



Implementing the Lenovo Storage V3700 V2 and V5030 Systems with IBM Spectrum Virtualize V8.1

Provides an overview of Lenovo Storage V3700 V2, V2 XP and V5030 systems

Introduces configuration setup for the Lenovo Storage V3700 V2, V2 XP and V5030 systems

Explains storage terminologies and concepts for various applications

Describes trouble shooting and monitoring of the Lenovo Storage V3700 V2, V2 XP and V5030





Implementing the Lenovo Storage V3700 V2 and V5030 Systems with IBM Spectrum Virtualize V8.1

January 2020

Note: Before using this information and the product it supports, read the information in “Notices” on page 797.

Last update on January 2020

© Copyright Lenovo 2020. All rights reserved.

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

Contents

Preface	xiii
Comments welcome	xiii
Do you have the latest version?	xiii
Summary of changes	xv
January 2020	xv
September 2018	xv
Chapter 1. Overview of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems	
1	1
1.1 Overview	2
1.2 Terminology	4
1.3 Models	5
1.4 Compatibility	8
1.5 Hardware	9
1.5.1 Control enclosure	10
1.5.2 Lenovo Storage V3700 V2	11
1.5.3 Lenovo Storage V3700 V2 XP	12
1.5.4 Lenovo Storage V5030	13
1.5.5 Expansion enclosure	14
1.5.6 Host interface cards	15
1.5.7 Disk drive types	16
1.6 Terms	16
1.6.1 Hosts	17
1.6.2 Node canister	17
1.6.3 I/O groups	17
1.6.4 Clustered system	18
1.6.5 RAID	18
1.6.6 Managed disks	19
1.6.7 Quorum disks	19
1.6.8 Storage pools	20
1.6.9 Volumes	22
1.6.10 iSCSI	25
1.6.11 Serial-attached SCSI	25
1.6.12 Fibre Channel	25
1.7 Features	26
1.7.1 Mirrored volumes	26
1.7.2 Thin Provisioning	26
1.7.3 Real-time Compression	27
1.7.4 Easy Tier	28
1.7.5 Storage Migration	28
1.7.6 FlashCopy	29
1.7.7 Remote Copy	29
1.7.8 IP replication	30
1.7.9 External virtualization	31
1.7.10 Encryption	31
1.8 Problem management and support	31
1.8.1 Support assistance	31
1.8.2 Event notifications	31

1.8.3	SNMP traps	32
1.8.4	Syslog messages	32
1.8.5	Call Home email	32
1.9	More information resources	33
Chapter 2	Initial configuration	35
2.1	Hardware installation planning	36
2.1.1	Procedure to install the SAS cables	37
2.2	SAN configuration planning	40
2.3	FC direct-attach planning	43
2.4	SAS direct-attach planning	45
2.5	LAN configuration planning	46
2.5.1	Management IP address considerations	47
2.5.2	Service IP address considerations	47
2.6	Host configuration planning	48
2.7	Miscellaneous configuration planning	49
2.8	System management	50
2.8.1	Graphical user interface (GUI)	50
2.8.2	Command-line interface (CLI)	51
2.9	First-time setup	52
2.10	Initial configuration	56
2.10.1	Adding enclosures after the initial configuration	66
2.10.2	Service Assistant Tool	71
Chapter 3	Graphical user interface overview	75
3.1	Overview of management software	76
3.1.1	Access to the storage management software	76
3.1.2	System pane layout	77
3.1.3	Navigation	81
3.1.4	Multiple selection	82
3.1.5	Status indicators area	84
3.2	Overview pane	85
3.3	Monitoring menu	85
3.3.1	System overview	87
3.3.2	System details	89
3.3.3	Events	91
3.3.4	Performance	92
3.3.5	Background Task	93
3.4	Pools menu	93
3.4.1	Pools	95
3.4.2	Child pools	98
3.4.3	Volumes by pool	99
3.4.4	Internal storage	101
3.4.5	External storage	101
3.4.6	MDisks by pools	104
3.4.7	System migration	105
3.5	Volumes menu	106
3.5.1	All volumes	108
3.5.2	Volumes by pool	110
3.5.3	Volumes by host	110
3.6	Hosts menu	111
3.6.1	Hosts	112
3.6.2	Host Clusters	113

3.6.3	Ports by host	115
3.6.4	Host mappings	116
3.6.5	Volumes by host	117
3.7	Copy services	117
3.7.1	FlashCopy	118
3.7.2	Consistency group	119
3.7.3	FlashCopy mappings	121
3.7.4	Remote copy	123
3.7.5	Partnerships	124
3.8	Access menu.	125
3.8.1	Users.	126
3.8.2	Audit Log option	127
3.9	Settings menu	128
3.9.1	Notifications	128
3.9.2	Network.	129
3.9.3	Security	130
3.9.4	System	133
3.9.5	Support	136
3.9.6	GUI preferences	138
Chapter 4.	Storage pools	139
4.1	Working with internal drives	140
4.1.1	Internal Storage window	140
4.1.2	Actions on internal drives	142
4.2	Working with storage pools	150
4.2.1	Creating storage pools	152
4.2.2	Actions on storage pools.	154
4.2.3	Child storage pools	159
4.3	Working with managed disks	162
4.3.1	Assigning managed disks to storage pools.	163
4.3.2	RAID configuration	169
4.3.3	Distributed RAID	170
4.3.4	RAID configuration presets	173
4.3.5	Actions on arrays	174
4.3.6	Actions on external MDisks.	177
4.3.7	More actions on MDisks	184
4.4	Working with external storage controllers	186
Chapter 5.	Host configuration	189
5.1	Host attachment overview.	190
5.2	Planning for direct-attached hosts.	191
5.2.1	Fibre Channel direct attachment to host systems.	191
5.2.2	FC direct attachment between nodes	191
5.3	Preparing the host operating system	191
5.3.1	Windows 2008 R2 and 2012 R2: Preparing for Fibre Channel attachment	192
5.3.2	Windows 2008 R2 and Windows 2012 R2: Preparing for iSCSI attachment . . .	198
5.3.3	Windows 2012 R2: Preparing for SAS attachment	204
5.3.4	VMware ESXi: Preparing for Fibre Channel attachment.	205
5.3.5	VMware ESXi: Preparing for iSCSI attachment	208
5.3.6	VMware ESXi: Preparing for SAS attachment	217
5.4	N-Port Virtualization ID (NPIV) Support	218
5.4.1	NPIV Prerequisites	220
5.4.2	Enabling NPIV on a new system.	221

5.4.3	Enabling NPIV on an existing system	224
5.5	Creating hosts by using the GUI	226
5.5.1	Creating Fibre Channel hosts	228
5.5.2	Configuring the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 for FC connectivity	234
5.5.3	Creating iSCSI hosts	236
5.5.4	Configuring the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 for iSCSI host connectivity	238
5.5.5	Creating SAS hosts	243
5.6	Host Clusters	245
5.6.1	Creating a host cluster	246
5.6.2	Adding a member to a host cluster	249
5.6.3	Listing a host cluster member	251
5.6.4	Assigning a volume to a Host Cluster	252
5.6.5	Remove volume mapping from a host cluster	255
5.6.6	Removing a host cluster member	258
5.6.7	Removing a host cluster	261
5.6.8	I/O throttling for hosts and Host Clusters	263
5.7	Proactive Host Failover	267
Chapter 6.	Volume configuration	269
6.1	Introduction to volumes	270
6.1.1	Image mode volumes	271
6.1.2	Managed mode volumes	272
6.1.3	Cache mode for volumes	273
6.1.4	Mirrored volumes	274
6.1.5	Thin-provisioned volumes	277
6.1.6	Compressed volumes	279
6.1.7	Volumes for various topologies	279
6.2	Create Volumes menu	280
6.3	Creating volumes using the Volume Creation	285
6.3.1	Creating Basic volumes using Volume Creation	285
6.3.2	Creating Mirrored volumes using Volume Creation	288
6.4	Mapping a volume to the host	291
6.5	Creating Custom volumes	293
6.5.1	Creating a custom thin-provisioned volume	294
6.5.2	Creating Custom Compressed volumes	297
6.5.3	Custom Mirrored Volumes	299
6.6	HyperSwap and the mkvolume command	301
6.6.1	Volume manipulation commands	304
6.7	Mapping Volumes to Host after volume creation	306
6.7.1	Mapping newly created volumes to the host using the wizard	306
6.8	Migrating a volume to another storage pool	310
6.9	Migrating volumes using the volume copy feature	313
6.10	I/O throttling	317
6.10.1	Define throttle on a volume	318
6.10.2	Remove a throttle from a volume	320
Chapter 7.	Storage migration	323
7.1	Storage migration wizard overview	324
7.2	Interoperation and compatibility	324
7.3	Storage migration wizard	325
7.3.1	External virtualization capability	325

7.3.2	Model and adapter card considerations	325
7.3.3	Overview of the storage migration wizard	326
7.3.4	Storage migration wizard tasks	327
Chapter 8.	Advanced host and volume administration	349
8.1	Advanced host administration	350
8.1.1	Modifying volume mappings	351
8.1.2	Unmapping volumes from a host	354
8.1.3	Renaming a host	357
8.1.4	Removing a host	360
8.1.5	Host properties	363
8.2	Adding and deleting host ports	367
8.2.1	Adding host port	367
8.2.2	Deleting a host port	371
8.3	Advanced volume administration	373
8.3.1	Advanced volume functions	374
8.3.2	Other actions are available for copies of volumes. For more information, see 8.5, “Advanced volume copy functions” on page 390. Unmapping volumes from all hosts 377	
8.3.3	Viewing which host is mapped to a volume	378
8.3.4	Renaming a volume	378
8.3.5	Shrinking a volume	378
8.3.6	Expanding a volume	380
8.3.7	Migrating a volume to another storage pool	380
8.3.8	Exporting to an image mode volume	381
8.3.9	Deleting a volume	383
8.3.10	Duplicating a volume	383
8.3.11	Adding a volume copy	385
8.4	Volume properties and volume copy properties	386
8.5	Advanced volume copy functions	390
8.5.1	Volume copy: Make Primary	391
8.5.2	Splitting into a new volume	393
8.5.3	Validate Volume Copies option	394
8.5.4	Delete volume copy option	397
8.5.5	Migrating volumes by using the volume copy features	398
8.6	Volumes by storage pool	399
8.7	Volumes by host	400
Chapter 9.	Advanced features for storage efficiency	403
9.1	Introduction	404
9.2	Easy Tier	404
9.2.1	Easy Tier overview	405
9.2.2	Tiered storage pools	406
9.2.3	Easy Tier process	408
9.2.4	I/O Monitoring	409
9.2.5	Data Placement Advisor	409
9.2.6	Data Migration Planner	409
9.2.7	Data Migrator	409
9.2.8	Easy Tier accelerated mode	410
9.2.9	Easy Tier operating modes	411
9.2.10	Easy Tier status	411
9.2.11	Storage Pool Balancing	412
9.2.12	Easy Tier rules	412

9.2.13	Creating multi-tiered pools: Enabling Easy Tier	414
9.2.14	Downloading Easy Tier I/O measurements.	425
9.2.15	Easy Tier I/O Measurement through the command-line interface.	427
9.2.16	IBM Storage Tier Advisor Tool	431
9.2.17	Processing heat log files.	431
9.3	Thin provisioning	432
9.3.1	Configuring a thin provisioned volume	433
9.3.2	Performance considerations	436
9.3.3	Limitations of virtual capacity	437
9.4	Real-time Compression Software	438
9.4.1	Common use cases	438
9.4.2	Real-time Compression concepts	439
9.4.3	Random Access Compression Engine	440
9.4.4	Random Access Compression Engine in stack	444
9.4.5	Data write flow	444
9.4.6	Data read flow.	445
9.4.7	Compression of existing data	445
9.4.8	Configuring compressed volumes.	446
9.4.9	Comprestimator	448
Chapter 10.	Copy services	451
10.1	FlashCopy	452
10.1.1	Business requirements for FlashCopy	452
10.1.2	Backup improvements with FlashCopy	452
10.1.3	Restore with FlashCopy	453
10.1.4	Moving and migrating data with FlashCopy	453
10.1.5	Application testing with FlashCopy	454
10.1.6	Host and application considerations to ensure FlashCopy integrity	454
10.1.7	FlashCopy attributes	454
10.1.8	Reverse FlashCopy	455
10.2	FlashCopy functional overview	456
10.3	Implementing FlashCopy	457
10.3.1	FlashCopy mappings	458
10.3.2	Multiple Target FlashCopy	458
10.3.3	Consistency Groups	459
10.3.4	FlashCopy indirection layer.	462
10.3.5	Grains and the FlashCopy bitmap.	462
10.3.6	Interaction and dependency between multiple target FlashCopy mappings.	464
10.3.7	Summary of the FlashCopy indirection layer algorithm.	465
10.3.8	Interaction with the cache	466
10.3.9	FlashCopy and image mode volumes.	467
10.3.10	FlashCopy mapping events	470
10.3.11	FlashCopy mapping states	472
10.3.12	Thin provisioned FlashCopy	474
10.3.13	Background copy	474
10.3.14	Serialization of I/O by FlashCopy	475
10.3.15	Event handling	476
10.3.16	Asynchronous notifications	477
10.3.17	Interoperation with Metro Mirror and Global Mirror	477
10.3.18	FlashCopy presets	478
10.4	Managing FlashCopy by using the GUI	480
10.4.1	Creating a FlashCopy mapping.	481
10.4.2	Single-click snapshot	493

10.4.3	Single-click clone	495
10.4.4	Single-click backup	496
10.4.5	Creating a FlashCopy Consistency Group	498
10.4.6	Creating FlashCopy mappings in a Consistency Group	499
10.4.7	Showing related volumes	504
10.4.8	Moving a FlashCopy mapping to a Consistency Group	505
10.4.9	Removing a FlashCopy mapping from a Consistency Group	506
10.4.10	Modifying a FlashCopy mapping	507
10.4.11	Renaming FlashCopy mapping	509
10.4.12	Renaming Consistency Group	509
10.4.13	Deleting FlashCopy mapping	510
10.4.14	Deleting FlashCopy Consistency Group	511
10.4.15	Starting FlashCopy process	512
10.4.16	Stopping FlashCopy process	512
10.5	Volume mirroring and migration options	513
10.6	Native IP replication	515
10.6.1	Native IP replication technology	515
10.6.2	Lenovo Storage V series System Layers	516
10.6.3	IP partnership limitations	518
10.6.4	VLAN support	519
10.6.5	IP partnership and terminology	520
10.6.6	States of IP partnership	521
10.6.7	Remote copy groups	522
10.7	Remote Copy	523
10.7.1	Multiple Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems mirroring	523
10.7.2	Importance of write ordering	526
10.7.3	Remote copy intercluster communication	528
10.7.4	Metro Mirror overview	529
10.7.5	Synchronous remote copy	530
10.7.6	Metro Mirror features	531
10.7.7	Metro Mirror attributes	531
10.7.8	Practical use of Metro Mirror	532
10.7.9	Global Mirror Overview	533
10.7.10	Asynchronous remote copy	533
10.7.11	Global Mirror features	535
10.7.12	Using Change Volumes with Global Mirror	537
10.7.13	Distribution of work among nodes	539
10.7.14	Background copy performance	540
10.7.15	Thin-provisioned background copy	540
10.7.16	Methods of synchronization	540
10.7.17	Practical use of Global Mirror	541
10.7.18	Valid combinations of FlashCopy, Metro Mirror, and Global Mirror	541
10.7.19	Remote Copy configuration limits	541
10.7.20	Remote Copy states and events	542
10.8	Consistency protection for Remote and Global mirror	549
10.9	Remote Copy commands	551
10.9.1	Remote Copy process	552
10.9.2	Listing available system partners	552
10.9.3	Changing the system parameters	553
10.9.4	System partnership	554
10.9.5	Creating a Metro Mirror/Global Mirror consistency group	555
10.9.6	Creating a Metro Mirror/Global Mirror relationship	556

