12 Create and Map Volumes

Select the ESXi host and click **Next** as shown in Figure 12-13.

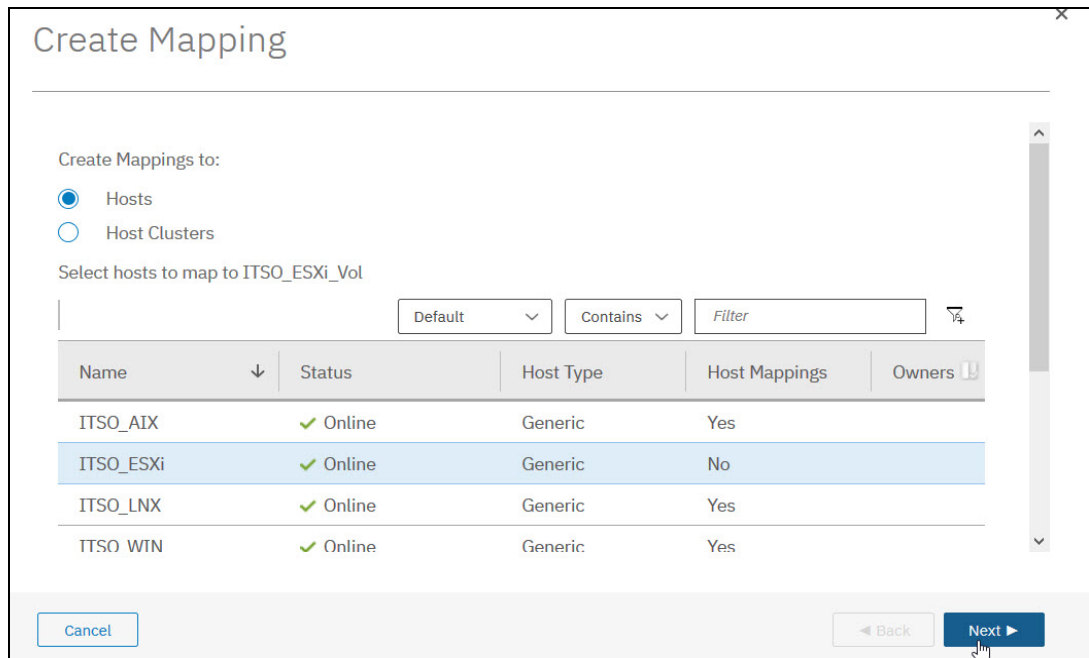


Figure 12-13 Create Mapping - Select host

Click **Map Volumes** to map the volume to the ESXi host as illustrated in Figure 12-14.

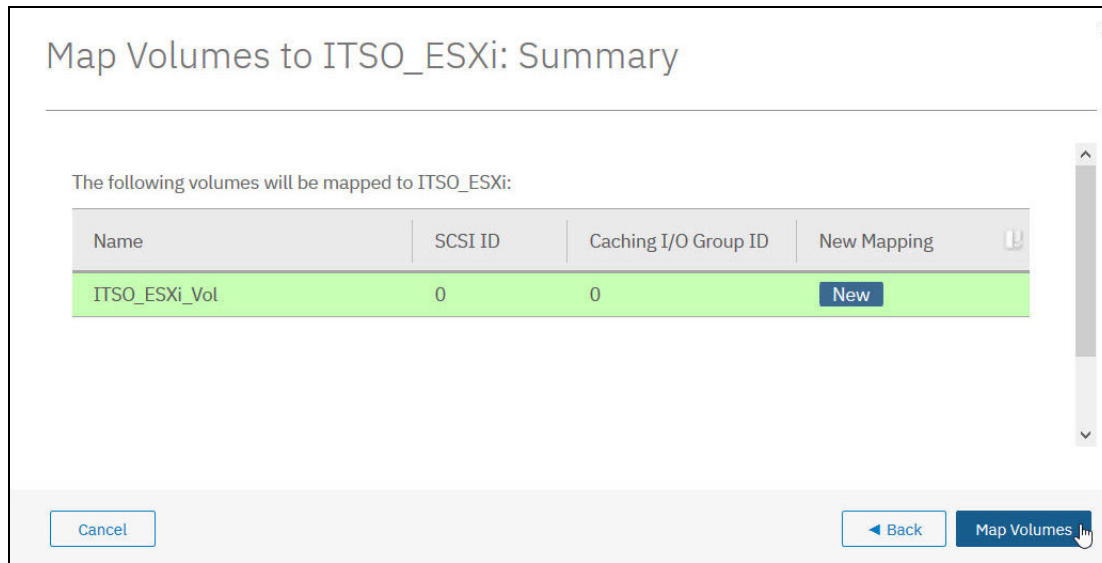


Figure 12-14 Create Mapping- Volume list

The volume is visible in the Volume pane as depicted in Figure 12-15.

Name	State	Synchronized	Pool	Protocol
ITSO_AIX_Data	✓ Online		Pool0	SCSI
ITSO_ESXi_Vol	✓ Online (formatting)		Pool0	SCSI

Figure 12-15 Volume list

To be able to migrate the VM files we need to create a data store on top of the FlashSystem 9100 volume. In the vSphere Client, select **Menu** → **Storage**, right-click the data center **ITSO-DC** and select **Storage** → **New data store**, as shown in Figure 12-16.

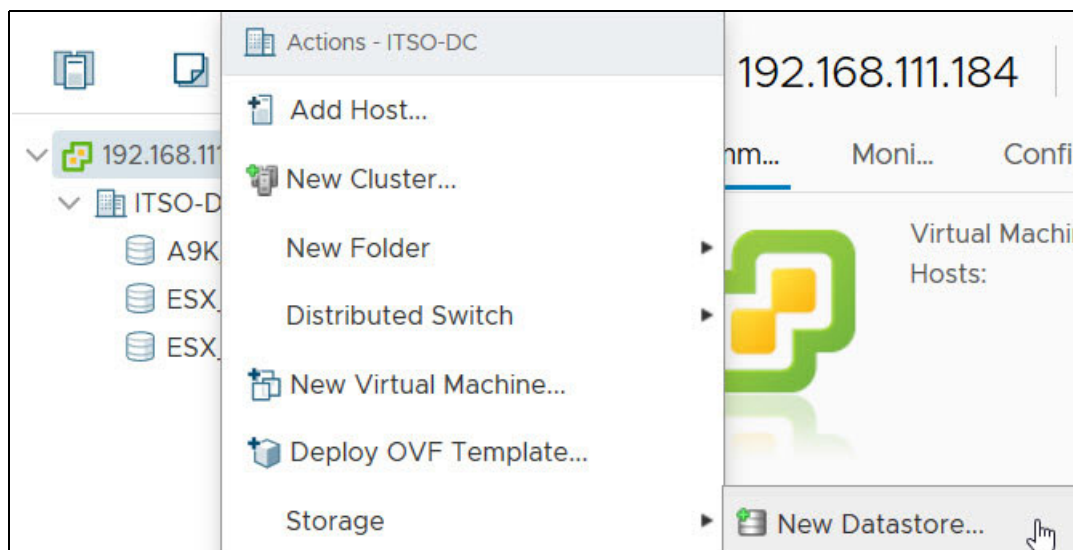


Figure 12-16 New data store

The New data store wizard opens as illustrated in Figure 12-17. Select VMFS as data store type and click **NEXT**.

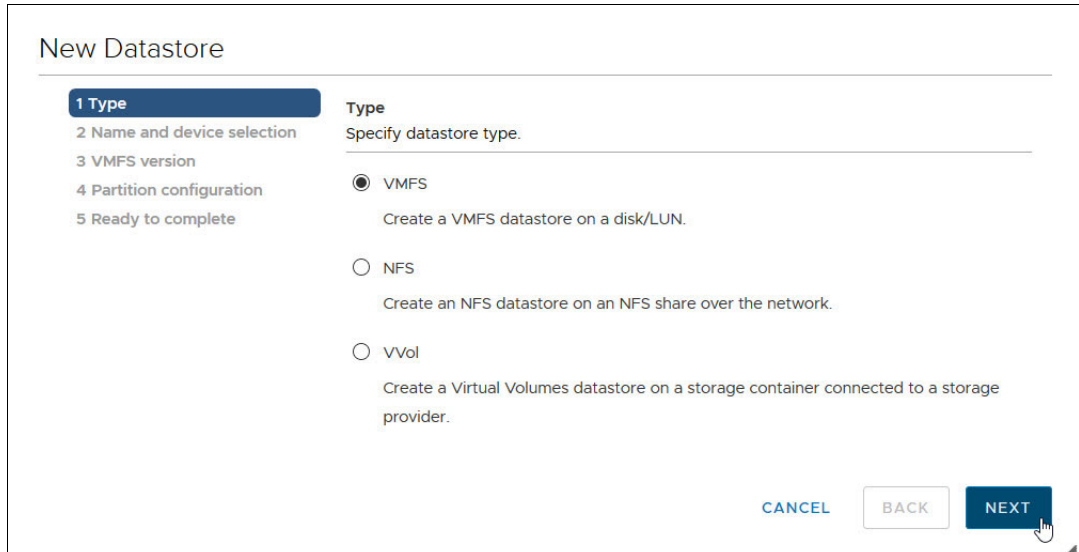


Figure 12-17 New data store - Type

Choose a name, select the ESXi host, select the device, and click **NEXT** as depicted in Figure 12-18. If the device is not visible, a rescan is necessary, so you would have to **CANCEL** the data store wizard, right-click data center and select **Storage** → **Rescan Storage** and then restart the data store wizard.