10.9.7	Changing Metro Mirror/Global Mirror relationship	556
10.9.8	Changing Metro Mirror/Global Mirror consistency group	556
10.9.9	Starting Metro Mirror/Global Mirror relationship	556
10.9.10	Stopping Metro Mirror/Global Mirror relationship	557
10.9.11	Starting Metro Mirror/Global Mirror consistency group	557
10.9.12	Stopping Metro Mirror/Global Mirror consistency group	558
10.9.13	Deleting Metro Mirror/Global Mirror relationship	558
10.9.14	Deleting Metro Mirror/Global Mirror consistency group	558
10.9.15	Reversing Metro Mirror/Global Mirror relationship	559
10.9.16	Reversing Metro Mirror/Global Mirror consistency group	559
10.10	Managing Remote Copy using the GUI	559
10.10.1	Creating Fibre Channel partnership	560
10.10.2	Creating stand-alone remote copy relationships	562
10.10.3	Creating Consistency Group	570
10.10.4	Renaming Consistency Group	577
10.10.5	Renaming remote copy relationship	578
10.10.6	Moving stand-alone remote copy relationship to Consistency Group	579
10.10.7	Removing remote copy relationship from Consistency Group	580
10.10.8	Starting remote copy relationship	581
10.10.9	Starting remote copy Consistency Group	582
10.10.10	Switching copy direction	583
10.10.11	Switching the copy direction for a Consistency Group	584
10.10.12	Stopping a remote copy relationship	585
10.10.13	Stopping Consistency Group	587
10.10.14	Deleting stand-alone remote copy relationships	588
10.10.15	Deleting Consistency Group	589
10.11	Troubleshooting remote copy	590
10.11.1	1920 error	590
10.11.2	1720 error	592
10.12	HyperSwap	593
10.12.1	Introduction to HyperSwap volumes	594
10.12.2	Failure scenarios	600
10.12.3	Current HyperSwap limitations	604
Chapter 11.	External storage virtualization	607
11.1	Planning for external storage virtualization	608
11.1.1	License for external storage virtualization	608
11.1.2	SAN configuration planning	608
11.1.3	External storage configuration planning	610
11.1.4	Guidelines for virtualizing external storage	610
11.2	Working with external storage	611
11.2.1	Adding external storage	611
11.2.2	Importing image mode volumes	615
11.2.3	Managing external storage	619
11.2.4	Removing external storage	623
Chapter 12.	RAS, monitoring, and troubleshooting	625
12.1	Reliability, availability, and serviceability features	626
12.2	System components	627
12.2.1	Enclosure midplane	627
12.2.2	Node canisters	627
12.2.3	Expansion canisters	637
12.2.4	Disk subsystem	639

12.2.5	Power supply units	643
12.3	Configuration backup	645
12.3.1	Generating a manual configuration backup by using the CLI	646
12.3.2	Downloading a configuration backup by using the GUI	646
12.4	System update	650
12.4.1	Updating node canister software	650
12.4.2	Updating the drive firmware	663
12.5	Monitoring	666
12.5.1	Email notifications and Call Home	667
12.6	Audit log	670
12.7	Event log	671
12.7.1	Managing the event log.	672
12.7.2	Alert handling and recommended actions.	676
12.8	Support Assistance	679
12.8.1	Configuring support assistance.	680
12.8.2	Set up Support Assistant	680
12.8.3	Disable Support Assistance	689
12.9	Collecting support information.	689
12.9.1	Collecting support information by using the GUI.	689
12.9.2	Automatic upload of Support Packages	689
12.9.3	Manual upload of Support Packages	693
12.9.4	Collecting support information by using the SAT	697
12.10	Powering off the system and shutting down the infrastructure	699
12.10.1	Powering off	699
12.10.2	Shutting down and starting up the infrastructure.	703
Chapter 13.	Encryption	705
13.1	Planning for encryption	706
13.2	Defining encryption of data at-rest	706
13.2.1	Encryption methods	709
13.2.2	Encryption keys.	710
13.2.3	Encryption licenses.	711
13.3	Activating encryption.	711
13.3.1	Obtaining an encryption license	711
13.3.2	Start activation process during initial system setup	712
13.3.3	Start activation process on a running system	714
13.3.4	Activate the license automatically.	715
13.3.5	Activate the license manually	717
13.4	Enabling encryption.	719
13.4.1	Starting the Enable Encryption wizard	720
13.4.2	Enabling encryption using USB flash drives	722
13.4.3	Enabling encryption using key servers	726
13.4.4	Enabling encryption using both providers	732
13.5	Configuring additional providers	739
13.5.1	Adding SKLM as a second provider	739
13.5.2	Adding USB flash drives as a second provider	742
13.6	Migrating between providers	743
13.6.1	Migration from USB flash drive provider to encryption key server	743
13.6.2	Migration from encryption key server to USB flash drive provider	744
13.7	Recovering from a provider loss	744
13.8	Using encryption	745
13.8.1	Encrypted pools	745
13.8.2	Encrypted child pools	747

13.8.3 Encrypted arrays	748
13.8.4 Encrypted MDisk	749
13.8.5 Encrypted volumes	751
13.8.6 Restrictions	753
13.9 Rekeying an encryption-enabled system	753
13.9.1 Rekeying using a key server	754
13.9.2 Rekeying using USB flash drives	756
13.10 Migrating between key providers	758
13.11 Disabling encryption	759
Appendix A. CLI setup and SAN Boot	761
Command-line interface	762
Basic setup	762
SAN Boot	774
Enabling SAN Boot for Windows	775
Enabling SAN Boot for VMware	775
Windows SAN Boot migration	775
Appendix B. Terminology	779
Notices	797
Trademarks	798

Preface

Organizations of all sizes face the challenge of managing massive volumes of increasingly valuable data. But storing this data can be costly, and extracting value from the data is becoming more difficult. IT organizations have limited resources but must stay responsive to dynamic environments and act quickly to consolidate, simplify, and optimize their IT infrastructures. The Lenovo® Storage V3700 V2, V2 XP and V5030 systems provide a smarter solution that is affordable, easy to use, and self-optimizing, which enables organizations to overcome these storage challenges.

These storage systems deliver efficient, entry-level configurations that are designed to meet the needs of small and midsize businesses. Designed to provide organizations with the ability to consolidate and share data at an affordable price, the Lenovo Storage V3700 V2, V2 XP and V5030 offer advanced software capabilities that are found in more expensive systems.

This book is intended for pre-sales and post-sales technical support professionals and storage administrators. It applies to the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 with IBM Spectrum Virtualize V8.1.

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 in one of the following ways:

- Use the online feedback form found at the web page for this document:

<http://lenovopress.com/lp0836>

- Send your comments in an email to:

comments@lenovopress.com

Do you have the latest version?

We update our books and papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Summary of changes

This section describes the changes made in this update and in previous updates. These updates might also include minor corrections and editorial changes that are not identified.

January 2020

- ▶ HyperSwap is not available on the Lenovo Storage V3700 V2 and V3700 V2 XP — 10.12, “HyperSwap” on page 593

September 2018

- ▶ Covers IBM Spectrum Virtualize V8.1
- ▶ Updated screenshots and descriptions of the user interfaces
- ▶ New and updated information on encryption
- ▶ New information on storage migration

Overview of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems

This chapter provides an overview of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 architecture and includes a brief explanation of storage virtualization.

Specifically, this chapter provides information about the following topics:

- ▶ 1.1, “Overview” on page 2
- ▶ 1.2, “Terminology” on page 4
- ▶ 1.3, “Models” on page 5
- ▶ 1.4, “Compatibility” on page 8
- ▶ 1.5, “Hardware” on page 9
- ▶ 1.6, “Terms” on page 16
- ▶ 1.7, “Features” on page 26
- ▶ 1.8, “Problem management and support” on page 31
- ▶ 1.9, “More information resources” on page 33

1.1 Overview

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 solution is a modular entry level and midrange storage solution. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 include the capability to virtualize their own internal Redundant Array of Independent Disk (RAID) storage and existing external storage area network (SAN)-attached storage (the Lenovo Storage V5030 only).

The three Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 offer a range of performance scalability and functional capabilities. Table 1-1 shows a summary of the features of these models.

Table 1-1 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models

	Lenovo V3700 V2	Lenovo V3700 V2 XP	Lenovo V5030
CPU cores	2	2	6
Cache	16 GB	Up to 32 GB	Up to 64 GB
Supported expansion enclosures	10	10	20
External storage virtualization	No	No	Yes
Compression	No	No	Yes
Encryption	No	Yes	Yes

For a more detailed comparison, see Table 1-3 on page 5.

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 feature the following benefits:

- ▶ Enterprise technology available to entry and midrange storage
- ▶ Expert administrators are not required
- ▶ Easy client setup and service
- ▶ Simple integration into the server environment
- ▶ Ability to grow the system incrementally as storage capacity and performance needs change

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 address the block storage requirements of small and midsize organizations. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 consist of one 2U control enclosure and, optionally, up to ten 2U expansion enclosures on the Lenovo Storage V3700 V2 and Lenovo Storage V3700 V2 XP systems and up to twenty 2U expansion enclosures on the Lenovo Storage V5030 systems. The Lenovo Storage V5030 systems are connected by serial-attached Small Computer Systems Interface (SCSI) (SAS) cables that make up one system that is called an *I/O group*.

With the Lenovo Storage V5030 systems, two I/O groups can be connected to form a cluster, providing a maximum of two control enclosures and 40 expansion enclosures. With the High Density expansion drawers you will be able to attach up to 16 expansion enclosures to a cluster.

The control and expansion enclosures are available in the following form factors, and they can be intermixed within an I/O group:

- ▶ 12 x 3.5-inch (8.89-centimeter) drives in a 2U unit
- ▶ 24 x 2.5-inch (6.35-centimeter) drives in a 2U unit

Two canisters are in each enclosure. Control enclosures contain two node canisters, and expansion enclosures contain two expansion canisters.

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support up to 1008 x 2.5 inch or 504 x 3.5 inch drives or a combination of both drive form factors for the internal storage in a two I/O group Lenovo Storage V5030 cluster.

SAS, Nearline (NL)-SAS, and solid-state drive (SSD) types are supported.

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are designed to accommodate the most common storage network technologies to enable easy implementation and management. It can be attached to hosts through a Fibre Channel (FC) SAN fabric, an Internet Small Computer System Interface (iSCSI) infrastructure, or SAS. Hosts can be attached directly or through a network.

Important: For more information about supported environments, configurations, and restrictions, see the Lenovo interoperability matrix, which is available at this web page:

<http://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v3700v2/6535/documentation>

<http://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v5030/6536/documentation>

For more information, see this web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/svc_inst_allplan_22qgvs.html

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are a virtualized storage solutions that groups their internal drives into RAID arrays, which are called managed disks (MDisks). MDisk can also be created on the Lenovo Storage V5030 systems by importing logical unit numbers (LUNs) from external FC SAN-attached storage. These MDisk are then grouped into storage pools. Volumes are created from these storage pools and provisioned out to hosts.

Storage pools are normally created with MDisk of the same drive type and drive capacity. Volumes can be moved non-disruptively between storage pools with differing performance characteristics. For example, a volume can be moved between a storage pool that is made up of NL-SAS drives to a storage pool that is made up of SAS drives to improve performance.

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems also provide several configuration options to simplify the implementation process. It also provides configuration presets and automated wizards that are called *Directed Maintenance Procedures* (DMP) to help resolve any events that might occur.

Included with an Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems are a simple and easy to use graphical user interface (GUI) to allow storage to be deployed quickly and efficiently. The GUI runs on any supported browser. The management GUI contains a series of preestablished configuration options that are called *presets* that use commonly used settings to quickly configure objects on the system. Presets are available for creating volumes and IBM FlashCopy mappings and for setting up a RAID configuration.

You can also use the command-line interface (CLI) to set up or control the system.

1.2 Terminology

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems use terminology that is consistent with the entire IBM Storwize for Lenovo family. The terms are defined in Table 1-2. More terms can be found in Appendix B, “Terminology” on page 779.

Table 1-2 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 terminology

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 term	Definition
Battery	Each control enclosure node canister in an Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 contain a battery.
Chain	Each control enclosure has either one or two chains, which are used to connect expansion enclosures to provide redundant connections to the inside drives.
Clone	A copy of a volume on a server at a particular point. The contents of the copy can be customized and the contents of the original volume are preserved.
Control enclosure	A hardware unit that includes a chassis, node canisters, drives, and power sources.
Data migration	Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can migrate data from existing external storage to its internal volumes.
Distributed RAID (DRAID)	No dedicated spare drives are in an array. The spare capacity is distributed across the array, which allows faster rebuild of the failed disk.
Drive	Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support a range of hard disk drives (HDDs) and Flash Drives.
Event	An occurrence that is significant to a task or system. Events can include the completion or failure of an operation, a user action, or the change in the state of a process.
Expansion canister	A hardware unit that includes the SAS interface hardware that enables the control enclosure hardware to use the drives of the expansion enclosure. Each expansion enclosure has two expansion canisters.
Expansion enclosure	A hardware unit that includes expansion canisters, drives, and power supply units.
External storage	MDisks that are SCSI logical units (LUs) that are presented by storage systems that are attached to and managed by the clustered system.
Fibre Channel port	Fibre Channel ports are connections for the hosts to get access to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 .
Host mapping	The process of controlling which hosts can access specific volumes within an Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 .
Internal storage	Array MDisks and drives that are held in enclosures that are part of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 .
iSCSI (Internet Small Computer System Interface)	Internet Protocol (IP)-based storage networking standard for linking data storage facilities.

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 term	Definition
Managed disk (MDisk)	A component of a storage pool that is managed by a clustered system. An MDisk is part of a RAID array of internal storage or a SCSI LU for external storage. An MDisk is not visible to a host system on the SAN.
Node canister	A hardware unit that includes the node hardware, fabric, and service interfaces, SAS expansion ports, and battery. Each control enclosure contains two node canisters.
PHY	A single SAS lane. Four PHYs are in each SAS cable.
Power Supply Unit	Each enclosure has two power supply units (PSU).
Quorum disk	A disk that contains a reserved area that is used exclusively for cluster management. The quorum disk is accessed when it is necessary to determine which half of the cluster continues to read and write data.
Serial-Attached SCSI (SAS) ports	SAS ports are connections for expansion enclosures and direct attachment of hosts to access the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 .
Snapshot	An image backup type that consists of a point-in-time view of a volume.
Storage pool	An amount of storage capacity that provides the capacity requirements for a volume.
Strand	The SAS connectivity of a set of drives within multiple enclosures. The enclosures can be control enclosures or expansion enclosures.
Thin provisioning or thin provisioned	The ability to define a storage unit (full system, storage pool, or volume) with a logical capacity size that is larger than the physical capacity that is assigned to that storage unit.
Traditional RAID (TRAIID)	Traditional Raid is uses the standard RAID levels.
Volume	A discrete unit of storage on disk, tape, or other data recording medium that supports a form of identifier and parameter list, such as a volume label or input/output control.
Worldwide port names	Each Fibre Channel port and SAS port is identified by its physical port number and worldwide port name (WWPN).

1.3 Models

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 platform consist of different models. Each model type supports a different set of features, as shown in Table 1-3.

Table 1-3 IBM Storwize V5000for Lenovo and Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 feature comparison

Feature	IBM V5000 for Lenovo	Lenovo V3700 V2	Lenovo V3700 V2 XP	Lenovo V5030
Cache	16 GB	16 GB	16 GB or 32 GB	32 GB or 64 GB

Feature	IBM V5000 for Lenovo	Lenovo V3700 V2	Lenovo V3700 V2 XP	Lenovo V5030
CPU	4 - core Ivy Bridge Xeon CPU 2GHz	2- core Broadwell-DE Celeron CPU 1.2GHz	2- core Broadwell-DE Xeon CPU 2.2GHz Hyper-threading	6 - core Broadwell-DE Xeon CPU 1.9GHz Hyper-threading
Compression	None	None	None	Licensed (with 64 GB cache only)
DRAID	Yes	Yes	Yes	Yes
SAS HW Encryption	None	None	Licensed	Licensed
External Virtualization	Licensed	Data Migration Only	Data Migration Only	Licensed
IBM Easy Tier	Licensed	Licensed	Licensed	Licensed
FlashCopy	Licensed	Licensed	Licensed	Licensed
Hyperswap	Yes	No	No	Yes
Remote Copy	Licensed	Licensed	Licensed	Licensed
Thin Provisioning	Yes	Yes	Yes	Yes
Traditional RAID	Yes	Yes	Yes	Yes
Volume Mirroring	Yes	Yes	Yes	Yes
VMware Virtual Volumes (VVols)	Yes	Yes	Yes	Yes

More information: For more information about the features, benefits, and specifications of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models, see the Lenovo Press product guides:

<https://lenovopress.com/lp0497-lenovo-storage-v3700-v2-and-v3700-v2-xp>

<https://lenovopress.com/lp0498-lenovo-storage-v5030>

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models are described in Table 1-4. All control enclosures have two node canisters. F models are expansion enclosures.

Table 1-4 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models

Model	Description	Cache	Drive Slots
One-year warranty			
6535-HC1	Lenovo V3700 V2large form factor (LFF) Control Enclosure	16 GB	12 x 3.5-inch
6535-HC4	Lenovo V3700 V2 small form factor (SFF) Control Enclosure	16 GB	24 x 2.5-inch
6535-HC2	Lenovo V3700 V2 XP LFF Control Enclosure	16 GB or 32 GB	12 x 3.5-inch
6535-HC5	Lenovo V3700 V2 XP SFF Control Enclosure	16 GB or 32 GB	24 x 2.5-inch

Model	Description	Cache	Drive Slots
6536-HC3	Lenovo V5030 LFF Control Enclosure	32 GB or 64 GB	12 x 3.5-inch
6536-HC6	Lenovo V5030 SFF Control Enclosure	32 GB or 64 GB	24 x 2.5-inch
6536-HC6	IBM Storwize V5030F All-Flash Array Control Enclosure	64GB	24 x 2.5-inch
6535-HC7	Lenovo V3700 V2 LFF Expansion Enclosure	N/A	12 x 3.5-inch
6535-HC8	Lenovo V3700 V2 SFF Expansion Enclosure	N/A	24 x 2.5-inch
6536-HC8	Lenovo V5030 SFF Expansion Enclosure	N/A	24 x 2.5-inch
6536-HC7	Lenovo Storage V5030 LFF Expansion Enclosure	N/A	12 x 3.5-inch
6536-HC8	Lenovo Storage V5030 SFF Expansion Enclosure	N/A	24 x 2.5-inch

The Lenovo Storage V5030 systems can be added to an existing IBM Storwize V5000 for Lenovo cluster to form a two-I/O group configuration. This configuration can be used as a migration mechanism to upgrade from the IBM Storwize V5000 for Lenovo to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. The IBM Storwize V5000 for Lenovo models are described in Table 1-5 for completeness.

Table 1-5 IBM Storwize V5000 for Lenovo models

Model	Cache	Drive slots
6194-12C	16 GB	12 x 3.5-inch
6194-24C	16 GB	24 x 2.5-inch
6194-12E	N/A	12 x 3.5-inch
6194-24E	N/A	24 x 2.5-inch

Figure 1-1 shows the front view of the LFF(12 x 3.5-inch) enclosures.



Figure 1-1 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 front view for 6535/6536 LFF(12 x 3.5 inch) enclosures

The drives are positioned in four columns of three horizontally mounted drive assemblies. The drive slots are numbered 1 - 12, starting at the upper left and moving left to right, top to bottom.

Figure 1-2 shows the front view of the SFF(24 x 2.5-inch) enclosures.




Figure 1-2 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 front view for 6535/6536 SFF(24 x 2.5-inch) enclosure

The drives are positioned in one row of 24 vertically mounted drive assemblies. The drive slots are numbered 1 - 24, starting from the left. A vertical center drive bay molding is between slots 12 and 13.

1.4 Compatibility

The Lenovo Storage V5030 system can be added into existing IBM Storwize V5000 for Lenovo clustered systems. All systems within a cluster must use the same version of IBM Storwize V5000 for Lenovo software, which is version 7.6.1 or later.

Restriction: The Lenovo Storage V3700 V2 and V3700 V2 XP are not compatible with IBM Storwize V5000 for Lenovo system as they are not able to join an existing I/O group.

A single Lenovo Storage V5030 control enclosure can be added to a single IBM Storwize V5000 for Lenovo cluster to bring the total number of I/O groups to two. They can be clustered by using either Fibre Channel or Fibre Channel over Ethernet (FCoE). The possible I/O group configuration options for all IBM Storwize V5000 for Lenovo and Lenovo V3700 V2, V3700 V2 XP and V5030 models are shown in Table 1-6.

Table 1-6 I/O group configurations

I/O group 0	I/O group 1
Lenovo Storage V3700 V2	N/A
Lenovo Storage V3700 V2 XP	N/A
Lenovo Storage V5030	N/A
Lenovo Storage V5030	Lenovo Storage V5030
Lenovo Storage V5030	IBM Storwize V5000 for Lenovo
IBM Storwize V5000for Lenovo	Lenovo Storage V5030
IBM Storwize V5000 for Lenovo	N/A
IBM Storwize V5000 for Lenovo	IBM Storwize V5000 for Lenovo

1.5 Hardware

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 solution is a modular storage system that is built on a common enclosure platform that is shared by the control enclosures and expansion enclosures.

Figure 1-3 shows an overview of hardware components of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 solution.

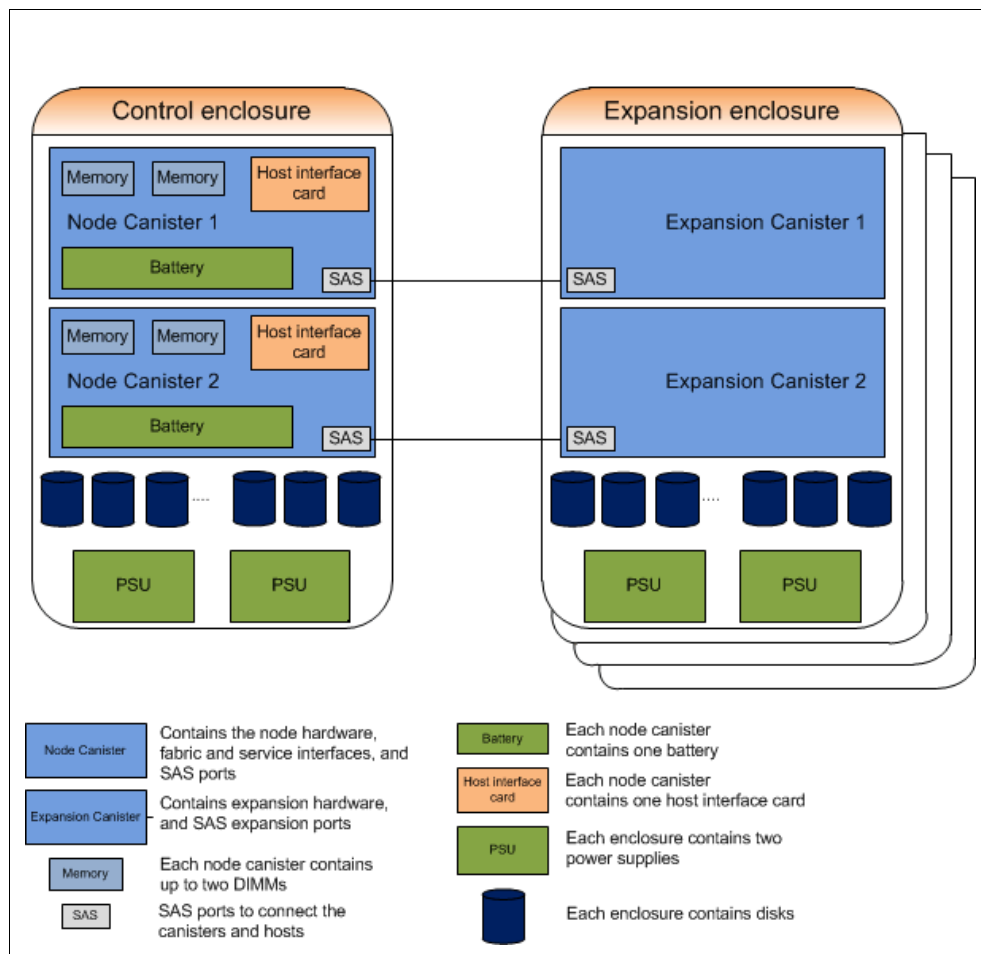


Figure 1-3 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 hardware components

Figure 1-4 shows the control enclosure rear view of a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 enclosure.



Figure 1-4 Lenovo V3700 V2 XP control enclosure rear view

In Figure 1-4 on page 9, you can see two power supply slots at the bottom of the enclosure. The power supplies are identical and exchangeable. Two canister slots are at the top of the chassis.

In Figure 1-5, you can see the rear view of an Lenovo Storage V3700, V3700 V2 XP, and V5030 expansion enclosure.




Figure 1-5 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 expansion enclosure rear view

You can see that the only difference between the control enclosure and the expansion enclosure is the canister. The canisters of the expansion enclosure have only two SAS ports.

For more information about the expansion enclosure, see 1.5.5, “Expansion enclosure” on page 14.

1.5.1 Control enclosure

Each Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems has one control enclosure that contains two node canisters (nodes), disk drives, and two power supplies.

The two node canisters act as a single processing unit and form an I/O group that is attached to the SAN fabric, an iSCSI infrastructure, or that is directly attached to hosts through FC or SAS. The pair of nodes is responsible for serving I/O to a volume. The two nodes provide a highly available fault-tolerant controller so that if one node fails, the surviving node automatically takes over. Nodes are deployed in pairs that are called *I/O groups*.

One node is designated as the configuration node, but each node in the control enclosure holds a copy of the control enclosure state information.

The Lenovo Storage V3700 V2 and Lenovo Storage V3700 V2 XP support a single I/O group. The Lenovo Storage V5030 supports two I/O groups in a clustered system.

The terms *node canister* and *node* are used interchangeably throughout this book.

The battery is used if power is lost. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems use this battery to power the canister while the cache data is written to the internal system flash. This memory dump is called a *fire hose memory dump*.

Note: The batteries of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are able to process two fire hose memory dump in a row. After this you will not be able to power up the system immediately. There is a need to wait until the batteries are charged over a level which allows them to run the next fire hose memory dump.

After the system is up again, this data is loaded back to the cache for destaging to the disks.

1.5.2 Lenovo Storage V3700 V2

Figure 1-6 shows a single Lenovo Storage V3700 V2 node canister.



Figure 1-6 Lenovo Storage V3700 V2 node canister

Each Lenovo Storage V3700 V2 node canister contains the following hardware:

- ▶ Battery
- ▶ Memory: 8 GB
- ▶ One 12 Gbps SAS port
- ▶ Two 10/100/1000 Mbps Ethernet ports
- ▶ One USB 2.0 port that is used to gather system information
- ▶ System flash
- ▶ Host interface card (HIC) slot (different options are possible)

Figure 1-6 shows the following features that are provided by the Lenovo Storage V3700 V2 node canister:

- ▶ Two 10/100/1000 Mbps Ethernet ports. Port 1 must be used for management, and port 2 can optionally be used for management. Port 2 serves as a technician port (as denoted by the white box with “T” in it) for system initialization and service.

Note: All three models use a technician port to perform initial setup. The implementation of the technician port varies between models: On Lenovo Storage V3700 V2 and V3700 V2 XP the second 1GbE port (labelled T) is initially enabled as a technician port. After cluster creation this port is disabled and can then be used for I/O and/or management. On Lenovo Storage V5030 the onboard 1GbE port (labelled T) is permanently enabled as a technician port. Connecting the technician port to the LAN will disable the port. The Lenovo Storage V3700 V2 and V3700 V2 XP technician port can be re-enabled after initial setup.

Commands used to enable / disable the technical port:

```
satask chserviceip -techport enable -force
```

```
satask chserviceip -techport disable
```

- ▶ Both ports can be used for iSCSI traffic and IP replication. For more information, see Chapter 5, “Host configuration” on page 189 and Chapter 10, “Copy services” on page 451.
- ▶ One USB port for gathering system information.

System initialization: Unlike the IBM Storwize V5000 for Lenovo, you must perform the system initialization of the Lenovo V3700 V2 by using the technician port instead of the USB port.

- One 12Gbs serial-attached SCSI (SAS 3.0) port to connect to the optional expansion enclosures. The Lenovo Storage V3700 V2 supports up to 10 expansion enclosures.

Important: The canister SAS port on the Lenovo Storage V3700 V2 does not support SAS host attachment. The Lenovo Storage V3700 V2 supports SAS hosts by using an optional host interface card. See 1.5.6, “Host interface cards” on page 15.

Do not use the port that is marked with a wrench. *This port is a service port only.*

1.5.3 Lenovo Storage V3700 V2 XP

Figure 1-7 shows a single Lenovo Storage V3700 V2 XP node canister.




Figure 1-7 Lenovo Storage V3700 V2 XP node canister

Each node canister contains the following hardware:

- Battery
- Memory: 8 GB upgradable to 16 GB
- Three 12 Gbps SAS ports)
- Two 10/100/1000 Mbps Ethernet ports
- One USB 2.0 port that is used to gather system information
- System flash
- HIC slot (different options are possible)

Figure 1-7 shows the following features that are provided by the Lenovo Storage V3700 V2 XP node canister:

- Two 10/100/1000 Mbps Ethernet ports. Port 1 must be used for management, and port 2 can optionally be used for management. Port 2 serves as a technician port (as denoted by the white box with “T” in it) for system initialization and service.

Note: All three models use a technician port to perform initial setup. The implementation of the technician port varies between models: On Lenovo Storage V3700 V2 and V3700 V2 XP the second 1GbE port (labelled T) is initially enabled as a technician port. After cluster creation this port is disabled and can then be used for I/O and/or management. On Lenovo Storage V5030 the onboard 1GbE port (labelled T) is permanently enabled as a technician port. Connecting the technician port to the LAN will disable the port. The Lenovo Storage V3700 V2 and V3700 V2 XP technician port can be re-enabled after initial setup.

Commands used to enable / disable the techport:

```
satask chserviceip -techport enable -force  
satask chserviceip -techport disable
```

- Both ports can be used for iSCSI traffic and IP replication. For more information, see Chapter 5, “Host configuration” on page 189 and Chapter 10, “Copy services” on page 451

- One USB port for gathering system information.

System initialization: Unlike the IBM Storwize V5000 for Lenovo, you must perform the system initialization of the Lenovo Storage V3700 V2 XP by using the technician port instead of the USB port.

- Three 12Gbps serial-attached SCSI (SAS 3.0) ports. The ports are numbered 1 - 3 from left to right. Port 1 is used to connect to the optional expansion enclosures. Ports 2 and 3 can be used to connect directly to SAS hosts. (Both 6G and 12G hosts are supported.) The Lenovo Storage V3700 V2 XP supports up to 10 expansion enclosures.

Service port: Do not use the port that is marked with a wrench. This port is a service port only.

1.5.4 Lenovo Storage V5030

Figure 1-8 shows a single Lenovo Storage V5030 node canister.



Figure 1-8 Lenovo Storage V5030 node canister

Each node canister contains the following hardware:

- Battery
- Memory: 16 GB upgradable to 32 GB
- Two 12 Gbps SAS ports
- One 10/100/1000 Mbps Ethernet technician port
- Two 1/10 Gbps Ethernet ports
- One USB 2.0 port that is used to gather system information
- System flash
- HIC slot (different options are possible)

Figure 1-8 shows the following features that are provided by the Lenovo Storage V5030 node canister:

- One Ethernet technician port (as denoted by the white box with “T” in it). This port can be used for system initialization and service only. For more information, see Chapter 1, “Overview of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems” on page 1. It cannot be used for anything else.
- Two 1/10 Gbps Ethernet ports. These ports are Copper 10GBASE-T with RJ45 connectors. Port 1 must be used for management. Port 2 can optionally be used for management. Both ports can be used for iSCSI traffic and IP replication. For more information, see Chapter 5, “Host configuration” on page 189 and Chapter 10, “Copy services” on page 451.

Important: The 1/10 Gbps Ethernet ports do not support speeds less than 1 Gbps (100 Mbps is not supported).

Ensure that you use the correct port connectors. The Lenovo Storage V5030 canister 10 Gbps connectors appear the same as the 1 Gbps connectors on the other IBM Storwize V5000 for Lenovo models. These RJ45 connectors differ from the optical small form-factor pluggable (SFP+) connectors on the optional 10 Gbps HIC. When you plan to implement the Lenovo Storage V5030, ensure that any network switches provide the correct connector type.

- One USB port to gather system information.

System initialization: Unlike the IBM Storwize V5000 for Lenovo, you must perform the system initialization of the Lenovo Storage V5030 by using the technician port instead of the USB port.

- Two 12Gbps serial-attached SCSI (SAS 3.0) ports. The ports are numbered 1 and 2 from left to right to connect to the optional expansion enclosures. The Lenovo Storage V5030 supports up to 20 expansion enclosures. Ten expansion enclosures can be connected to each port.

Important: The canister SAS ports on the Lenovo Storage V5030 do not support SAS host attachment. The Lenovo Storage V5030 supports SAS hosts by using an HIC. See 1.5.6, “Host interface cards” on page 15.

Do not use the port that is marked with a wrench. This port is a service port only.

1.5.5 Expansion enclosure

The optional Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 expansion enclosure contain two expansion canisters, disk drives, and two power supplies. Two types of expansion enclosures are available: large form factor (LFF) Expansion Enclosure, a small form factor (SFF) Expansion Enclosure,

Figure 1-9 shows the rear of the expansion enclosure.




Figure 1-9 Expansion enclosure of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030

The expansion enclosure power supplies are the same as the control enclosure power supplies. A single power lead connector is on each power supply unit.