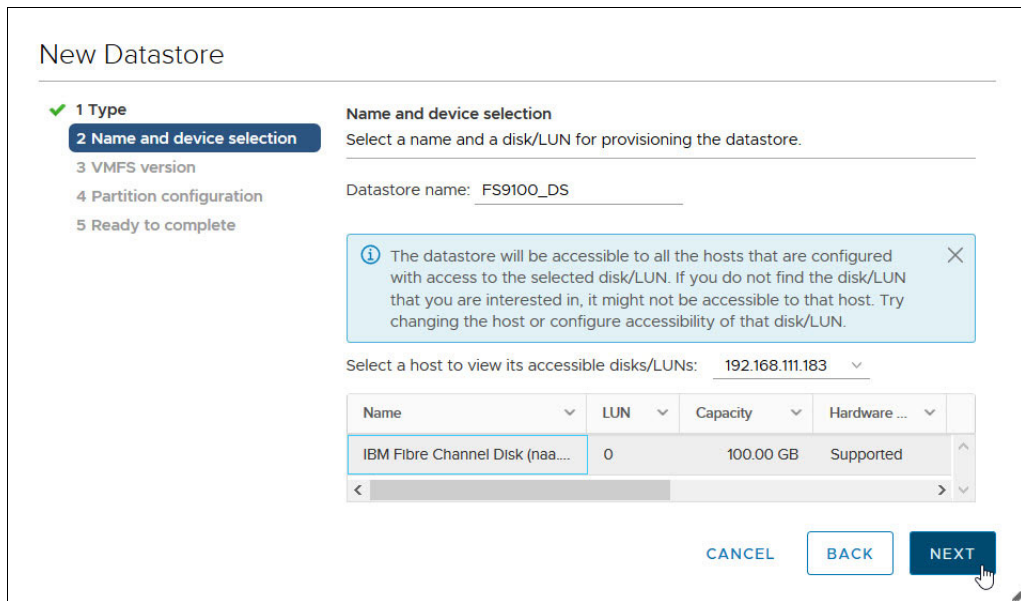


Figure 12-18 New data store - Name and Device

Select the VMFS version to use. The version can be higher than the source data store. Click **NEXT** as shown in Figure 12-19.

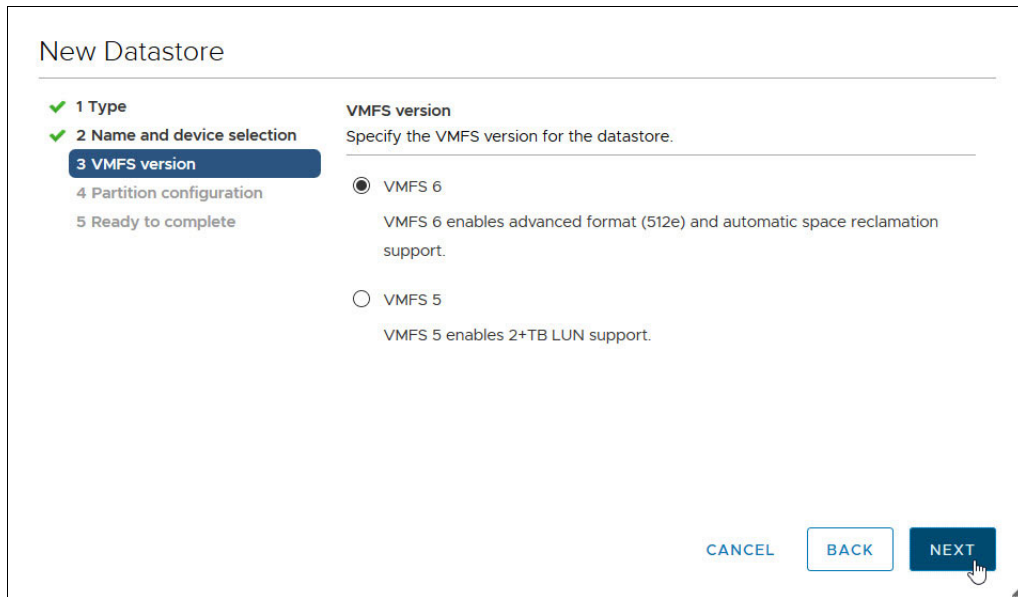


Figure 12-19 New data store - VMFS version

Select Use all available partitions and click **NEXT** as indicated in Figure 12-20.

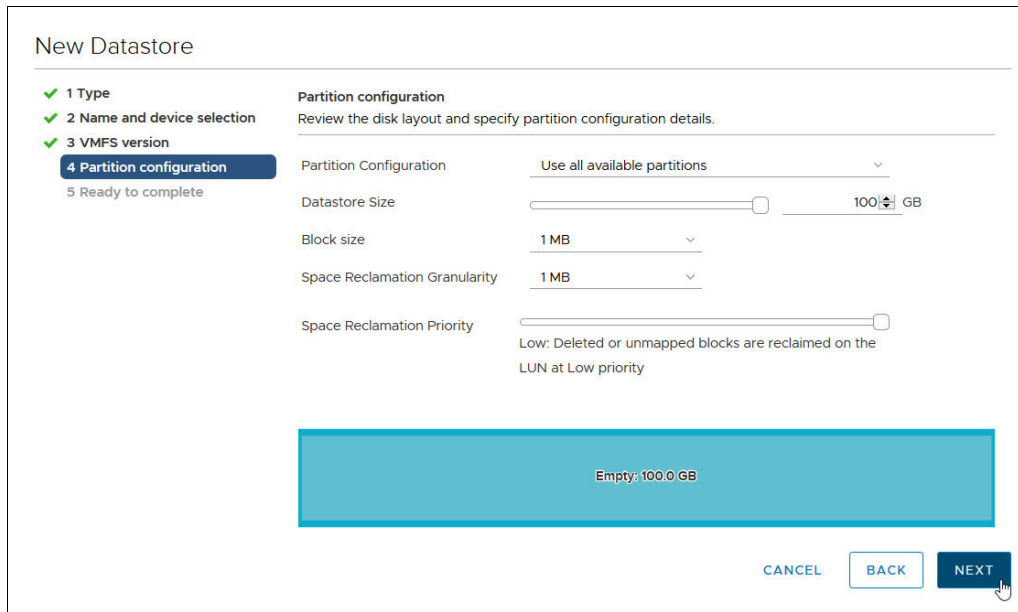


Figure 12-20 New data store - Partition Configuration

The summary is shown, click **FINISH** to create the target data store as depicted in Figure 12-21.

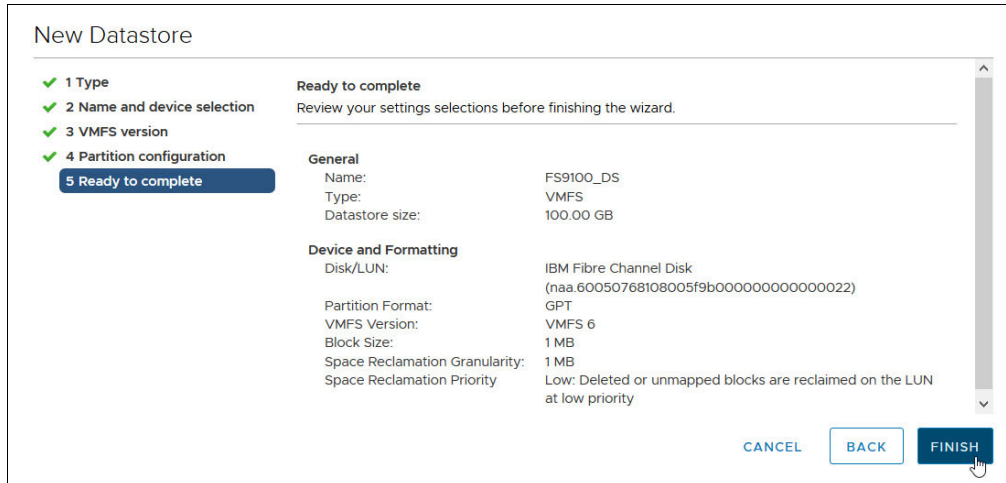


Figure 12-21 New data store - Ready to complete

The preparation steps for the Storage vMotion are complete. To do the Storage vMotion, select **Menu** → **Hosts and Clusters** as shown in Figure 12-22.

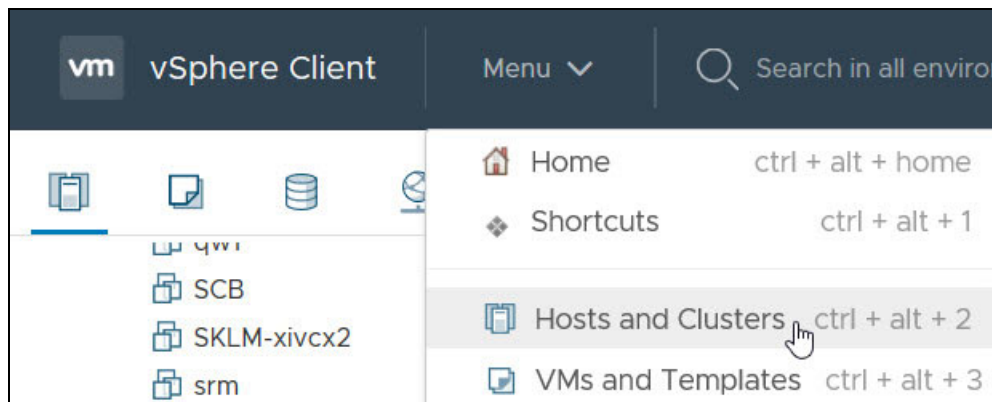


Figure 12-22 vSphere Client Hosts and Clusters

Right-click **WIN2016** VM and select **Migrate** as illustrated in Figure 12-23 on page 202.

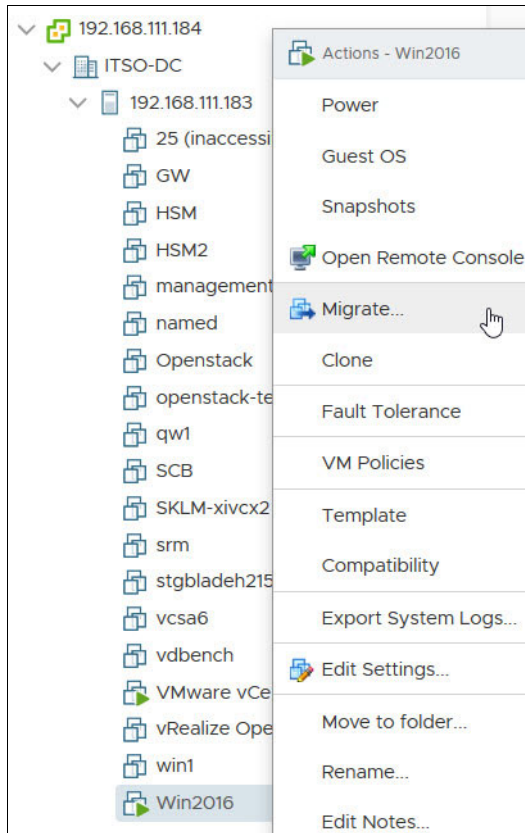


Figure 12-23 vSphere Client Migrate

The Migration wizard opens, select **Change storage only**, which will do a Storage vMotion, click **NEXT** as depicted in Figure 12-24.

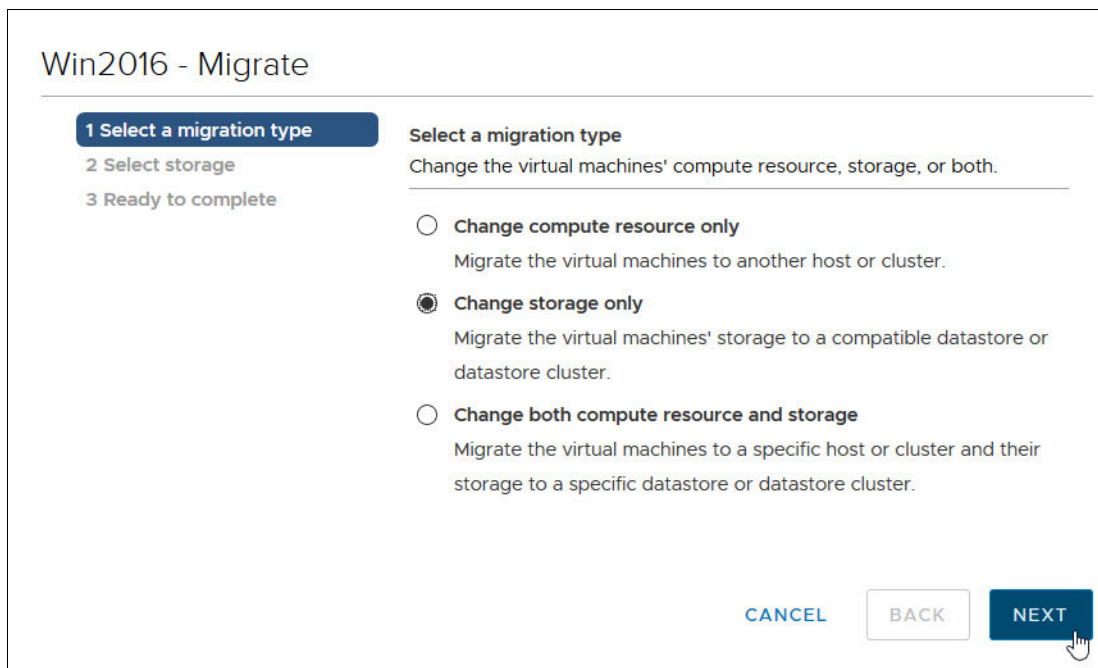


Figure 12-24 vSphere Client Migrate - Type

Select the target data store **FS9100_DS** and click **NEXT** as shown in Figure 12-25.

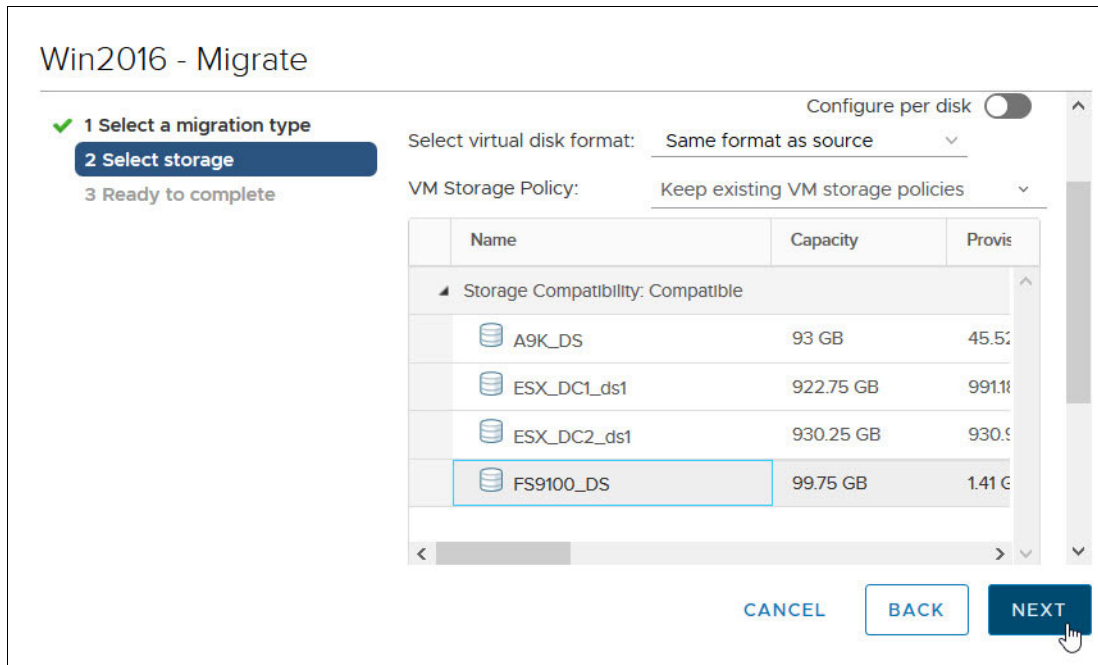


Figure 12-25 vSphere Client Migrate - Select storage

The summary page is shown, click **FINISH** to start the Storage vMotion, as illustrated in Figure 12-26.

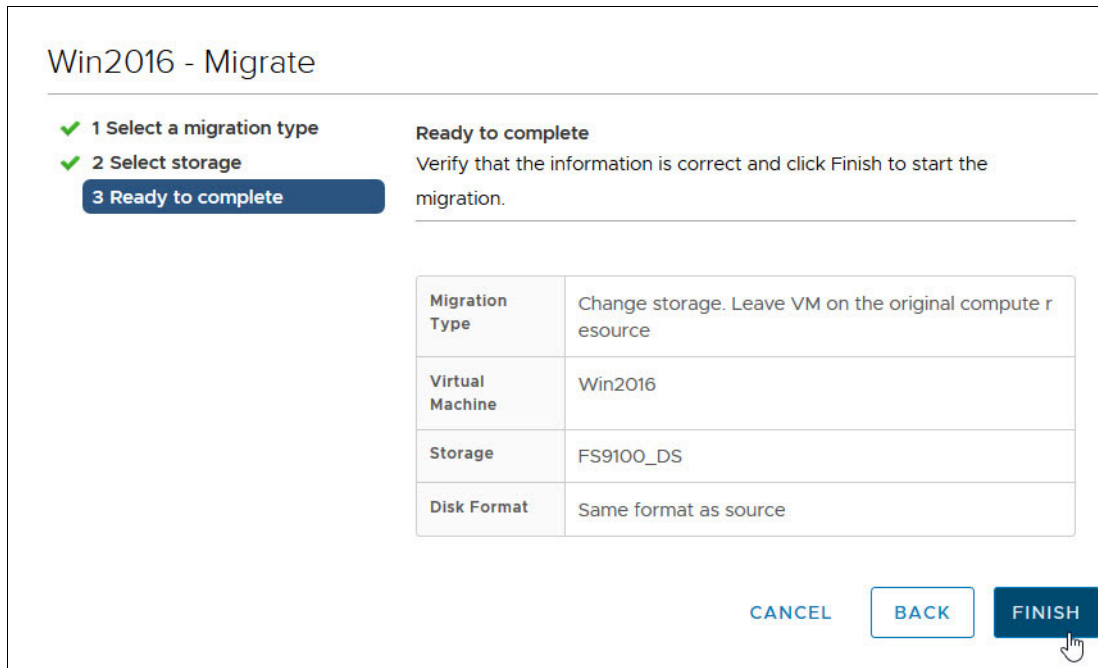


Figure 12-26 vSphere Client Migrate - Ready to complete

The Storage vMotion is running, go to **Recent Tasks** to see the running task as depicted in Figure 12-27.