Each expansion canister provides two SAS interfaces that are used to connect to the control enclosure and any further optional expansion enclosures. The ports are numbered 1 on the left and 2 on the right. SAS port 1 is the IN port, and SAS port 2 is the OUT port.

The use of SAS connector 1 is mandatory because the expansion enclosure must be attached to a control enclosure or another expansion enclosure further up in the chain. SAS connector 2 is optional because it is used to attach to further expansion enclosures down the chain.

The Lenovo Storage V3700 V2 and Lenovo Storage V3700 V2 XP support a single chain of up to 10 expansion enclosures that attach to the control enclosure. The Lenovo Storage V5030 supports up to 40 expansion enclosures in a configuration that consists of two control enclosures, which are each attached to 20 expansion enclosures in two separate chains.

Table 1-7 shows the maximum number of supported expansion enclosures and the drive limits for each model.

Table 1-7 Expansion enclosure and drive limits

	V3700 V2	V3700 V2 XP	V5030
Maximum number of supported expansion enclosures	10	10	40
Maximum number of supported drives	392	392	1520

Each port includes two LEDs to show the status. The first LED indicates the link status and the second LED indicates the fault status.

For more information about LED and ports, see this web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/v3700_system_leds.html

Restriction: The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 expansion enclosures can be used with an Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 control enclosure only. The IBM Storwize V5000 for Lenovo expansion enclosures cannot be used with an Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 control enclosure.

1.5.6 Host interface cards

All Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support Ethernet ports as standard for iSCSI connectivity. For the Lenovo Storage V3700 V2 and Lenovo Storage V3700 V2 XP, these Ethernet ports are 1 GbE ports. For the Lenovo Storage V5030, these Ethernet ports are 10 GbE ports. The Lenovo Storage V3700 V2 XP also includes 12 Gb SAS ports for host connectivity as a standard.

Additional host connectivity options are available through an optional adapter card. Table 1-8 shows the available configurations for a single control enclosure.

Table 1-8 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 configurations available

	1 Gb Ethernet (iSCSI)	10 Gb Ethernet Copper 10GBASE-T (iSCSI)	12 Gb SAS	16 Gb FC	10 Gb Ethernet Optical SFP+ iSCSI/FCoE

Lenovo Storage V5030	8 ports (with optional adapter card).	4 ports (standard).	8 ports (with optional adapter card).	8 ports (with optional adapter card).	8 ports (with optional adapter card).
Lenovo Storage V3700 V2 XP	4 ports (standard). Additional 8 ports (with optional adapter card)	N/A	4 ports (standard). Additional 8 ports (with optional adapter card).	8 ports (with optional adapter card).	8 ports (with optional adapter card).
Lenovo Storage V3700 V2	4 ports (standard) Additional 8 ports (with optional adapter card)	N/A	8 ports (with optional adapter card)	8 ports (with optional adapter card)	8 ports (with optional adapter card)

Optional adapter cards: Only one pair of identical adapter cards is allowed for each control enclosure.

1.5.7 Disk drive types

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 enclosures support Flash Drives, SAS, and Nearline SAS drive types. Each drive has two ports (two PHYs) to provide fully redundant access from each node canister. I/O can be issued down both paths simultaneously.

Table 1-9 shows the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 disk drive types, disk revolutions per minute (RPMs), and sizes that are available at the time of writing.

Table 1-9 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 disk drive types

Drive type		RPM	Size
2.5-inch form factor	Flash Drive	N/A	400 GB, 800 GB, 1.6 TB, and 3.2 TB
2.5-inch form factor	Read Intensive (RI) Flash Drive	N/A	1.92 TB, 3.84 TB and 7.68 TB
2.5-inch form factor	SAS	10,000	900 GB, 1.2 TB, and 1.8 TB
2.5-inch form factor	SAS	15,000	300 GB, 600 GB and 900 GB
2.5-inch form factor	Nearline SAS	7,200	2 TB
3.5-inch form factor	SAS	10,000	900 GB, 1.2 TB, and 1.8 TB ^a
3.5-inch form factor	SAS	15,000	300 GB, 600 GB and 900 GB ^a
3.5-inch form factor	Nearline SAS	7,200	4 TB, 6 TB, 8 TB and 10 TB

a. 2.5-inch drive in a 3.5-inch drive carrier

1.6 Terms

In this section, we introduce the terms that are used for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 throughout this book.

1.6.1 Hosts

A *host* system is a server that is connected to Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 through a Fibre Channel connection, an iSCSI connection, or an SAS connection.

Hosts are defined on Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 by identifying their WWPNs for Fibre Channel and SAS hosts. The iSCSI hosts are identified by using their iSCSI names. The iSCSI names can be iSCSI qualified names (IQNs) or extended unique identifiers (EUIs). For more information, see Chapter 5, “Host configuration” on page 189.

Hosts can be Fibre Channel-attached through an existing Fibre Channel network infrastructure or direct-attached, iSCSI-attached through an existing IP network, or directly attached through SAS.

1.6.2 Node canister

A *node canister* provides host interfaces, management interfaces, and SAS interfaces to the control enclosure. A node canister has the cache memory, the internal storage to store software and logs, and the processing power to run the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 virtualization and management software. A clustered system consists of one or two node pairs. Each node pair forms one I/O group. I/O groups are explained in 1.6.3, “I/O groups” on page 17.

One of the nodes within the system, which is known as the *configuration node*, manages configuration activity for the clustered system. If this node fails, the system nominates the other node to become the configuration node.

1.6.3 I/O groups

Within Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, one or two pairs of node canisters are known as *I/O groups*. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 supports either two-node or four-node canisters in a clustered system, which provides either one or two I/O groups, depending on the model. See Table 1-6 on page 8 for more details.

When a host server performs I/O to one of its volumes, all of the I/Os for a specific volume are directed to the I/O group. Also, under normal conditions, the I/Os for that specific volume are always processed by the same node within the I/O group.

When a host server performs I/O to one of its volumes, all of the I/O for that volume is directed to the I/O group where the volume was defined. Under normal conditions, these I/Os are also always processed by the same node within that I/O group.

Both nodes of the I/O group act as preferred nodes for their own specific subset of the total number of volumes that the I/O group presents to the host servers (a maximum of 2,048 volumes for each host). However, both nodes also act as a failover node for the partner node within the I/O group. Therefore, a node takes over the I/O workload from its partner node (if required) without affecting the server’s application.

In a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 environment (which uses active-active architecture), the I/O handling for a volume can be managed by both nodes of the I/O group. The I/O groups must be connected to the SAN so that all hosts can access all nodes. The hosts must use multipath device drivers to handle this capability.

Up to 256 host server objects can be defined to one-I/O group or 512 host server objects can be defined in a two-I/O group system. More information about I/O groups is in Chapter 6, “Volume configuration” on page 269.

Important: The active/active architecture provides the availability to process I/Os for both controller nodes and allows the application to continue to run smoothly, even if the server has only one access route or path to the storage controller. This type of architecture eliminates the path/LUN thrashing that is typical of an active/passive architecture.

1.6.4 Clustered system

A *clustered system* consists of one or two pairs of node canisters. Each pair forms an I/O group. All configuration, monitoring, and service tasks are performed at the system level. The configuration settings are replicated across all node canisters in the clustered system. To facilitate these tasks, one or two management IP addresses are set for the clustered system. By using this configuration, you can manage the clustered system as a single entity.

A process exists to back up the system configuration data on to disk so that the clustered system can be restored in a disaster. This method does not back up application data. Only Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems configuration information is backed up.

System configuration backup: After the system configuration is backed up, save the backup data on to your local hard disk (or at the least outside of the SAN). If you are unable to access the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, you do not have access to the backup data if it is on the SAN. Perform this configuration backup after each configuration change to be safe.

The system can be configured by using the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 management software (GUI), CLI, or USB key.

1.6.5 RAID

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 contain several internal drive objects, but these drives cannot be directly added to the storage pools. Drives need to be included in a Redundant Array of Independent Disks (*RAID*) to provide protection against the failure of individual drives.

These drives are referred to as *members* of the array. Each array has a RAID level. RAID levels provide various degrees of redundancy and performance. The maximum number of members in the array varies based on the RAID level.

Traditional RAID (TRAIID) has the concept of hot spare drives. When an array member drive fails, the system automatically replaces the failed member with a hot spare drive and rebuilds the array to restore its redundancy. Candidate and spare drives can be manually exchanged with array members.

Apart from traditional disk arrays, spectrum virtualize software V7.6 introduced Distributed RAID. Distributed RAID improves recovery time of failed disk drives in an array by the distribution of spare capacity between primary disks, rather than dedicating a whole spare drive for replacement.

Details about traditional and distributed RAID arrays are described in details in Chapter 4, “Storage pools” on page 139.

1.6.6 Managed disks

A *managed disk* (MDisk) refers to the unit of storage that Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 virtualizes. This unit can be a logical volume on an external storage array that is presented to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 or a (traditional or distributed) RAID array that consists of internal drives. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can then allocate these MDisks into storage pools.

An MDisk is invisible to a host system on the storage area network because it is internal to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems.

An MDisk features the following modes:

- ▶ Array

Array mode MDisks are constructed from internal drives by using the RAID functionality. Array MDisks are always associated with storage pools.

- ▶ Unmanaged

LUNs that are presented by external storage systems to Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are discovered as unmanaged MDisks. The MDisk is not a member of any storage pools, which means that it is not used by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems.

- ▶ Managed

Managed MDisks are LUNs, which are presented by external storage systems to a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, that are assigned to a storage pool and provide extents so that volumes can use them. Any data that might be on these LUNs when they are added is lost.

- ▶ Image

Image MDisks are LUNs that are presented by external storage systems to a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. For more information, see Chapter 6, “Volume configuration” on page 269.

1.6.7 Quorum disks

A *quorum disk* is an MDisk that contains a reserved area for use exclusively by the system. In the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, internal drives can be considered as quorum candidates. The clustered system uses quorum disks to break a tie when exactly half the nodes in the system remain after a SAN failure.

The clustered system automatically forms the quorum disk by taking a small amount of space from an MDisk. It allocates space from up to three different MDisks for redundancy, although only one quorum disk is active.

To avoid the possibility of losing all of the quorum disks because of a failure of a single storage system if the environment has multiple storage systems, you need to allocate the quorum disk on different storage systems. You can manage the quorum disks by using the CLI.

IP quorum base support provides an alternative for IBM Storwize V5000 for Lenovo HyperSwap implementations. Instead of Fibre Channel storage on a third site, the IP network is used for communication between the IP quorum application and node canisters in the system to cope with tie-break situations if the inter-site link fails. The IP quorum application is a Java application that runs on a host at the third site. The IP quorum application enables the use of a lower-cost IP network-attached host as a quorum disk for simplified implementation and operation.

Note: IP Quorum allows the user to replace a third-site Fibre Channel attached quorum disk with an IP Quorum application. The Java application runs on a Linux host and is used to resolve split-brain situations. Quorum disks are still required in sites 1 and 2 for cookie crumb and meta data. The application can also be used with clusters in standard a standard topology configuration - but the primary use case is a customer with a cluster split over two sites (stretched or HyperSwap). You need Java to run the IP quorum. Your Network must provide as least < 80ms round-trip latency and all node need a service ip address and all service ip addresses must be ping able from the quorum host.

1.6.8 Storage pools

A *storage pool* (up to 1028 per system) is a collection of MDisk (up to 128) that are grouped to provide capacity for volumes. All MDisk in the pool are split into extents of the same size. Volumes are then allocated out of the storage pool and are mapped to a host system.

MDisk can be added to a storage pool at any time to increase the capacity of the pool. MDisk can belong in only one storage pool. For more information, see Chapter 4, “Storage pools” on page 139.

Each MDisk in the storage pool is divided into a number of extents. The size of the extent is selected by the administrator when the storage pool is created and cannot be changed later. The size of the extent ranges from 16 MB - 8 GB.

Default extent size: The GUI of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 has a default extent size value of 1024 MB when you define a new storage pool.

The extent size directly affects the maximum volume size and storage capacity of the clustered system.

A system can manage 2^{22} (4,194,304) extents. For example, with a 16 MB extent size, the system can manage up to $16 \text{ MB} \times 4,194,304 = 64 \text{ TB}$ of storage.

The effect of extent size on the maximum volume and cluster size is shown in Table 1-10.

Table 1-10 Maximum volume and cluster capacity by extent size

Extent size	Maximum volume capacity for normal volumes (GB)	Maximum storage capacity of cluster
16	2048 (2 TB)	64 TB
32	4096 (4 TB)	128 TB
64	8192 (8 TB)	256 TB
128	16384 (16 TB)	512 TB
256	32768 (32 TB)	1 PB
512	65536 (64 TB)	2 PB
1024	131072 (128 TB)	4 PB
2048	262144 (256 TB)	8 PB
4096	262144 (256 TB)	16 PB
8192	262144 (256 TB)	32 PB

Use the same extent size for all storage pools in a clustered system. This rule is a prerequisite if you want to migrate a volume between two storage pools. If the storage pool extent sizes are not the same, you must use volume mirroring to copy volumes between storage pools, as described in Chapter 4, “Storage pools” on page 139.

You can set a threshold warning for a storage pool that automatically issues a warning alert when the used capacity of the storage pool exceeds the set limit.

Child storage pools

Instead of being created directly from MDisk, *child pools* are created from existing capacity that is allocated to a parent pool. As with parent pools, volumes can be created that specifically use the capacity that is allocated to the child pool. Parent pools grow automatically as more MDisk are allocated to them. However, child pools provide a fixed capacity pool of storage. You can use a child pool to manage a quota of storage for a particular purpose.

Child pools can be created by using the management GUI, CLI, or IBM Spectrum Control when you create VMware vSphere virtual volumes. For more information about child pools, see Chapter 4, “Storage pools” on page 139.

Single-tiered storage pool

MDisk that are used in a single-tiered storage pool must have the following characteristics to prevent performance problems and other problems:

- ▶ They must have the same hardware characteristics, for example, the same RAID type, RAID array size, disk type, and disk revolutions per minute (RPMs).
- ▶ The disk subsystems that provide the MDisk must have similar characteristics, for example, maximum input/output operations per second (IOPS), response time, cache, and throughput.
- ▶ You need to use MDisk of the same size and ensure that the MDisk provide the same number of extents. If this configuration is not feasible, you must check the distribution of the volumes’ extents in that storage pool.

Multi-tiered storage pool

A *multi-tiered storage pool* has a mix of MDisk with more than one type of disk, for example, a storage pool that contains a mix of generic_hdd *and* generic_ssd MDisk.

A multi-tiered storage pool contains MDisk with different characteristics unlike the single-tiered storage pool. MDisk with similar characteristics then form the tiers within the pool. However, each tier needs to have MDisk of the same size and that provide the same number of extents.

A multi-tiered storage pool is used to enable automatic migration of extents between disk tiers by using the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Easy Tier function, as described in Chapter 9, “Advanced features for storage efficiency” on page 403.

This functionality can help improve the performance of host volumes on the IBM Storwize V5000 for Lenovo.

1.6.9 Volumes

A *volume* is a logical disk that is presented to a host system by the clustered system. In our virtualized environment, the host system has a volume that is mapped to it by Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 translates this volume into a number of extents, which are allocated across MDisk. The advantage with storage virtualization is that the host is decoupled from the underlying storage, so the virtualization appliance can move around the extents without affecting the host system.

The host system cannot directly access the underlying MDisk in the same manner as it can access RAID arrays in a traditional storage environment.

The following types of volumes are available:

- Striped

A striped volume is allocated one extent in turn from each MDisk in the storage pool. This process continues until the space that is required for the volume is satisfied.

It also is possible to supply a list of MDisk to use.

Figure 1-10 shows how a striped volume is allocated, assuming that 10 extents are required.