Recent Tasks		Alarms		
Task Name	Target	Status	Initiator	Queued F
Relocate virtual machine	Win2016	31%	VSPHERE.LOCAL\Administrator	4 ms

Figure 12-27 vSphere Client Recent Tasks

After the Storage vMotion has finished, the Win2016 VM is running on the FlashSystem 9100 data store as shown in Figure 12-28.

The screenshot shows the vSphere Client interface for a Windows 2016 VM. The left-hand navigation pane shows a tree view of the environment, with the VM 'Win2016' selected under the host '192.168.111.183'. The main panel displays the VM's status as 'Powered On' and provides details such as Guest OS (Microsoft Windows Server), Compatibility (ESXi 6.7 and later), and IP Address (192.168.111.183). A yellow warning banner indicates that VMware Tools is not installed. The 'Related Objects' section lists the Host (192.168.111.183), Networks (VM Network), and Storage (FS9100_DS).

Figure 12-28 Windows 2016 VM - After Storage vMotion

The source data store is now either available for reuse or for removal. If you remove all data stores of and the source storage system from the vSphere environment, it is suggested to remove Third-Party Software for your source storage system from vSphere as well.



Monitoring the migration process

Monitoring in a storage environment is crucial and it is part of what usually is called *storage governance*.

With a robust and reliable storage monitoring system, you can save significant money and minimize pain in your operation, by monitoring and predicting usage bottlenecks in your storage environment.

This chapter provides suggestions and the basic concepts of how to implement a storage monitoring system for migrations to the IBM Spectrum Virtualize family of products using specific functions or external IBM tools.

This chapter includes the following sections:

- ▶ Monitoring with the Spectrum Virtualize GUI
- ▶ Performance monitoring
- ▶ Capacity metrics for block storage systems
- ▶ Important performance metrics
- ▶ Performance Monitoring of a host-based migration

Note: Although the following sections discuss various tools and methods to monitor your storage environment, proper planning along with monitoring your storage environment for physical space depletion are critical. Data stored on different storage systems might not get the same space savings between storage products. Be sure to monitor your space usage during your migrations to ensure you do not run out of physical capacity.

13.1 Monitoring with the Spectrum Virtualize GUI

In a Spectrum Virtualize environment, the Spectrum Virtualize management GUI is the primary tool that is used to administer, service and monitor your storage system. The built in performance monitoring, provides an easy way to look at system or node level metrics to help identify impacts during data migrations.

To go to the performance view, go to **Monitoring** → **Performance** in the Spectrum Virtualize GUI main window. Within this window, you will find CPU Utilization, as well as Volumes, Interfaces, and MDisks chart data, that can show either IOPS or MBps statistics. Use the drop-downs along the top to change System versus Node level statistics and the IOPS versus MBps view. The data is shown in 5 second increments, with the last 5 minutes of data shown.

Within each of the 4 charts shown, there are additional options that can be toggled.

CPU Utilization - Shows the percentage of CPU Utilization broken down by system process and compression.

Volumes - Shows the total read and write traffic, along with read and write latency for volumes. Can be shown with IOPS or MBps at either the system or node view.

Interfaces - Shows the total IOPS or MBps for each type of interface used. In the IOPS view, you can see the total Fibre Channel (FC) traffic, iSCSI traffic, serial-attached SCSI (SAS) traffic or IP remote copy traffic. In the MBps view, you also get the total compressed IP remote copy traffic metric.

MDisks - Shows the total read and write traffic, along with read and write latency for at the MDisk level. Can be shown with IOPS or MBps at either the system or node view.

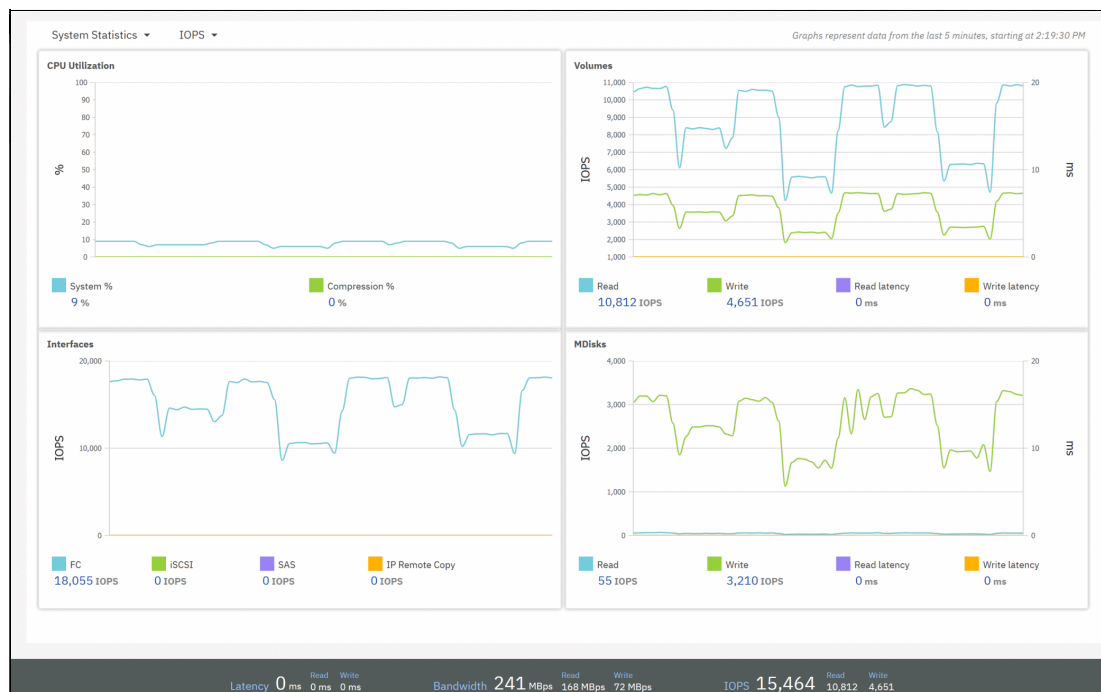


Figure 13-1 Spectrum Virtualize Performance GUI Charts

13.1.1 Monitoring using quotas and alerts

The Spectrum Virtualize GUI shows space usage of storage pools and volume usage for Thin Provisioned or Compressed Volumes that can be monitored by setting some specific quota alerts. Alerts for events can be setup using SNMP, Syslog, SMTP.

See event reporting in IBM Knowledge Center:

https://www.ibm.com/support/knowledgecenter/STSLR9_8.3.1/com.ibm.fs9200_831.doc/svc_svcundererrorreport_17k0nf.html

Storage Pool

During storage pool configuration, you can set a warning such that when the pool capacity reaches this quota setting, an alert is issued. This setting generates a warning when the used capacity in the storage pool first exceeds the specified threshold.

Volumes

Thin provisioned and compressed volumes near their size limits are monitored at specified thresholds to preserve data integrity. If a volume can be shrunk to below the recommended new limit, you are advised to do so. If volume capacity cannot be reduced to meet the recommended limit, you are advised to create a non-compressed mirror of the data (if one does not exist) and delete the primary copy.

13.2 Performance monitoring

Monitoring performance and the ability to collect historical performance metric statistics is almost compulsory for any storage system, and is true for Spectrum Virtualize based storage systems as well.

The next sections show what performance analysis tools are integrated with Spectrum Virtualize systems, and what IBM external tools are available to collect performance statistics to allow historical retention as well.

Tip: Considerations for every migration within any environment might have different metrics to consider and monitor. Understanding prior workloads and performance expectations is a good baseline as additional workload gets put onto a storage system

Performance statistics are useful to debug or prevent some potential bottlenecks.

Use these statistics to check that the new storage system performs as expected and also, if needed, to compare with similar statistics from the system being replaced.

The sections that follow summarize how to monitor your system, using the Spectrum Virtualize GUI, Spectrum Control, or Storage Insights. For details, see IBM Knowledge Center documentation for the respective products.

13.2.1 Performance monitoring with the GUI

In the Spectrum Virtualize GUI, real-time performance statistics provide short-term status information for your systems. The statistics are shown as graphs in the management GUI.

You can use system statistics to monitor the bandwidth of all the volumes, interfaces, and MDisks that are being used on your system. You can also monitor the overall CPU utilization for the system. These statistics summarize the overall performance health of the system and can be used to monitor trends in bandwidth and CPU utilization.

You can monitor changes to stable values or differences between related statistics, such as the latency between volumes and MDisks. These differences can then be further evaluated by performance diagnostic tools.

Additionally, with system-level statistics, you can quickly view bandwidth of volumes, interfaces, and MDisks. Each of these graphs displays the current bandwidth in megabytes per second and a view of bandwidth over time.

Each data point can be accessed to determine its individual bandwidth use and to evaluate whether a specific data point might represent performance impacts. For example, you can monitor the interfaces, such as Fibre Channel or SAS interfaces, to determine whether the host data-transfer rate is different from the expected rate.

The CPU utilization graph shows the current percentage of CPU usage and specific data points on the graph that show peaks in utilization. If compression is being used, you can monitor the amount of CPU resources that are being used for compression and the amount that is available to the rest of the system.