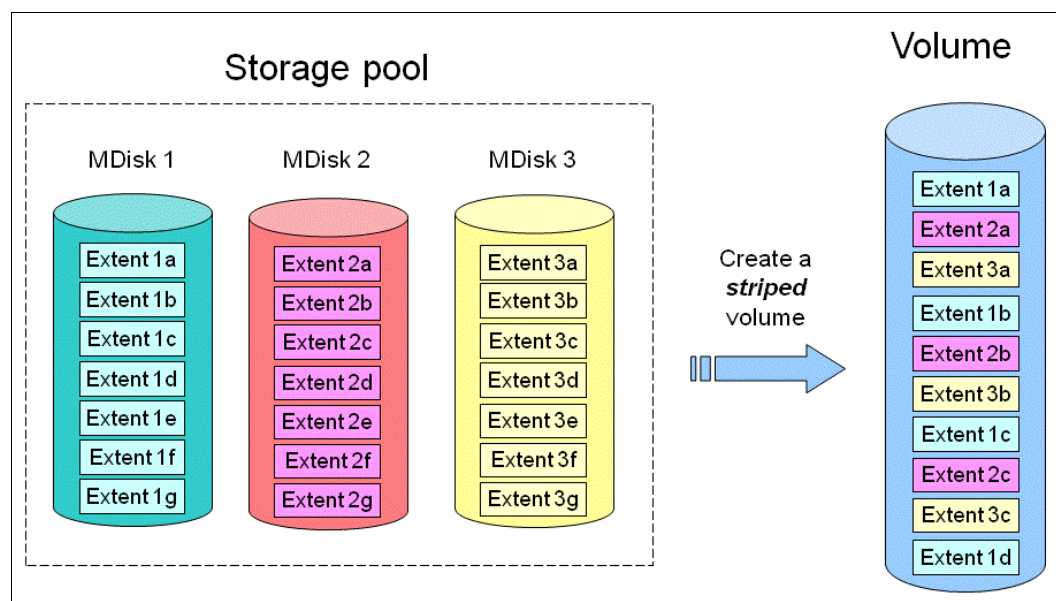


Figure 1-10 Striped volume

- Sequential

A sequential volume is a volume in which the extents are allocated one after the other from one MDisk to the next MDisk, as shown in Figure 1-11.




Figure 1-11 Sequential volume

► Image mode

Image mode volumes are special volumes that have a direct relationship with one MDisk. They are used to migrate existing data into and out of the clustered system to or from external FC SAN-attached storage.

When the image mode volume is created, a direct mapping is made between extents that are on the MDisk and the extents that are on the volume. The logical block address (LBA) x on the MDisk is the same as the LBA x on the volume, which ensures that the data on the MDisk is preserved as it is brought into the clustered system, as shown in Figure 1-12.




Figure 1-12 Image mode volume

Certain virtualization functions are not available for image mode volumes, so it is often useful to migrate the volume into a new storage pool. After it is migrated, the MDisk becomes a managed MDisk.

If you want to migrate data from an existing storage subsystem, use the storage migration wizard, which guides you through the process.

For more information, see Chapter 7, “Storage migration” on page 323.

If you add an MDisk that contains data to a storage pool, any data on the MDisk is lost. If you are presenting externally virtualized LUNs that contain data to an Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, import them as image mode volumes to ensure data integrity or use the migration wizard.

1.6.10 iSCSI

iSCSI is an alternative method of attaching hosts to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. The iSCSI function is a software function that is provided by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 code, not hardware.

In the simplest terms, iSCSI allows the transport of SCSI commands and data over an Internet Protocol network that is based on IP routers and Ethernet switches. iSCSI is a block-level protocol that encapsulates SCSI commands into TCP/IP packets and uses an existing IP network instead of requiring FC host bus adapters (HBAs) and a SAN fabric infrastructure.

Concepts of names and addresses are carefully separated in iSCSI.

An iSCSI name is a location-independent, permanent identifier for an iSCSI node. An iSCSI node has one iSCSI name, which stays constant for the life of the node. The terms *initiator name* and *target name* also refer to an iSCSI name.

An iSCSI address specifies the iSCSI name of an iSCSI node and a location of that node. The address consists of a host name or IP address, a TCP port number (for the target), and the iSCSI name of the node. An iSCSI node can have any number of addresses, which can change at any time, particularly if they are assigned by way of Dynamic Host Configuration Protocol (DHCP). An IBM Storwize V5000 for Lenovo node represents an iSCSI node and provides statically allocated IP addresses.

Each iSCSI node, that is, an initiator or target, has a unique IQN, which can have a size of up to 255 bytes. The IQN is formed according to the rules that were adopted for Internet nodes. The IQNs can be abbreviated by using a descriptive name, which is known as an *alias*. An alias can be assigned to an initiator or a target.

For more information about configuring iSCSI, see Chapter 4, “Storage pools” on page 139.

1.6.11 Serial-attached SCSI

The serial-attached SCSI (SAS) standard is an alternative method of attaching hosts to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support direct SAS host attachment to address easy-to-use, affordable storage needs. Each SAS port device has a worldwide unique 64-bit SAS address and operates at 12 Gbps.

1.6.12 Fibre Channel

Fibre Channel (FC) is the traditional method that is used for data center storage connectivity. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support FC connectivity at speeds of 4, 8, and 16 Gbps. Fibre Channel Protocol is used to encapsulate SCSI commands over the FC network. Each device in the network has a unique 64-bit worldwide port name (WWPN). The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support FC connections directly to a host server or to external FC switched fabrics.

1.7 Features

In this section, we describe the features of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. Different models offer a different range of features. See Table 1-3 on page 5 for a comparison.

1.7.1 Mirrored volumes

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 provide a function that is called *storage volume mirroring*, which enables a volume to have two physical copies. Each volume copy can belong to a different storage pool and be on a different physical storage system to provide a high-availability (HA) solution. Each mirrored copy can be either a generic, thin-provisioned, or compressed volume copy.

When a host system issues a write to a mirrored volume, Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 write the data to both copies. When a host system issues a read to a mirrored volume, Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 request it from the primary. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 automatically use the alternative copy without any outage for the host system. When the mirrored volume copy is repaired, Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 synchronize the data again.

A mirrored volume can be converted into a non-mirrored volume by deleting one copy or by splitting away one copy to create a non-mirrored volume.

The use of mirrored volumes can also assist with migrating volumes between storage pools that have different extent sizes. Mirrored volumes can also provide a mechanism to migrate fully allocated volumes to thin-provisioned or compressed volumes without any host outages.

The Volume Mirroring feature is included as part of the base software, and no license is required.

1.7.2 Thin Provisioning

Volumes can be configured to be *thin-provisioned* or *fully allocated*. A thin-provisioned volume behaves as though it were a fully allocated volume in terms of read/write I/O. However, when a volume is created, the user specifies two capacities: the real capacity of the volume and its virtual capacity.

The *real capacity* determines the quantity of MDisk extents that are allocated for the volume. The *virtual capacity* is the capacity of the volume that is reported to Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and to the host servers.

The real capacity is used to store the user data and the metadata for the thin-provisioned volume. The real capacity can be specified as an absolute value or a percentage of the virtual capacity.

The Thin Provisioning feature can be used on its own to create over-allocated volumes, or it can be used with FlashCopy. Thin-provisioned volumes can be used with the mirrored volume feature, also.

A thin-provisioned volume can be configured to *auto expand*, which causes the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 to automatically expand the real capacity of a thin-provisioned volume as it gets used. This feature prevents the volume from going offline. Auto expand attempts to maintain a fixed amount of unused real capacity on the volume. This amount is known as the *contingency capacity*. When the thin-provisioned volume is initially created, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 initially allocate only 2% of the virtual capacity in real physical storage. The contingency capacity and auto expand features seek to preserve this 2% of free space as the volume grows.

If the user modifies the real capacity, the contingency capacity is reset to be the difference between the used capacity and real capacity. In this way, the autoexpand feature does not cause the real capacity to grow much beyond the virtual capacity.

A volume that is created with a zero contingency capacity goes offline when it must expand. A volume with a non-zero contingency capacity stays online until it is used up.

To support the auto expansion of thin-provisioned volumes, the volumes themselves have a configurable warning capacity. When the used free capacity of the volume exceeds the warning capacity, a warning is logged. For example, if a warning of 80% is specified, the warning is logged when 20% of the free capacity remains. This approach is similar to the capacity warning that is available on storage pools.

A thin-provisioned volume can be converted to either a fully allocated volume or compressed volume by using volume mirroring (and vice versa).

The Thin Provisioning feature is included as part of the base software, and no license is required.

1.7.3 Real-time Compression

The Lenovo Storage V5030 model can create compressed volumes, allowing more data to be stored in the same physical space. IBM Real-time Compression (RtC) can be used for primary active volumes and with mirroring and replication (FlashCopy/Remote Copy). RtC is available on the Lenovo Storage V5030 model only.

Existing volumes can take advantage of Real-time Compression to result in an immediate capacity saving. An existing volume can be converted to a compressed volume by creating a compressed volume copy of the original volume followed by deleting the original volume.

No changes to the existing environment are required to take advantage of RtC. It is transparent to hosts while the compression occurs within the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems.

Software-only compression: The use of RtC on the Lenovo Storage V5030 requires dedicated CPU resources from the node canisters. If more performance is required for deploying RtC, consider purchasing the IBM Storwize V7000 for Lenovo system. The IBM Storwize V7000 for Lenovo system uses dedicated hardware options for compression acceleration.

The Lenovo Storage V5030 model has the additional memory upgrade (32 GB for each node canister). When the first compressed volume is created 4 of the 6 CPU cores are allocated to RtC. Of the 32GB of memory on each node canister roughly 9-10 GB is allocated to RtC. There are no hardware compression accelerators as in the IBM Storwize V7000 for Lenovo. The actual LZ4 compression is done by the CPUs as was the case with the IBM Storwize V7000 for Lenovo. Table 1-11 on page 28 shows how the cores are used with RtC.

Table 1-11 Cores usage with RtC

Model	Compression Disabled		Compression Enabled	
	Normal Processing	RtC	Normal Processing	RtC
Lenovo Storage V3700 V2	2 cores	NA	NA	NA
Lenovo Storage V3700 V2 XP	2 cores	NA	NA	NA
Lenovo Storage V5030	6 cores	0 cores	2 cores	4 cores + HT

The faster CPU with more cores, the extra memory and the hyper-threading capability of the Lenovo Storage V5030 as well as improvements to RtC software results in good performance for smaller customer configurations common to the market segment this product is intended to serve. The feature is licensed per enclosure. Conversely, Real-time Compression is not available on the Lenovo V3700 V2 or V3700 V2 XP model.

1.7.4 Easy Tier

IBM Easy Tier provides a mechanism to seamlessly migrate extents to the most appropriate tier within the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 solution. This migration can be to different tiers of internal drives within Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 or to external storage systems that are virtualized by Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, for example, an IBM FlashSystem 900.

The Easy Tier function can be turned on or turned off at the storage pool and volume level.

You can demonstrate the potential benefit of Easy Tier in your environment before you install Flash drives by using the IBM Storage Advisor Tool. For more information about Easy Tier, see Chapter 9, “Advanced features for storage efficiency” on page 403.

The IBM Easy Tier feature is licensed per enclosure.

1.7.5 Storage Migration

By using the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Storage Migration feature, you can easily move data from other existing Fibre Channel-attached external storage to the internal capacity of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. You can migrate data from other storage to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems to realize the benefits of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 with features, such as the easy-to-use GUI, internal virtualization, thin provisioning, and copy services.

The Storage Migration feature is included in the base software, and no license is required.

1.7.6 FlashCopy

The FlashCopy feature copies a source volume on to a target volume. The original contents of the target volume is lost. After the copy operation starts, the target volume has the contents of the source volume as it existed at a single point in time. Although the copy operation completes in the background, the resulting data at the target appears as though the copy was made instantaneously. FlashCopy is sometimes described as an instance of a *time-zero* (T0) copy or *point-in-time* (PiT) copy technology.

FlashCopy can be performed on multiple source and target volumes. FlashCopy permits the management operations to be coordinated so that a common single point in time is chosen for copying target volumes from their respective source volumes.

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 also permit multiple target volumes to be FlashCopies from the same source volume. This capability can be used to create images from separate points in time for the source volume, and to create multiple images from a source volume at a common point in time. Source and target volumes can be any volume type (generic, thin-provisioned, or compressed).

Reverse FlashCopy enables target volumes to become restore points for the source volume without breaking the FlashCopy relationship and without waiting for the original copy operation to complete. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support multiple targets and multiple rollback points.

The FlashCopy feature is licensed per enclosure.

For more information about FlashCopy copy services, see Chapter 10, “Copy services” on page 451.

1.7.7 Remote Copy

Remote Copy can be implemented in one of two modes, synchronous or asynchronous.

With the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, Metro Mirror and Global Mirror are the IBM branded terms for the functions that are synchronous Remote Copy and asynchronous Remote Copy.

By using the Metro Mirror and Global Mirror copy services features, you can set up a relationship between two volumes so that updates that are made by an application to one volume are mirrored on the other volume. The volumes can be in the same system or on two different systems.

For both Metro Mirror and Global Mirror copy types, one volume is designated as the primary and the other volume is designated as the secondary. Host applications write data to the primary volume, and updates to the primary volume are copied to the secondary volume. Normally, host applications do not perform I/O operations to the secondary volume.

The Metro Mirror feature provides a synchronous copy process. When a host writes to the primary volume, it does not receive confirmation of I/O completion until the write operation completes for the copy on the primary and secondary volumes. This design ensures that the secondary volume is always up-to-date with the primary volume if a failover operation must be performed.

The Global Mirror feature provides an asynchronous copy process. When a host writes to the primary volume, confirmation of I/O completion is received before the write operation completes for the copy on the secondary volume. If a failover operation is performed, the application must recover and apply any updates that were not committed to the secondary volume. If I/O operations on the primary volume are paused for a brief time, the secondary volume can become an exact match of the primary volume.

Global Mirror can operate with or without cycling. When it is operating without cycling, write operations are applied to the secondary volume as soon as possible after they are applied to the primary volume. The secondary volume is less than 1 second behind the primary volume, which minimizes the amount of data that must be recovered in a failover. However, this approach requires that a high-bandwidth link is provisioned between the two sites.

When Global Mirror operates with cycling mode, changes are tracked and where needed copied to intermediate change volumes. Changes are transmitted to the secondary site periodically. The secondary volumes are much further behind the primary volume, and more data must be recovered in a failover. Because the data transfer can be smoothed over a longer time period, lower bandwidth is required to provide an effective solution.

For more information about the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 copy services, see Chapter 10, “Copy services” on page 451.

The Lenovo Remote Copy feature is licensed for each enclosure.

1.7.8 IP replication

IP replication enables the use of lower-cost Ethernet connections for remote mirroring. The capability is available as a chargeable option on all IBM Storwize for Lenovo and Lenovo Storage V series family systems.

The function is transparent to servers and applications in the same way that traditional Fibre Channel-based mirroring is transparent. All remote mirroring modes (Metro Mirror, Global Mirror, and Global Mirror with Change Volumes) are supported.

Configuration of the system is straightforward. The IBM Storwize for Lenovo and Lenovo Storage V series systems normally find each other in the network, and they can be selected from the GUI.

IP replication includes Bridgeworks SANSlide network optimization technology, and it is available at no additional charge. Remember, Remote Mirror is a chargeable option but the price does not change with IP replication. Existing Remote Mirror users have access to the function at no additional charge.

IP connections that are used for replication can have long *latency* (the time to transmit a signal from one end to the other), which can be caused by distance or by many “hops” between switches and other appliances in the network. Traditional replication solutions transmit data, wait for a response, and then transmit more data, which can result in network utilization as low as 20% (based on IBM measurements). And this scenario gets worse the longer the latency.

Bridgeworks SANSlide technology that is integrated with the IBM Storwize for Lenovo and Lenovo Storage V series families require no separate appliances, no additional cost, and no configuration steps. It uses artificial intelligence (AI) technology to transmit multiple data streams in parallel, adjusting automatically to changing network environments and workloads. SANSlide improves network bandwidth utilization up to 3x so clients can deploy a less costly

network infrastructure or take advantage of faster data transfer to speed up replication cycles, improve remote data currency, and enjoy faster recovery.

IP replication can be configured to use any of the available 1 GbE or 10 GbE Ethernet ports (apart from the technician port) on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. See Table 1-8 on page 15 for port configuration options.

Copy services configuration limits

For the most up-to-date list of these limits, see the following support page:

<https://support.lenovo.com/us/en/solutions/ht505190>

1.7.9 External virtualization

By using this feature, you can consolidate FC SAN-attached disk controllers from various vendors into pools of storage. In this way, the storage administrator can manage and provision storage to applications from a single user interface and use a common set of advanced functions across all of the storage systems under the control of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. External virtualization is only available for the Lenovo Storage V5030.

The External Virtualization feature is licensed per disk enclosure.

1.7.10 Encryption

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 provide optional encryption of data-at-rest functionality, which protects against the potential exposure of sensitive user data and user metadata that is stored on discarded, lost, or stolen storage devices. Encryption can be enabled and configured only on the LenovoV3700 V2 XP and Lenovo Storage V5030 enclosures that support encryption. LenovoV3700 V2 does not offer encryption functionality.

Encryption is a licensed feature that requires a license key to enable it before it can be used.

1.8 Problem management and support

In this section, we introduce problem management and support topics.

1.8.1 Support assistance

To use Support assistance, you must have access to the internet. Support assistance enables support personnel to access the system to complete troubleshooting and maintenance tasks. You can configure either local support assistance, where support personnel visit your site to fix problems with the system, or remote support assistance. Both local and remote support assistance uses secure connections to protect data exchange between the support center and system. More access controls can be added by the system administrator. You can use the management GUI or the command-line interface to view support assistance settings.

1.8.2 Event notifications

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can use Simple Network Management Protocol (SNMP) traps, syslog messages, and e-mail to notify you and the IBM Support

Center when significant events are detected. Any combination of these notification methods can be used simultaneously.

You can configure Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 to send different types of notification to specific recipients and choose the alerts that are important to you. When you configure Call Home to the Lenovo Support Center, all events are sent through email only.

1.8.3 SNMP traps

SNMP is a standard protocol for managing networks and exchanging messages. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 send. You can use the management GUI or the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 CLI to configure and modify your SNMP settings.

You can use the Management Information Base (MIB) file for SNMP to configure a network management program to receive SNMP messages that are sent by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. This file can be used with SNMP messages from all versions of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 software.

1.8.4 Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IP network can be IPv4 or IPv6. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can send syslog messages that notify personnel about an event. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can transmit syslog messages in expanded or concise format. You can use a syslog manager to view the syslog messages that Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 send. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the management GUI or the CLI to configure and modify your syslog settings.

1.8.5 Call Home email

The Call Home feature transmits operational and error-related data to you and Lenovo through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification email. When configured, this function alerts Lenovo service personnel about hardware failures and potentially serious configuration or environmental issues. You can use the Call Home function if you have a maintenance contract with Lenovo or if the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are within the warranty period.

To send email, you must configure at least one SMTP server. You can specify as many as five other SMTP servers for backup purposes. The SMTP server must accept the relaying of email from the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 clustered system IP address. You can then use the management GUI or the CLI to configure the email settings, including contact information and email recipients. Set the reply address to a valid email address. Send a test email to check that all connections and infrastructure are set up correctly. You can disable the Call Home function at any time by using the management GUI or the CLI.

1.9 More information resources

For more information about Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, see the following web pages:

- The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 home page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/lenovo_vseries.html

- Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Information Center:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.common.nav.doc/overview_storage_vseries.html

The Online Information Center also includes a Learning and Tutorial section where you can obtain videos that describe the use and implementation of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

Initial configuration

This chapter describes the initial configuration steps for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

Specifically, this chapter provides information about the following topics:

- ▶ 2.1, “Hardware installation planning” on page 36
- ▶ 2.2, “SAN configuration planning” on page 40
- ▶ 2.3, “FC direct-attach planning” on page 43
- ▶ 2.4, “SAS direct-attach planning” on page 45
- ▶ 2.5, “LAN configuration planning” on page 46
- ▶ 2.6, “Host configuration planning” on page 48
- ▶ 2.7, “Miscellaneous configuration planning” on page 49
- ▶ 2.8, “System management” on page 50
- ▶ 2.9, “First-time setup” on page 52
- ▶ 2.10, “Initial configuration” on page 56

2.1 Hardware installation planning

After you verify that you have all of the hardware components that you purchased, it is important to perform the correct planning before the actual physical installation. The following checklist of requirements can be used to plan your installation:

- ❑ Install the hardware as described in Chapter 2 of the *Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Quick Installation Guide*. The document is available at these sites:
<http://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v3700/v2/6535/documentation>
<http://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v5030/6536/documentation>
- ❑ An appropriate 19-inch rack must be available. Depending on the number of enclosures to install, more than one might be required. Each enclosure measures 2 U. A single Lenovo Storage V3700 V2 or Lenovo Storage V3700 V2 XP control enclosure supports up to 10 expansion enclosures. A single Lenovo Storage V5030 control enclosure supports up to 20 expansion enclosures.
- ❑ Redundant power outlets must be in the rack for each of the two power cords that are required for each enclosure to be installed. Several power outlets are required, depending on the number of enclosures to be installed. The power cords conform to the IEC320 C13/C14 standards.
- ❑ A minimum of four Fibre Channel ports that are attached to redundant fabrics are required. For dual I/O group systems, a minimum of eight Fibre Channel ports are required.

Fibre Channel ports: Fibre Channel (FC) ports are required only if you are using FC hosts or clustered systems that are arranged as two I/O groups. You can use the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 with Ethernet-only cabling for Internet Small Computer System Interface (iSCSI) hosts or use serial-attached SCSI (SAS) cabling for hosts that are directly attached.

- ❑ For the Lenovo Storage V3700 V2 XP system, up to two hosts can be directly connected by using SAS ports 2 and 3 on each node canister, with SFF-8644 mini SAS HD cabling.
- ❑ You must have a minimum of two Ethernet ports on the LAN, with four preferred for more redundancy or iSCSI host access.
- ❑ You must have a minimum of two Ethernet cable drops, with four preferred for more redundancy or iSCSI host access. If you have two I/O groups, you must have a minimum of four Ethernet cable drops. Ethernet port 1 on each node canister must be connected to the LAN, with port two as optional.

LAN connectivity: Port 1 on each node canister must be connected to the same physical local area network (LAN) or be configured in the same virtual LAN (VLAN) and be on the same subnet or set of subnets.

Technician port: On the Lenovo Storage V3700 V2 and V3700 V2 XP models, Port 2 is the *technician port*, which is used for system initialization and service. Port 2 must not be connected to the LAN until the system initialization or service is complete.

The Lenovo Storage V5030 model has a dedicated technician port.

- ❑ The 10 Gb Ethernet (copper) ports of a Lenovo Storage V5030 system require a Category 6A shielded cable that is terminated with an 8P8C modular connector (RJ45 compatible connector) to function at 10 Gb.
- ❑ Verify that the default IP addresses that are configured on Ethernet port 1 on each of the node canisters (192.168.70.121 on node 1 and 192.168.70.122 on node 2) do not conflict with existing IP addresses on the LAN. The default mask that is used with these IP addresses is 255.255.255.0, and the default gateway address that is used is 192.168.70.1.
- ❑ You need a minimum of three IPv4 or IPv6 IP addresses for systems that are arranged as one I/O group and minimum of five if you have two I/O groups. One is for the clustered system and is used by the administrator for management, and one for each node canister for service access as needed.

IP addresses: An additional IP address must be used for backup configuration access. This other IP address allows a second system IP address to be configured on port 2 of either node canister, which the storage administrator can also use for management of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 system.

- ❑ A minimum of one and up to eight IPv4 or IPv6 addresses are needed if iSCSI-attached hosts access volumes from the Lenovo Storage V Series.
- ❑ At least two 0.6-meter (1.96 feet), 1.5-meter (4.9 feet), or 3-meter (9.8 feet) 12 Gb mini-SAS cables are required for each expansion enclosure. The length of the cables depends on the physical rack location of the expansion enclosure relative to the control enclosures or other expansion enclosures.

2.1.1 Procedure to install the SAS cables

We show the procedures to install the SAS cables for the different models.

Lenovo Storage V3700 V2 and V3700 V2 XP

The Lenovo Storage V3700 V2 and Lenovo Storage V3700 V2 XP support up to 10 expansion enclosures in a single chain. To install the cables, complete the following steps:

1. By using the supplied SAS cables, connect the control enclosure to the first expansion enclosure:
 - a. Connect SAS port 1 of the left node canister in the control enclosure to SAS port 1 of the left expansion canister in the first expansion enclosure.
 - b. Connect SAS port 1 of the right node canister in the control enclosure to SAS port 1 of the right expansion canister in the first expansion enclosure.
2. To connect a second expansion enclosure, use the supplied SAS cables to connect it to the previous enclosure in the chain:
 - a. Connect SAS port 2 of the left canister in the previous expansion enclosure to SAS port 1 of the left expansion canister in the next expansion enclosure.
 - b. Connect SAS port 2 of the right canister in the previous expansion enclosure to SAS port 1 of the right expansion canister in the next expansion enclosure.
3. Repeat the previous steps until all expansion enclosures are connected.

Figure 2-1 shows how to cable a Lenovo Storage V3700 V2.




Figure 2-1 Lenovo Storage V3700 V2 SAS cabling

Figure 2-2 shows how to cable a Lenovo Storage V3700 V2 XP.




Figure 2-2 Lenovo Storage V3700 V2 XP SAS cabling

Lenovo Storage V5030

The Lenovo Storage V5030 supports up to 20 expansion enclosures per I/O group in two SAS chains of 10. Up to 40 expansion enclosures can be supported in a two I/O group configuration. To install the cables, complete the following steps:

1. By using the supplied SAS cables, connect the control enclosure to first expansion enclosure by using the first chain:
 - a. Connect SAS port 1 of the left node canister in the control enclosure to SAS port 1 of the left expansion canister in the first expansion enclosure.
 - b. Connect SAS port 1 of the right node canister in the control enclosure to SAS port 1 of the right expansion canister in the first expansion enclosure.
2. To connect a second expansion enclosure, use the supplied SAS cables to connect the control enclosure to expansion enclosure by using the second chain:
 - a. Connect SAS port 2 of the left node canister in the control enclosure to SAS port 1 of the left expansion canister in the second expansion enclosure.
 - b. Connect SAS port 2 of the right node canister in the control enclosure to SAS port 1 of the right expansion canister in the second expansion enclosure.
3. To connect additional expansion enclosures, alternate connecting them between chain one and chain two to keep the configuration balanced:
 - a. Connect SAS port 2 of the left canister in the previous expansion enclosure to SAS port 1 of the left expansion canister in the next expansion enclosure.
 - b. Connect SAS port 2 of the right canister in the previous expansion enclosure to SAS port 1 of the right expansion canister in the next expansion enclosure.
4. Repeat the steps until all expansion enclosures are connected.

Figure 2-3 shows how to cable a Lenovo Storage V5030.




Figure 2-3 Lenovo Storage V5030 cabling

2.2 SAN configuration planning

To connect a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 system to a Fibre Channel (FC) SAN, you must first install the optional 16 Gb FC adapters in every node canister that you connect. Ensure that you use the correct fibre cables to connect the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 to the Fibre Channel SAN. With the FC cards installed, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can be used to interface with FC hosts, external storage controllers, and other IBM Storwize for Lenovo and Lenovo Storage V series systems that are visible on the SAN fabric.

The Lenovo Storage V3700 V2 and V3700 V2 XP models support a single I/O group only and can migrate from external storage controllers only. The Lenovo Storage V5030 supports up to two I/O groups that form a cluster over the FC fabric. The Lenovo Storage V5030 also supports full virtualization of external storage controllers.

The advised SAN configuration consists of a minimum of two fabrics that encompass all host ports and any ports on external storage systems that are to be virtualized by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 ports must have the same number of cables that are connected, and they must be evenly split between the two fabrics to provide redundancy if one of the fabrics goes offline (planned or unplanned).

Zoning must be implemented after the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, hosts, and optional external storage systems are connected to the SAN fabrics. To enable the node canisters to communicate with each other in band, create a zone with only the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 WWPNs (two from each node canister) on each of the two fabrics.

If an external storage system is to be virtualized, create a zone in each fabric with the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 worldwide port names (WWPNs) (two from each node canister) with up to a maximum of eight WWPNs from the external storage system. Assume that every host has a Fibre Channel connection to each fabric. Create a zone with the host WWPN and one WWPN from each node canister in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems in each fabric.

Important: It is critical that only one initiator host bus adapter (HBA) is in any zone.

For load balancing between the node ports on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, alternate the host Fibre Channel ports between the ports of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

A maximum of eight paths through the SAN are allowed from each host to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. Hosts where this number is exceeded are not supported. The restriction limits the number of paths that the multipathing driver must resolve. A host with only two HBAs must not exceed this limit with the correct zoning in a dual fabric SAN.

Maximum ports or WWPNs: The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support a maximum of 16 ports or WWPNs from a virtualized external storage system.

Figure 2-4 shows how to cable devices to the SAN. Optionally, ports 3 and 4 can be connected to SAN fabrics to provide additional redundancy and throughput. Refer to this example as the zoning is described.

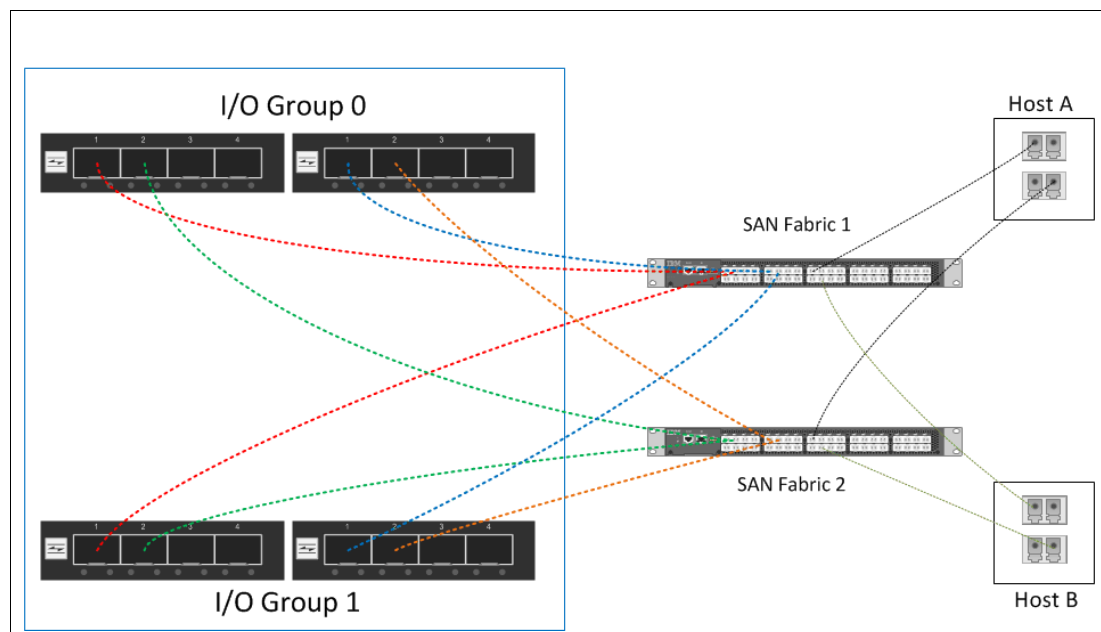


Figure 2-4 Lenovo Storage V3700 V2, V3700 V2 XP and V5030 SAN cabling and zoning diagram

Create a host/Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 zone for each server that volumes are mapped to and from the clustered system, as shown in the following examples in Figure 2-4 on page 41:

- ▶ Zone Host A port 1 (HBA 1) with all node canister ports 1
- ▶ Zone Host A port 2 (HBA 2) with all node canister ports 2
- ▶ Zone Host B port 1 (HBA 1) with all node canister ports 3
- ▶ Zone Host B port 2 (HBA 2) with all node canister ports 4

Similar zones must be created for all other hosts with volumes on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 I/O groups.

Verify the interoperability with which the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 connects to SAN switches or directors by following the requirements that are provided at these web pages:

- ▶ Lenovo Storage V3700 V2 and V3700 V2 XP
<http://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v3700-v2/6535/documentation>
- ▶ Lenovo Storage V5030
<http://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v5030/6536/documentation>

Switches or directors are at the firmware levels that are supported by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

Important: The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 port login maximum that is listed in the restriction document must not be exceeded. The document is available at this web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/svc_webgetstartovr_21pax3.html

Connectivity issues: If any connectivity issues occur between the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 ports and the Brocade SAN switches or directors at 8 Gbps, see this web page for the correct setting of the **fillword port config** parameter in the Brocade operating system:

<https://ibm.biz/Bdrb4g>

2.3 FC direct-attach planning

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can be used with a direct-attach Fibre Channel host configuration. The advised configuration for direct attachment is at least one Fibre Channel cable from the host that is connected to each node of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 to provide redundancy if one of the nodes goes offline, as shown in Figure 2-5.




Figure 2-5 Lenovo Storage V3700 V2, V3700 V2 XP and V5030 FC direct-attach host configuration

If your direct-attach Fibre Channel host requires connectivity to both of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 I/O groups (the Lenovo Storage V5030), we suggest that at least one Fibre Channel cable is used from the host to each of the node canisters of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, as shown in Figure 2-6. This suggestion also applies to a cluster where one I/O group is an IBM Storwize V5000 for Lenovo and the other I/O group is a Lenovo Storage V5030.