The Volumes and MDisks graphs on the performance window show four metrics: Read, Write, Read latency, and Write latency. You can use these metrics to help determine the overall performance health of the volumes and MDisks on your system. Consistent unexpected results can indicate errors in configuration, system faults, or connectivity issues.

Each graph represents 5 minutes of collected statistics, updated every 5 seconds, and provides a means of assessing the overall performance of your system, as shown in Figure 13-1 on page 206.

You can also obtain a quick overview by using the GUI option **System** → **Dashboard**, as shown in Figure 13-2.

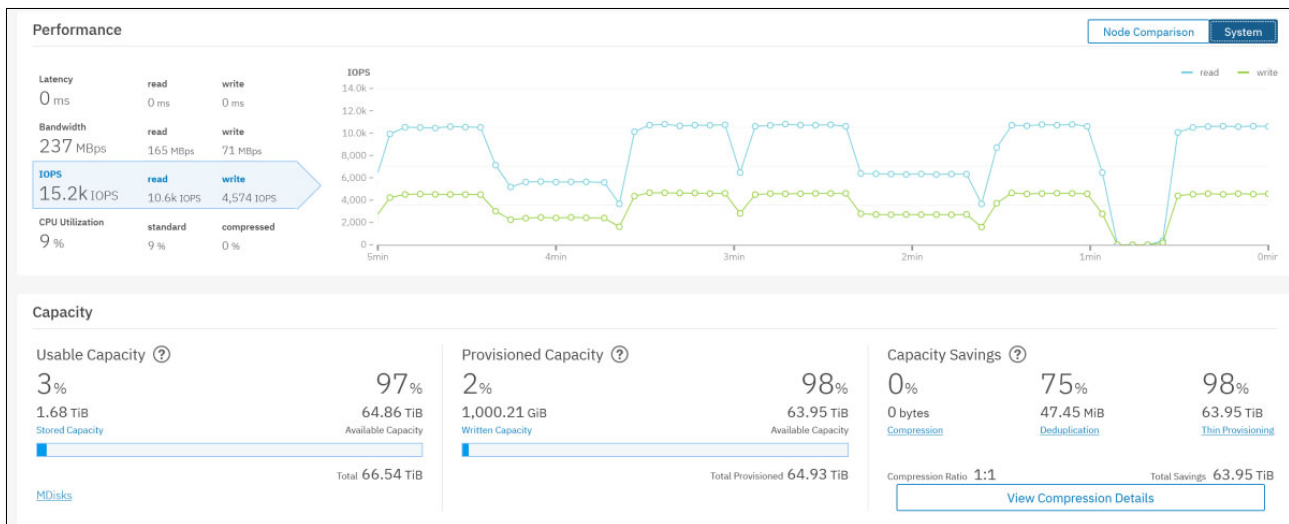


Figure 13-2 System → Dashboard

13.2.2 Performance monitoring with IBM Spectrum Control

IBM Spectrum Control is an on-premises storage management, monitoring, and reporting solution. It leverages the metadata that it collects about supported storage devices to provide services such as custom alerting, analytics, and replication management. Both IBM Spectrum Control and IBM Storage Insights monitor storage systems, but IBM Spectrum Control also monitors hypervisors, fabrics, and switches to provide you with unique analytics and insights into the topology of your storage network. It also provides more granular collection of performance data, with 1-minute intervals rather than the 5-minute intervals in IBM Storage Insights or IBM Storage Insights Pro.

To learn more about the capabilities of IBM Spectrum Control, check out the dedicated IBM Knowledge Center for detailed information at [IBM Spectrum Control documentation](#).

IBM Spectrum Control provides a large amount of detailed information about Spectrum Virtualize systems. The next sections provide some basic suggestions about what metrics need to be monitored and analyzed to debug potential bottleneck problems.

For more information about the installation, configuration, and administration of IBM Spectrum Control (including how to add a storage system), see these websites:

- ▶ IBM Spectrum Control 5.3.6 Limitations and known issues

<https://www.ibm.com/support/pages/536-limitations-and-known-issues-ibm-spectrum-control>

- ▶ Installing IBM Spectrum Control 5.3.6

https://www.ibm.com/support/knowledgecenter/SS5R93_5.3.6/com.ibm.spectrum.sc.doc/fqz0_t_installing_main.html

Note: IBM Spectrum Control 5.3.x or higher is advised for monitoring the IBM FlashSystem family.

13.2.3 Performance monitoring with IBM Storage Insights

IBM Storage Insights is an off-premise, IBM Cloud service that provides cognitive support capabilities, monitoring, and reporting for storage systems. Because it's an IBM Cloud service, getting started is simple and upgrades are handled automatically.

By leveraging the IBM Cloud infrastructure, IBM Support can monitor your storage environment to help minimize the time to resolution of problems and collect diagnostic packages without requiring you to manually upload them.

IBM Storage Insights Key features are:

- ▶ Rapid results when you need them
- ▶ Single pane view across your storage environment
- ▶ Performance analyses at your fingertips
- ▶ Valuable insight from predictive analytics
- ▶ Two editions that meet your needs
- ▶ Simplified, comprehensive and proactive product support

Figure 13-3 shows an IBM Storage Insights example window.

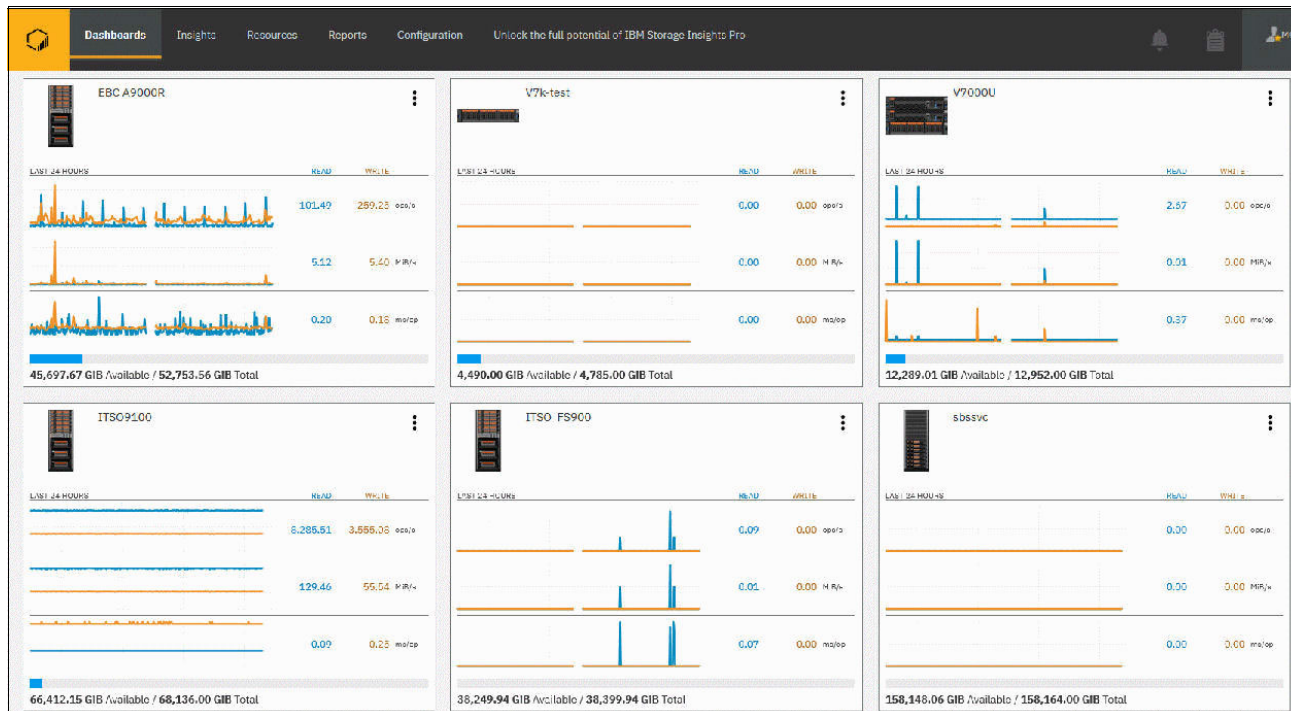


Figure 13-3 IBM Storage Insights

IBM Storage Insights Pro dashboard

The performance dashboard provides Key Performance Indicators for the critical monitoring metrics. These guidelines do not represent the maximum operating limits of the related components, but are rather suggested limits that are selected with an emphasis on maintaining a stable and predictable performance profile.

The dashboard displays the *Last 24 hours* from the active viewing time and date. Selecting an individual element from the chart overlays the corresponding 24 hours for the previous day and seven days prior. This display allows for an immediate historical comparison of the respective metric. The day of reference can also be changed to allow historical comparison of previous days.

These dashboards provide two critical functions:

- ▶ Provides an “at-a-glance” view of all the critical Spectrum Virtualize monitoring metrics.
- ▶ Provides a historical comparison of the current metric profile with previous days that enables rapid detection of anomalous workloads and behaviors.

Key Performance Indicators

You can view the key metrics that are outside of a standard range for storage systems that run Spectrum Virtualize software by using the performance guidelines. The guidelines were established by a historical analysis of storage environments.

Most of the performance charts show an orange line that indicates the best practice value for the metric. These guidelines are established as the levels that allow for a diverse set of workload characteristics while maintaining a stable performance profile. The other lines on each chart represent the measured values for the metric for the resources on your storage system: I/O groups, ports, or canisters.