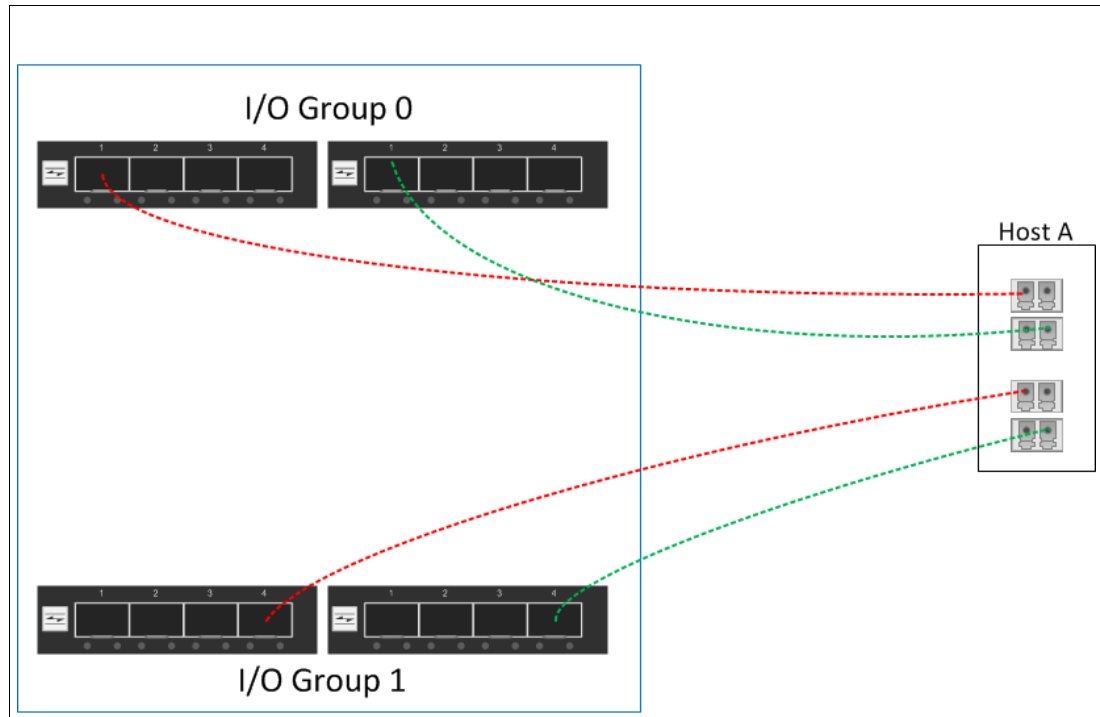


Figure 2-6 Lenovo Storage V5030 FC direct-attach host configuration to I/O groups

Verify direct-attach interoperability with the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and the supported server operating systems by following the requirements that are provided at these sites:

<http://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v3700v2/6535/documentation>

<http://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v5030/6536/documentation>

2.4 SAS direct-attach planning

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 allow SAS host attachment by using an optional SAS card that must be installed in both node canisters. In addition, the Lenovo V3700 V2 XP has two onboard SAS ports for host attachment. The SAS expansion ports cannot be used for host attachment. Figure 2-7, Figure 2-8, and Figure 2-9 show the SAS ports that can be used for host attachment in yellow for each Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models.




Figure 2-7 Lenovo Storage V3700 V2 SAS host attachment




Figure 2-8 Lenovo Storage V3700 V2 XP SAS host attachment




Figure 2-9 Lenovo Storage V5030 SAS host attachment

Inserting cables: You can insert the cables upside down despite the keyway. Ensure that the blue tag on the SAS connector is underneath when you insert the cables.

We suggest that each SAS host is connected to both node canisters because this approach provides redundancy in a path or canister failure.

2.5 LAN configuration planning

Two Ethernet ports per node canister are available for connection to the LAN on an Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems.

Use Ethernet port 1 to access the management graphical user interface (GUI), the service assistant GUI for the node canister, and iSCSI host attachment. Port 2 can be used for the management GUI and iSCSI host attachment.

Each node canister in a control enclosure connects over an Ethernet cable from Ethernet port 1 of the canister to an enabled port on your Ethernet switch or router. Optionally, you can attach an Ethernet cable from Ethernet port 2 on the canister to your Ethernet network.

Configuring IP addresses: No issue exists with the configuration of multiple IPv4 or IPv6 addresses on an Ethernet port or with the use of the same Ethernet port for management and iSCSI access.

However, you cannot use the same IP address for management and iSCSI host use.

Table 2-1 shows possible IP configuration options of the Ethernet ports on the Lenovo Storage V3700 V2,V3700 V2 XP, and V5030 systems.

Table 2-1 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 IP address configuration options per node canister

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 management node canister 1		Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 partner node canister 2	
IPv4/6 management address	Ethernet port 1	IPv4/6 service address	Ethernet port 1
IPv4/6 service address		IPv4/6 iSCSI address	
IPv4/6 iSCSI address			
IPv4/6 management address	Ethernet port 2	IPv4/6 iSCSI address	Ethernet port 2
IPv4/6 iSCSI address			

IP management addresses: The IP management address that is shown on node canister 1 in Table 2-1 is an address on the configuration node. If a failover occurs, this address transfers to node canister 2, and this node canister becomes the new configuration node. The management addresses are managed by the configuration node canister only (1 or 2, and in this case, node canister 1).

Technician port: On the Lenovo Storage V3700 V2 and V3700 V2 XP models, port 2 serves as the technician port, which is used for system initialization and service. Port 2 must not be connected to the LAN until the system initialization or service is complete.

The Lenovo Storage V5030 model has a dedicated technician port.

2.5.1 Management IP address considerations

Because Ethernet port 1 from each node canister must connect to the LAN, a single management IP address for the clustered system is configured as part of the initial setup of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems.

The management IP address is associated with one of the node canisters in the clustered system and that node then becomes the configuration node. If this node goes offline (planned or unplanned), the management IP address fails over to the other node's Ethernet port 1.

For more clustered system management redundancy, you need to connect Ethernet port 2 on each of the node canisters to the LAN, which allows for a backup management IP address to be configured for access, if necessary.

Figure 2-10 shows a logical view of the Ethernet ports that are available for the configuration of the one or two management IP addresses. These IP addresses are for the clustered system and associated with only one node, which is then considered the configuration node.

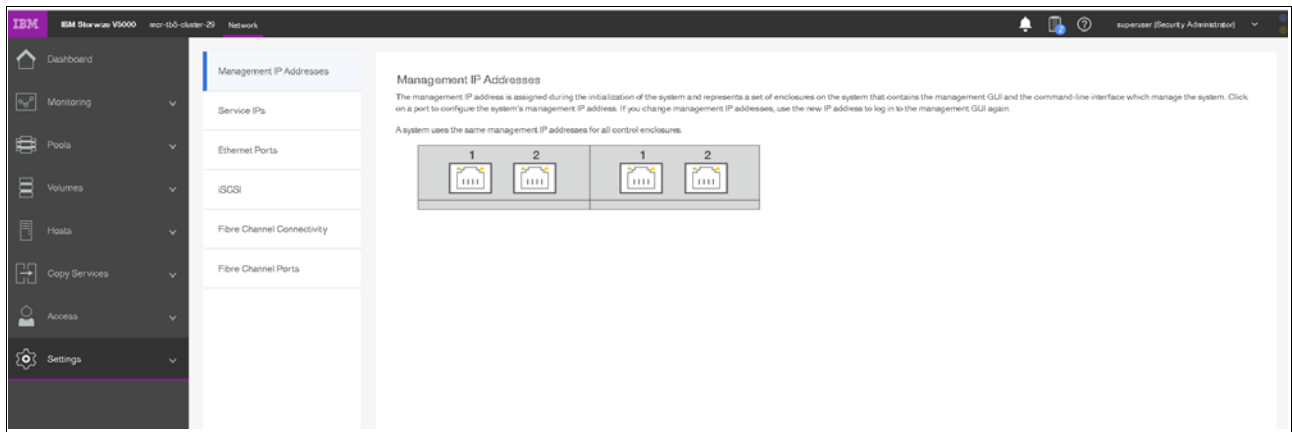


Figure 2-10 Ethernet ports that are available for configuration

2.5.2 Service IP address considerations

Ethernet port 1 on each node canister is used for system management and for service access, when required. In normal operation, the service IP addresses are not needed. However, if a node canister problem occurs, it might be necessary for service personnel to log on to the node to perform service actions.

Figure 2-11 shows a logical view of the Ethernet ports that are available for the configuration of the service IP addresses. Only port one on each node can be configured with a service IP address.

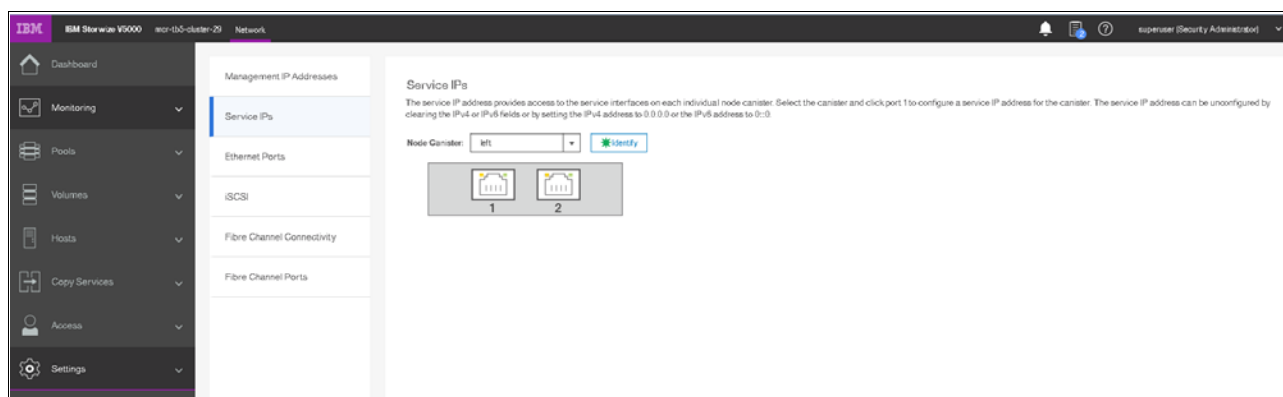


Figure 2-11 Service IP addresses that are available for configuration

2.6 Host configuration planning

Hosts must have two Fibre Channel connections for redundancy, but the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 also support hosts with a single HBA port connection. However, if that HBA loses its link to the SAN fabric or the fabric fails, the host loses access to its volumes. Even with a single connection to the SAN, the host has multiple paths to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 volumes because that single connection must be zoned with at least one Fibre Channel port per node. Therefore, multipath software is required. Multipath software is required also for direct-attach SAS hosts. They can connect by using a single host port, but for redundancy, two SAS connections per host are advised.

If two connections per host are used, multipath software is also required on the host. If an iSCSI host is deployed, it also requires multipath software. All node canisters must be configured and connected to the network so that any iSCSI hosts see at least two paths to volumes. Multipath software is required to manage these paths.

Various operating systems are supported by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. For more information about various configurations supported, check the Lenovo interoperability matrix web page at the following address:

<https://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v3700v2/6535/documentation>

<https://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v5030/6536/documentation>

For more information, see Chapter 5, “Host configuration” on page 189.

2.7 Miscellaneous configuration planning

During the initial setup of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, the installation wizard asks for various information that needs to be available during the installation process. Several of these fields are mandatory to complete the initial configuration.

Collect the information in the following checklist *before* the initial setup is performed. The date and time can be manually entered, but to keep the clock synchronized, use a Network Time Protocol (NTP) service:

- ☐ Document the LAN NTP server IP address that is used for the synchronization of devices.
- ☐ To send alerts to storage administrators and to set up Call Home to Lenovo for service and support, you need the following information:
 - ☐ Name of the primary storage administrator for Lenovo to contact, if necessary.
 - ☐ Email address of the storage administrator for Lenovo to contact, if necessary.
 - ☐ Phone number of the storage administrator for Lenovo to contact, if necessary.
 - ☐ Physical location of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems for Lenovo service (for example, Building 22, first floor).
 - ☐ Simple Mail Transfer Protocol (SMTP) or email server address to direct alerts to and from the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.
 - ☐ For the Call Home service to work, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems must have access to an SMTP server on the LAN that can forward emails to the default Lenovo service address.
 - ☐ Email address of local administrators that must be notified of alerts.
 - ☐ IP address of Simple Network Management Protocol (SNMP) server to direct alerts to, if required (for example, operations or Help desk).

After the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 initial configuration, you might want to add more users who can manage the system. You can create as many users as you need, but the following roles generally are configured for users:

- ▶ Security Admin
- ▶ Administrator
- ▶ CopyOperator
- ▶ Service
- ▶ Monitor

The user in the Security Admin role can perform any function on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

The user in the Administrator role can perform any function on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, except manage users.

User creation: The Security Admin role is the only role that has the create users function. Limit this role to as few users as possible.

The user in the CopyOperator role can view anything in the system, but the user can configure and manage only the copy functions of the FlashCopy capabilities.

The user in the Monitor role can view object and system configuration information but cannot configure, manage, or modify any system resource.

The only other role that is available is the service role, which is used if you create a user ID for the Lenovo service support representative (SSR). With this user role, Lenovo service personnel can view anything on the system (as with the Monitor role) and perform service-related commands, such as adding a node back to the system after it is serviced or including disks that were excluded.

2.8 System management

The graphical user interface (GUI) is used to configure, manage, and troubleshoot the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems. It is used primarily to configure Redundant Array of Independent Disks (RAID) arrays and logical drives, assign logical drives to hosts, replace and rebuild failed disk drives, and expand the logical drives.

It allows for troubleshooting and management tasks, such as checking the status of the storage server components, updating the firmware, and managing the storage server.

The GUI also offers advanced functions, such as FlashCopy, Volume Mirroring, Remote Mirroring, and Easy Tier. A command-line interface (CLI) for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems also are available.

This section describes system management by using the GUI and CLI.

2.8.1 Graphical user interface (GUI)

A web browser is used for GUI access. You must use a supported web browser to access the management GUI. At the time of writing, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support the following browsers:

- ▶ Mozilla Firefox 41
- ▶ Mozilla Firefox Extended Support Release (ESR) 38
- ▶ Microsoft Internet Explorer (IE) 11 and Microsoft Edge
- ▶ Google Chrome 46

Supported web browsers: Follow this link to find more information about supported browsers and to check the latest supported levels:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/svc_webbrowserreqmts_3rdhu7.html

Complete the following steps to open the management GUI from any web browser:

1. Browse to one of the following locations:
 - `http(s)://host name of your cluster/`
 - `http(s)://cluster IP address of your cluster/`
(An example is `https://192.168.70.120.`)
2. Use the password that you created during system setup to authenticate with the superuser or any additional accounts that you created. The default user name and password for the management GUI is shown:
 - User name: superuser
 - Password: passw0rd

Note: The 0 character in the password is the number zero, not the letter O.

For more information, see Chapter 3, “Graphical user interface overview” on page 75.

After you complete the initial configuration that is described in 2.10, “Initial configuration” on page 56, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Systems overview window opens, as shown in Figure 2-12.

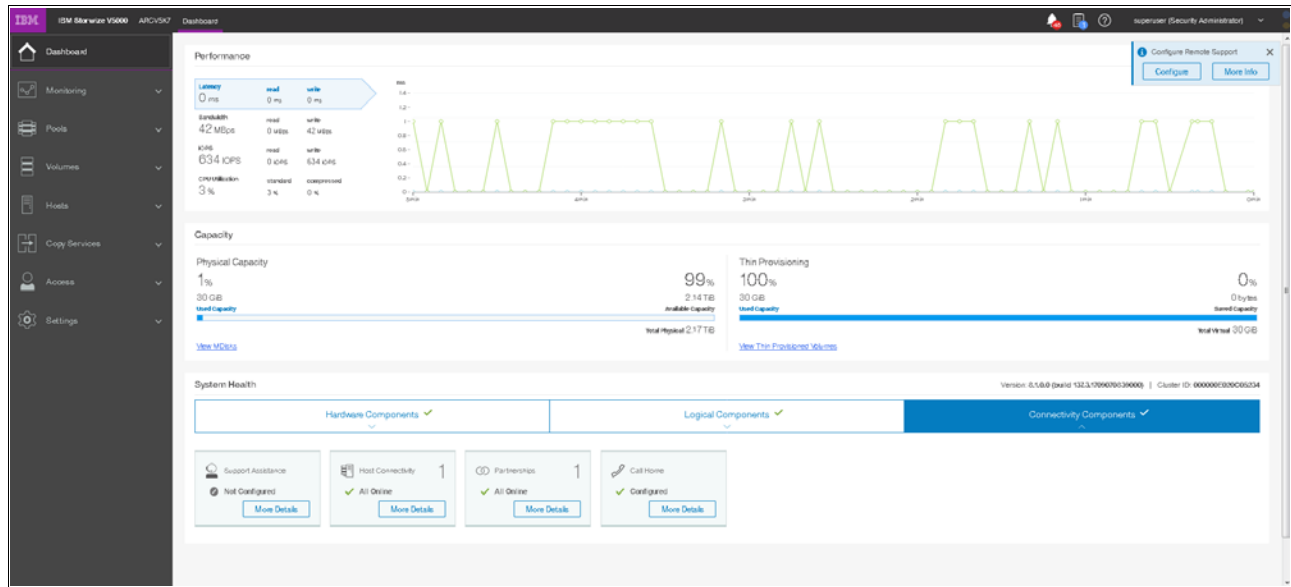


Figure 2-12 Setup wizard: Overview window

2.8.2 Command-line interface (CLI)

The command-line interface (CLI) is a flexible tool for system management that uses the Secure Shell (SSH) protocol. A public/private SSH key pair is optional for SSH access. The storage system can be managed by using the CLI, as shown in Example 2-1.