You can use the lines to compare how close to potentially becoming overloaded your resources are. If your storage system is responding poorly and the charts indicate overloaded resources, you might have to better balance the workload. You can balance the workload between the canisters of the cluster, potentially adding more I/O groups to the cluster, or move some workload to other storage systems.

Figure 13-4 shows the Key Performance Indicators in the Dashboard.

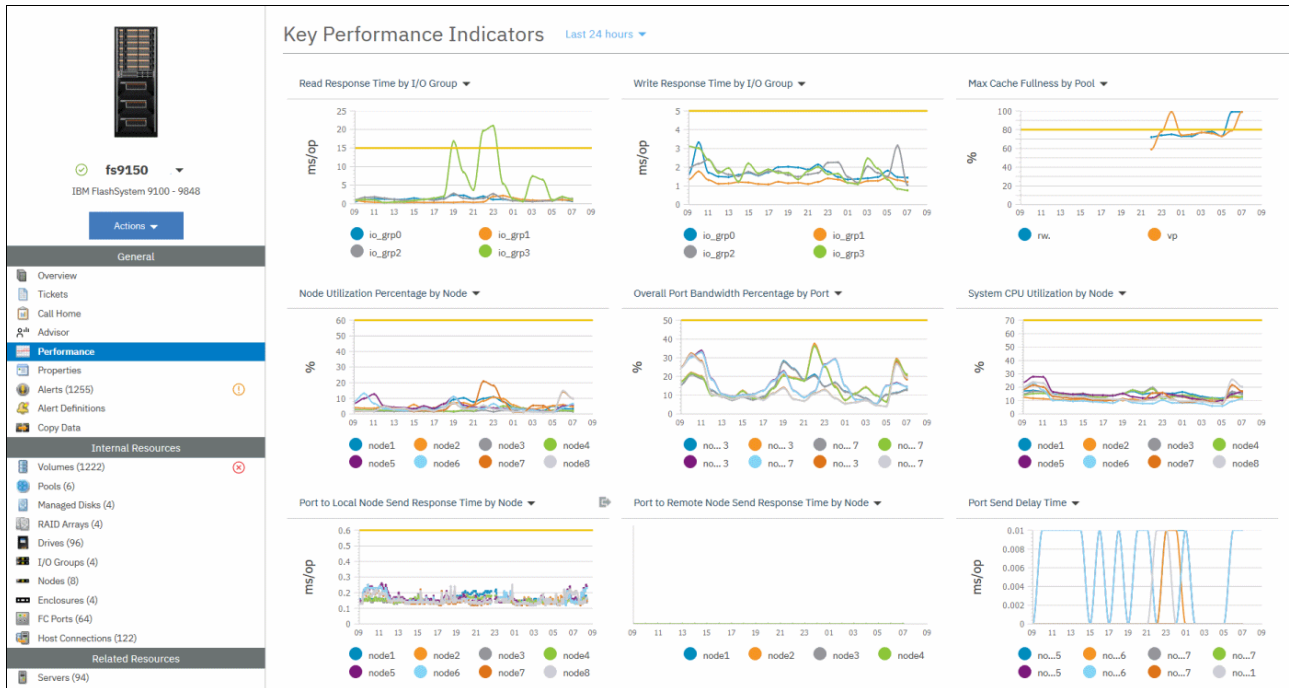


Figure 13-4 DashBoard- Key Performance Indicators

The charts show the hourly performance data measured for each resource on the selected day. Use the following charts to compare the workloads on your storage system with the best practice guidelines:

- ▶ **Max Cache Fullness by Pool:** Add and track cache fullness metrics to identify the pools that are experiencing heavy cache usage. This heavy usage might cause latency issues such as write operations being queued and volumes with slow response times.
- ▶ **Node Utilization Percentage by Node:** Compare the guideline value for this metric, for example, 60% utilization, with the measured value from your system. The average of the bandwidth percentages of those ports in the node that are actively used for host and MDisk send and receive operations. The average is weighted by port speed and adjusted according to the technology limitations of the node hardware. For clusters without FC ports this chart is empty (or when no host I/O is going on).
- ▶ **Overall Port Bandwidth Percentage by Port:** Compare the guideline value for this metric, for example, 50%, with the measured value from your system. Because a cluster can have many ports, the chart shows only the eight ports with the highest average bandwidth over the selected day.
- ▶ **Port Send Delay Time:** The average number of milliseconds of delay that occur on the port for each send operation. The reason for these delays might be a lack of buffer credits. You cannot view zero buffer credit performance metrics for 16 Gbps Fibre Channel ports on resources that run IBM Spectrum Virtualize. Use the Port Send Delay Time metric if the Zero Buffer Credit Timer (8 Gbps) metric is not available.

- ▶ **Port to Local Node Send Response Time by Node:** Compare the guideline value for this metric, for example, 0.6 ms/op, with the measured value from your system. This is an important metric for a good performing cluster.
- ▶ **Port to Remote Node Send Response Time by Node:** Because latencies for copy-services operations can vary widely, a guideline is not established for this metric. Use this chart to identify any discrepancies between the data rates of different nodes.
- ▶ **Read Response Time by I/O Group:** Compare the guideline value for this metric, for example, 15 ms/op, with the measured value from your system. It means, when you see this constantly being exceeded, something might be wrong with the hardware. Other causes remain possible.
- ▶ **System CPU Utilization by Node:** Compare the guideline value for this metric, for example, 70% utilization, with the measured value from your system.
- ▶ **Write Response Time by I/O Group:** Compare the guideline value for this metric, for example, 5 ms/op, with the measured value from your system.

Note: The guidelines are not thresholds, and they are not related to the alerting feature in IBM Storage Insights Pro.

Licensing and editions of IBM Storage Insights

Several editions of IBM Storage Insights enable you to select the capabilities that serve your needs best. Licensing is implemented through different subscription levels.

- ▶ The no initial cost version is called **IBM Storage Insights** and provides a unified view of a storage environment with a diagnostic events feed, an integrated support experience, and key capacity and performance metrics.
- ▶ The **capacity-based**, subscription version is called **IBM Storage Insights Pro** and includes **all** the **features** of IBM Storage Insights plus a more comprehensive view of the performance, capacity, and health of storage resources.

For more information about product comparison, see the IBM Knowledge Center for Storage Insights:

https://www.ibm.com/support/knowledgecenter/SSQRB8/com.ibm.spectrum.si.doc/prd_ovw_versions.html

Tip: Alerts are a good way to be notified of conditions and potential problems that are detected on your storage. If you use IBM Spectrum Control and IBM Storage Insights for IBM Spectrum Control together to enhance your monitoring capabilities, you are advised to define alerts in one of the offerings and not both.

13.3 Capacity metrics for block storage systems

Effective and exact capacity management is based on fundamental knowledge of capacity metrics in each FlashSystem product. Data Reduction pool, Thin Provisioning, compression and deduplication are adding a lot of metrics to the Spectrum Virtualize GUI, Spectrum Control as well as Storage Insights. The capacity metrics in this section are based on Storage Insights.

The Capacity section on the Dashboard provides an overall view of system capacity. This section displays physical capacity, volume capacity, and capacity savings. All of those characteristics are important when planning a storage migration to make sure that you provision enough capacity for the new system.

Physical capacity indicates the total capacity in all storage on the system. Physical capacity includes all the storage the system can virtualize and assign to pools. Physical capacity is displayed in a bar graph and divided into three categories: Stored Capacity, Available Capacity, and Total Physical. If the system supports self-compressing drives, certain system configurations make determining accurate physical capacity on the system difficult.

For example, if the system contains self-compressed drives and data reduction pools without compression enabled, the system cannot determine the accurate amount of physical capacity that is used on the system. In this case, over provisioning and losing access to write operations is possible. If this condition is detected by the system, the Physical Capacity section of the Dashboard page displays a message instead of capacity information.

To recover from this condition, you need to ensure that all thin-provisioned volumes and thin-provisioned volumes that are deduplicated are migrated to volumes with compression enabled in the data reduction pools. Alternatively, you can migrate the volumes to fully allocated volumes and use the drive compression to save capacity.

We will discuss the different values and how they help to determine the capacity utilization, trends in capacity and space usage, and more importantly can prevent an out of space situation in the environment. In order to not run out of space you have to understand the different level of components, such as arrays, pools, system, and so on.

You should understand which limits exist for each of them so that you understand that although one pool might be fine, another pool(s) could run out of storage and just monitoring at the system level is not appropriate if you have two or more pools.

Figure 13-5 shows how to interpret the capacity and savings in a storage environment.

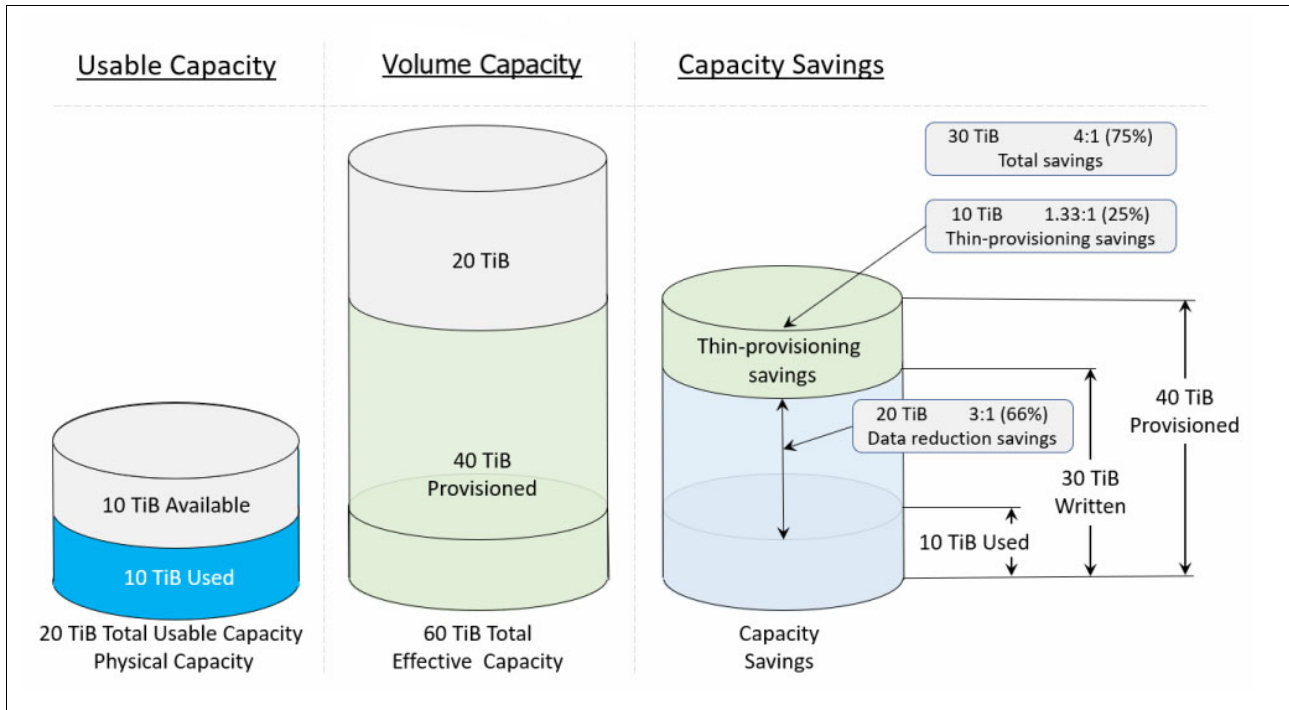


Figure 13-5 Understanding capacity information

Lists of the capacity and space usage metrics that you can add to charts are provided in the Storage system capacity metrics section.

13.3.1 Storage system capacity metrics

For each system in Storage Insights, the Capacity by Pool or Capacity by Volume options can show additional information. Figure 13-6 shows the Capacity, Volume Capacity, and Capacity Savings of an IBM FlashSystem 9100.




Figure 13-6 Capacity Overview

Details by pool or volume can be displayed as shown in the View capacity by pool Figure 13-7 or View capacity by volume Figure 13-8.




Figure 13-7 View capacity by pool

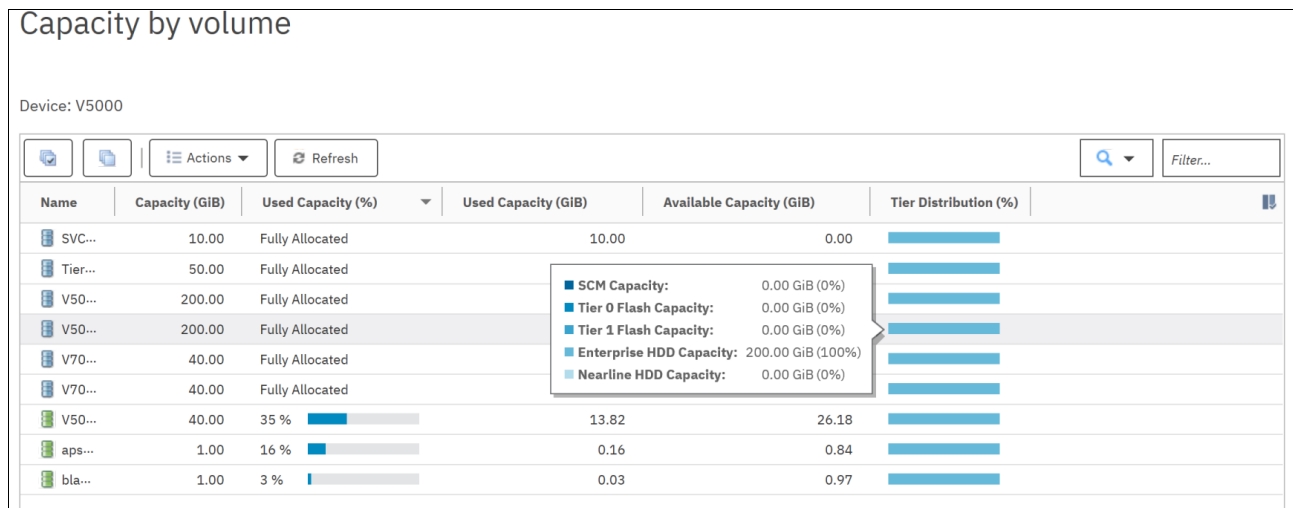


Figure 13-8 View capacity by volume highlighting one volumes Tier Distribution

A complete list of Storage system capacity metrics, Pool capacity metrics, and Volume capacity metrics can be found at the following web page:

https://www.ibm.com/support/knowledgecenter/en/SSQR88/com.ibm.spectrum.si.doc/mgr_capacity_space_metrics_block.html#mgr_capacity_space_metrics_block_VolumeCapacityMetrics

13.4 Important performance metrics

The following metrics are some of the most important metrics that need to be analyzed to understand a performance problem in IBM Spectrum Virtualize products. Those metrics are valid to analyze the front end (by node, by host, or by volume) or the back end (by MDisk or by storage pool):

Terminology: R/W stands for Read and Write operations.

- ▶ **I/O Rate R/W:** The term *I/O* is used to describe any program, operation, or device that transfers data to or from a computer, and to or from a peripheral device. Typically measured in IOPS.
- ▶ **Data Rate R/W:** The data transfer rate (DTR) is the amount of digital data that is moved from one place to another in a specific time. This metric is the amount of data moved from a host to a specific storage device. Typically measured in MBPS.
- ▶ **Response time R/W:** In case of storage Subsystem, this is the time used to complete an I/O operation. Typically measured in ms.
- ▶ **Cache Hit R/W:** This is the percentage of times where read data or write data can be found in cache or, for writes, can find cache free space that it can be written to.
- ▶ **Average Data Block Size R/W:** The block size is the unit of work for the file system. Every read and write is done in full multiples of the block size.
- ▶ **Port-to-Local Node Queue Time (Send):** The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local cluster. A good scenario has less than 1 ms on average.
- ▶ **Port Protocol Errors (Port Send Delay Time):** The average number of milliseconds of delay that occur on the port for each send operation. The reason for these delays might be a lack of buffer credits. Use the Port Send Delay Time metric if the Zero Buffer Credit Timer metric is not available.
- ▶ **Port data rate (send and receive):** The average number of data in MBps for operations in which the port receives or sends data.
- ▶ **Port Protocol Errors (Port Send Delay I/O Percentage):** The percentage of send operations where a delay occurred, relative to the total number of send operations that were measured for the port.
- ▶ **Global Mirror (Overlapping Write Percentage):** The percentage of overlapping write operations that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel, and so they are excluded from this value.

Note: The *Overall Host Attributed Response Time Percentage* is also an important metric, which should be used in conjunction with IBM Spectrum Control V5.3.3 or higher. Previous versions had a calculation error.

Many others metrics are supplied to IBM Spectrum Control from IBM Spectrum Virtualize systems. For more information about all metrics, see:

[Performance metrics for resources that run IBM Spectrum Virtualize](#)

13.5 Performance Monitoring of a host-based migration

For host-based migrations, in addition to performance monitoring of the Spectrum Virtualize system, it is important to monitor the load on the host as well. Depending on the operating system you can monitor performance with host tools like *iostat*, *perfmon*, *nmon*, *topas* or *esxtop*.

The following performance monitoring example was done during a VMware Storage vMotion. The workload was produced by IOmeter running on the local disk of a Windows 2016 VM.

Before VMware Storage vMotion

The Windows VM is running on IBM FlashSystem A9000R volume. The workload of the VM is shown in the IBM Hyper-Scale Manager (HSM) Dashboard as illustrated in Figure 13-9.



Figure 13-9 HSM Dashboard before migration

The load on the A9000R in Storage Insights (no initial cost version) is shown in Figure 13-10.

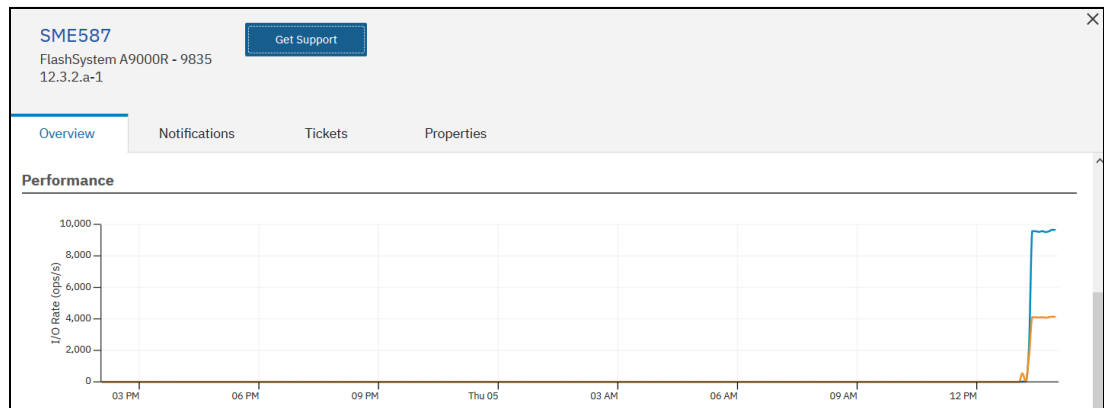


Figure 13-10 Storage Insights GUI: A9000R before migration

The load as shown in *esxtop* on the A9000R volume (*naa.6001738c7c805c230000000000386c0*) is illustrated in Figure 13-11. More details about the usage of *esxtop* can be found on the following web page:

<https://kb.vmware.com/s/article/1008205>

DEVICE	PATH/WORLD/PARTITION	DQLEN	WQLEN	ACTV	QUBED	%USD	LOAD	CMDS/s	READS/s	WRITES/s
mpx.vmhba1:C0:T1:L0	-	31	-	0	0	0	0.00	0.00	0.00	0.00
naa.6001738c7c805c230000000000000000	-	30	-	0	0	0	0.00	0.00	0.00	0.00
naa.6001738c7c805c23000000000000386c0	-	30	-	4	0	13	0.13	13427.54	9365.46	4061.70
naa.600507604085ea381d6f2604b3d51262	-	64	-	0	0	0	0.00	12.40	0.00	12.40
naa.600507604085ea381d6f2604b3d51db6	-	64	-	0	0	0	0.00	0.00	0.00	0.00
naa.60050768108005f9b000000000000022	-	30	-	0	0	0	0.00	50.54	0.00	1.14

Figure 13-11 *esxtop* before migration

During Storage vMotion

After the Storage vMotion has started as described in Chapter 12., “Host-based migrations in VMware Storage vMotion” on page 189, writes to the FlashSystem 9100 volume (naa.60050768108005f9b00000000000022) have started as shown in Figure 13-12. The Storage vMotion increases the workload for the ESXi host, in addition to the normal load the ESXi host must read the written blocks from the A9000R volume and write it to the FlashSystem 9100 volume.

DEVICE	PATH/WORLD/PARTITION	DQLEN	WQLEN	ACTV	QUED	%USD	LOAD	CMDS/s	READS/s	WRITES/s
mpx.vmhbal:C0:T1:L0	-	31	-	0	0	0	0.00	0.00	0.00	0.00
naa.6001738c7c805c230000000000000000	-	30	-	0	0	0	0.00	0.00	0.00	0.00
naa.6001738c7c805c23000000000000386c0	-	30	-	12	0	40	0.40	16168.64	12997.55	3167.18
naa.600507604085ea381d6f2604b3d51262	-	64	-	0	0	0	0.00	14.09	6.85	7.04
naa.600507604085ea381d6f2604b3d51db6	-	64	-	0	0	0	0.00	0.00	0.00	0.00
naa.60050768108005f9b000000000000022	-	30	-	5	0	16	0.17	6097.84	215.58	5742.20

Figure 13-12 esxtop during migration

After Storage vMotion

After the migration has finished, the complete workload runs on the FlashSystem 9100 volume as illustrated in Figure 13-13.

DEVICE	PATH/WORLD/PARTITION	DQLEN	WQLEN	ACTV	QUED	%USD	LOAD	CMDS/s	READS/s	WRITES/s
mpx.vmhbal:C0:T1:L0	-	31	-	0	0	0	0.00	0.00	0.00	0.00
naa.6001738c7c805c230000000000000000	-	30	-	0	0	0	0.00	0.00	0.00	0.00
naa.6001738c7c805c23000000000000386c0	-	30	-	0	0	0	0.00	0.00	0.00	0.00
naa.600507604085ea381d6f2604b3d51262	-	64	-	0	0	0	0.00	33.84	0.00	33.84
naa.600507604085ea381d6f2604b3d51db6	-	64	-	0	0	0	0.00	0.00	0.00	0.00
naa.60050768108005f9b000000000000022	-	30	-	4	0	13	0.13	11852.95	8327.19	3525.37

Figure 13-13 esxtop after migration

In the FlashSystem 9100 (FS9100) GUI the three phases of the migration are visible as depicted in Figure 13-14. In the first phase (before) there are no IOPS to the FS9100 issued, during the second phase (migration) only write IOPS are issued from the ESXi host, in the last phase (after) the FS9100 serves all IOPS (read and write) to the ESXi host.

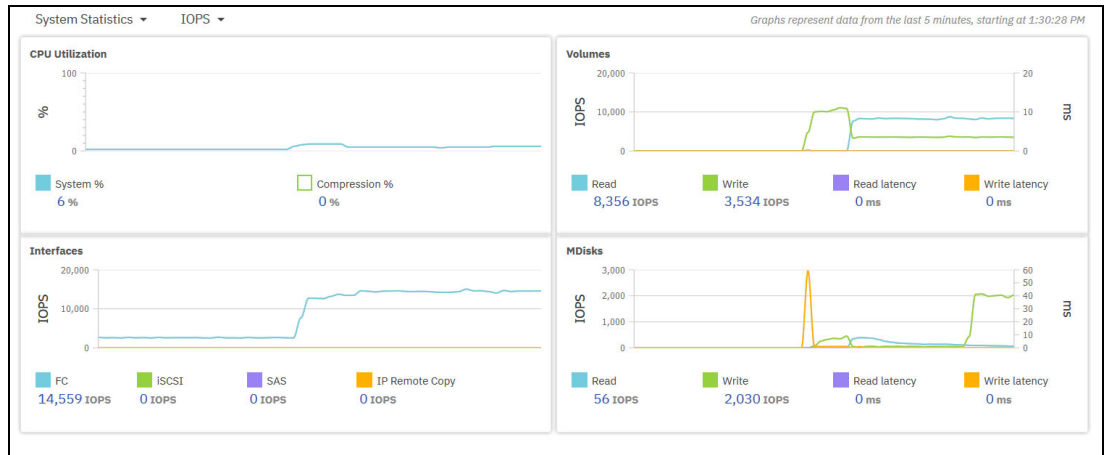


Figure 13-14 FlashSystem 9100 GUI after migration

The workload on FS9100 as visible in Storage Insights is shown in Figure 13-15.

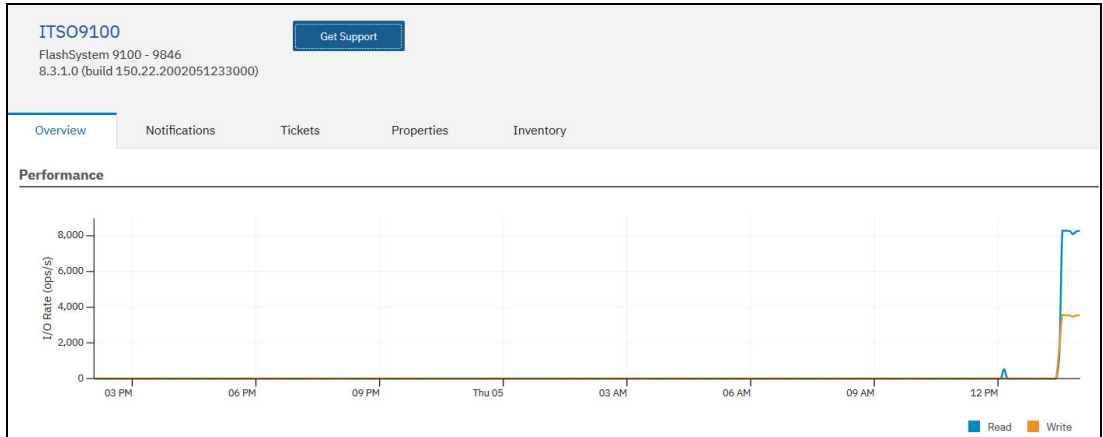


Figure 13-15 Storage Insights GUI: FS9100 after migration

Related publications

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

IBM Redbooks

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

- ▶ *IBM FlashSystem 9100 Architecture, Performance, and Implementation*, SG24-8425
- ▶ *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521
- ▶ *IBM PowerHA SystemMirror V7.2.3 for IBM AIX and V7.22 for Linux*, SG24-8434
- ▶ *IBM Spectrum Virtualize for Public Cloud on AWS Implementation Guide*, REDP-5534
- ▶ *IBM FlashSystem 9200 and 9100 Best Practices and Performance Guidelines*, SG24-8448
- ▶ *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V8.2.1*, SG24-7933
- ▶ *Business Process Management Deployment Guide Using IBM Business Process Manager V8.5*, SG24-8175
- ▶ *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401

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

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- ▶ IBM FlashSystem 9x00 Knowledge Center
<https://www.ibm.com/support/knowledgecenter/STSLR9>
- ▶ IBM SAN Volume Controller Knowledge Center
<https://www.ibm.com/support/knowledgecenter/en/STPVGU>
- ▶ IBM FlashSystem 900 Knowledge Center
<https://www.ibm.com/support/knowledgecenter/STKMQB>
- ▶ IBM FlashSystem A9000R Knowledge Center
<https://www.ibm.com/support/knowledgecenter/en/STJKN5>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

Block Storage Migration in Open Environments

(0.2"spine)
0.17"->0.473"
90->249 pages



SG24-8470-00

ISBN 0738458708

Printed in U.S.A.

Get connected