Example 2-1 System management by using the command-line interface

```
Lenovo Storage:V5030:superuser>svcinfo lsenclosureslot
enclosure_id slot_id port_1_status port_2_status drive_present drive_id
1 1 online online yes 10
1 2 online online yes 11
1 3 online online yes 15
1 4 online online yes 16
1 5 online online yes 12
1 6 online online yes 4
1 7 online online yes 7
1 8 online online yes 8
1 9 online online yes 9
1 10 online online yes 5
1 11 online online yes 18
1 12 online online yes 14
1 13 online online yes 13
1 14 online online yes 2
1 15 online online yes 6
1 16 online online yes 3
```

1	17	online	online	yes	1
1	18	online	online	yes	0
1	19	online	online	yes	20
1	20	online	online	no	
1	21	online	online	yes	19
1	22	online	online	yes	21
1	23	online	online	yes	22
1	24	online	online	yes	17
2	1	online	online	yes	25
2	2	online	online	yes	27
2	3	online	online	no	
2	4	online	online	yes	31
2	5	online	online	yes	24
2	6	online	online	yes	26
2	7	online	online	yes	33
2	8	online	online	yes	32
2	9	online	online	yes	23
2	10	online	online	yes	28
2	11	online	online	yes	29
2	12	online	online	yes	30

Lenovo Storage:V5030:superuser>

You can set up the initial Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems by using the process and tools that are described in 2.9, “First-time setup” on page 52.

2.9 First-time setup

This section describes how to set up a first-time Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 service and system.

Before you set up the initial Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, ensure that the system is powered on.

Power on: See the following information to check the power status of the system:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/v3700_system_leds.html

Set up the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems by using the technician Ethernet port:

1. Configure an Ethernet port on the personal computer to enable the Dynamic Host Configuration Protocol (DHCP) configuration of its IP address and Domain Name System (DNS) settings.

If you do not use DHCP, you must manually configure the personal computer. Specify the static IPv4 address 192.168.0.2, subnet mask 255.255.255.0, gateway 192.168.0.1, and DNS 192.168.0.1.

2. Locate the Ethernet port that is labeled T on the rear of the node canister.

On the Lenovo Storage V3700 V2 and V3700 V2 XP systems, the second Ethernet port is also used as the technician port, as shown in Figure 2-13 and Figure 2-14.

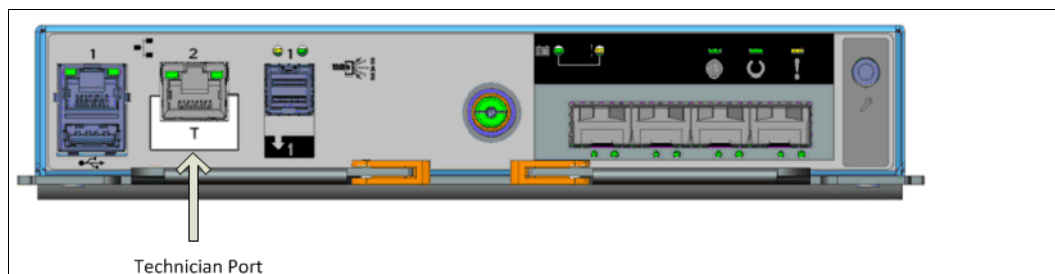


Figure 2-13 Lenovo Storage V3700 V2 technician port

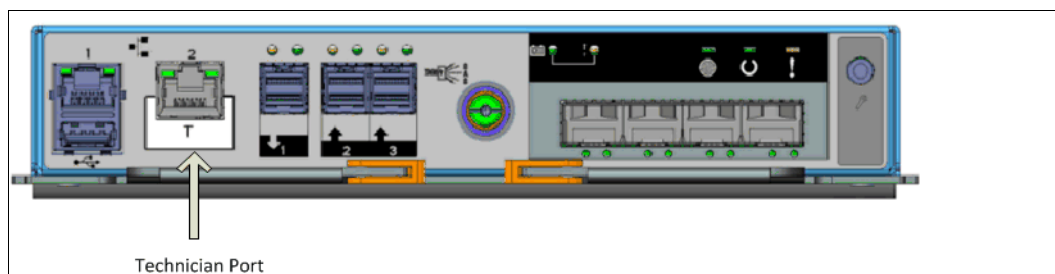


Figure 2-14 Lenovo Storage V3700 V2 XP technician port

The Lenovo Storage V5030 system uses a dedicated technician port, which is shown in Figure 2-15.

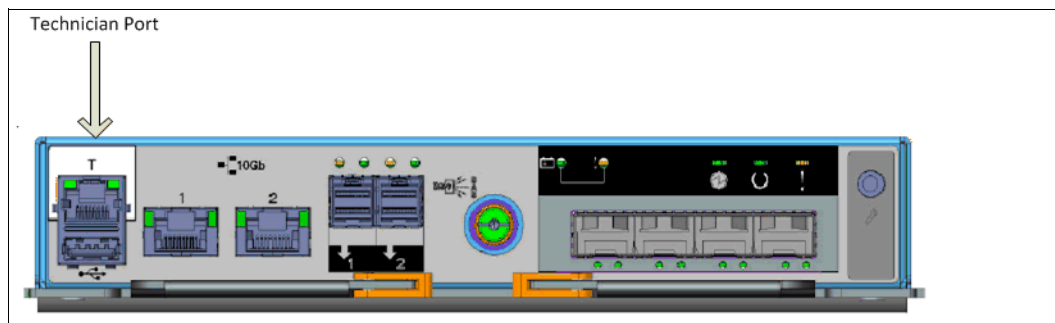


Figure 2-15 Lenovo Storage V5030 technician port

3. Connect an Ethernet cable between the port of the personal computer that is configured in step 2 and the technician port. After the connection is made, the system automatically

configures the IP address and DNS settings for the personal computer if DHCP is available. If it is not available, the system uses the values that you provided in step 1.

4. After the Ethernet port of the personal computer connects, open a supported browser and browse to address `http://install`. (If you do not have DHCP, open a supported browser and go to this static IP address: 192.168.0.1.) The browser is automatically directed to the initialization tool as shown in Figure 2-16.

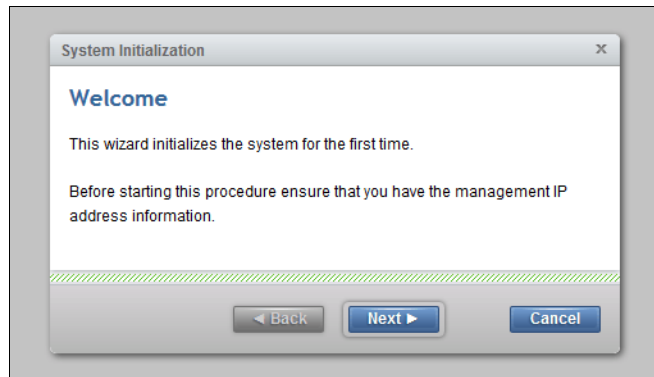


Figure 2-16 System initialization: Welcome

5. If you experience a problem when you try to connect due to a change in system states, wait 5 - 10 seconds and try again.
6. Click **Next**, as shown in Figure 2-17.

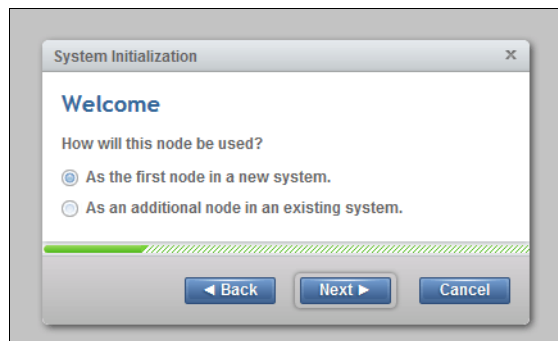


Figure 2-17 System initialization node usage

7. Choose the first option to set up the node as a new system and click **Next** to continue to the window that is shown in Figure 2-18.




Figure 2-18 System initialization: Create a New System

8. Complete all of the fields with the networking details for managing the system and click **Next**. When the task completes, as shown in Figure 2-19, click **Close**.




Figure 2-19 System initialization: Cluster creation

Note: The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI show the CLI as you go through the configuration steps.

9. The system takes approximately 10 minutes to reboot and reconfigure the Web Server as shown in Figure 2-20 on page 56. After this time, click **Next** to proceed to the final step.




Figure 2-20 System Initialization: Restarting Web Server

10. After you complete the initialization process, disconnect the cable between the personal computer and the technician port as instructed in Figure 2-21. Reestablish the connection to the customer network and click **Next** to be redirected to the management address that you provided to configure the system initially.

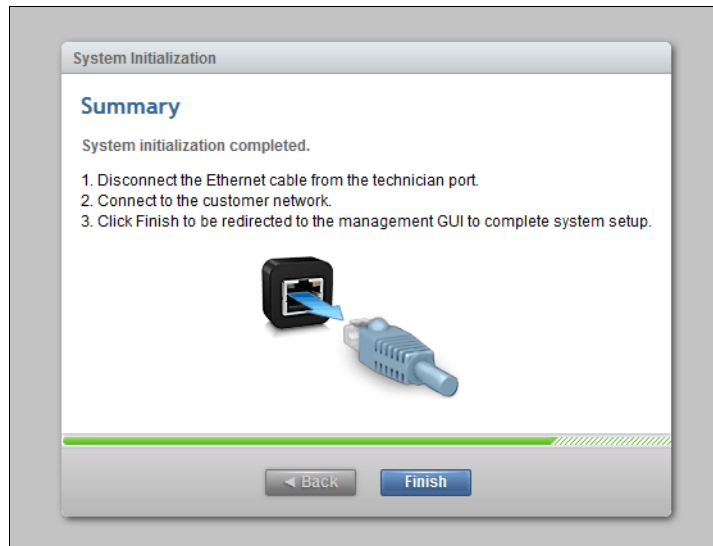


Figure 2-21 System initialization: Completion summary

2.10 Initial configuration

This section describes how to complete the initial configuration, including the following tasks:

- ▶ System components verification
- ▶ Email event notifications
- ▶ System name, date, and time settings
- ▶ License functions
- ▶ Initial storage configuration
- ▶ Initial configuration summary

If you completed the initial setup, that wizard automatically redirects you to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI. Otherwise, complete the following steps to complete the initial configuration process:

1. Start the service configuration wizard by using a web browser on a workstation and point it to the system management IP address that was defined in Figure 2-18 on page 55.

2. Enter a new secure password twice for the superuser user as shown in Figure 2-22 and click **Log in**.

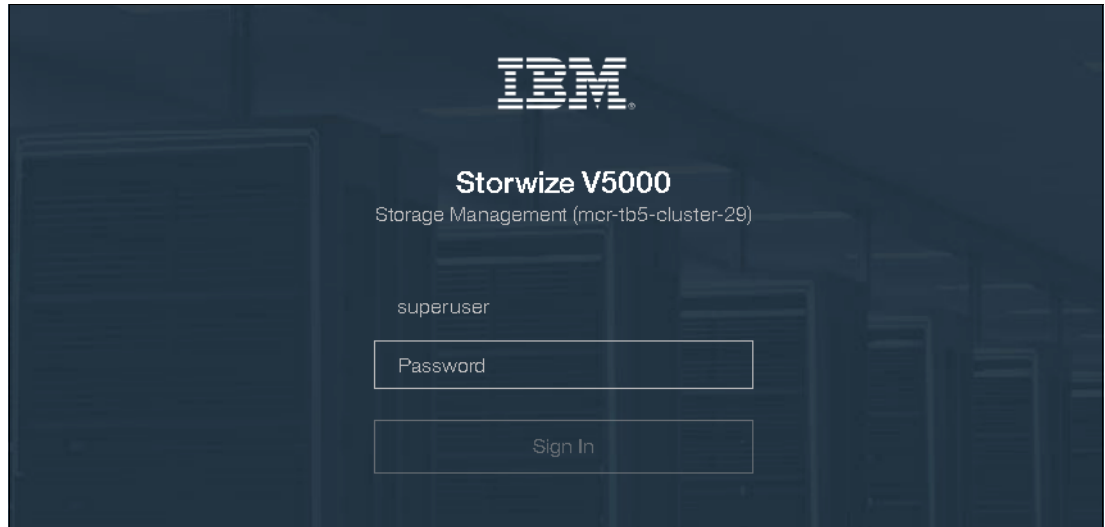


Figure 2-22 Setup wizard: Password prompt

3. Verify the prerequisites in the Welcome window as shown in Figure 2-23 and click **Next**.

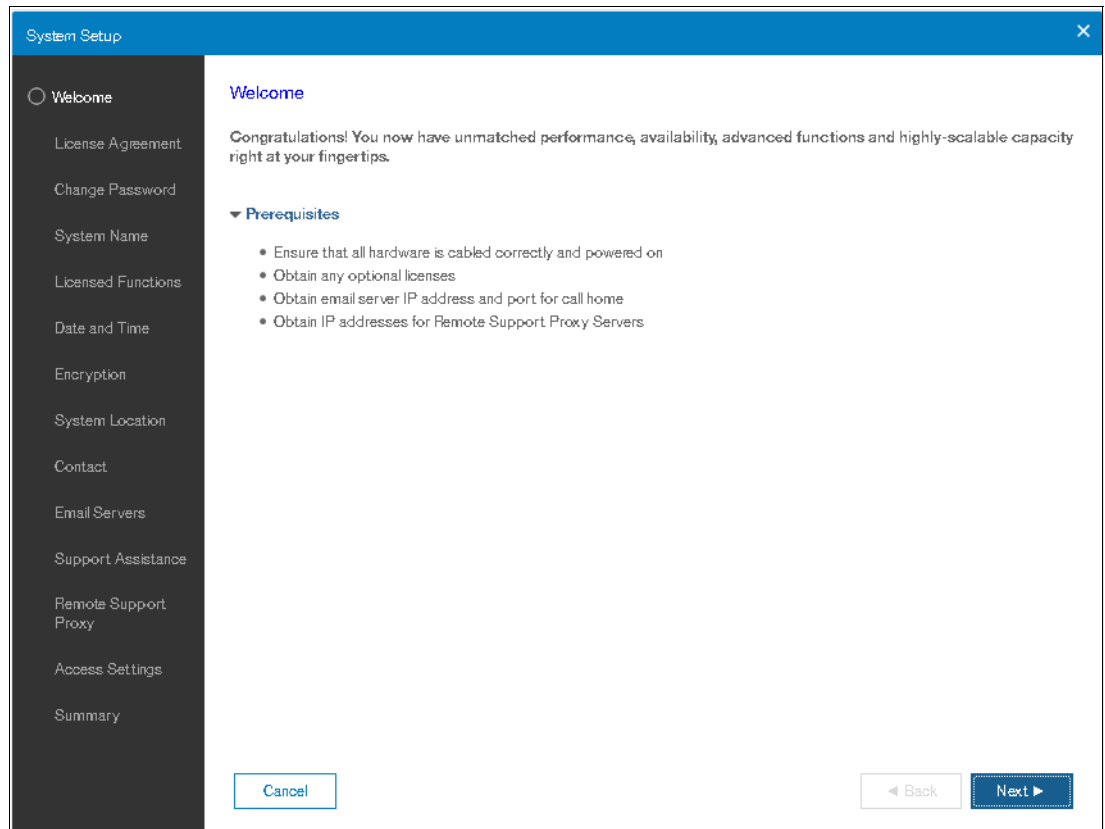


Figure 2-23 Setup wizard: Welcome

4. Accept the license agreement after reading it carefully as shown in Figure 2-24 on page 58 and click **Next**.

System Setup

- ✓ Welcome
- License Agreement
- Change Password
- System Name
- Licensed Functions
- Date and Time
- Encryption
- System Location
- Contact
- Email Servers
- Support Assistance
- Remote Support Proxy
- Access Settings
- Summary

Read the license agreement carefully.

License Addendum Java Notices Non-IBM Licenses Additional Licenses and Notices

International License Agreement for Early Release of Programs

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, CLICKING ON AN "ACCEPT" BUTTON, OR OTHERWISE USING THE PROGRAM, LICENSEE AGREES TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF LICENSEE, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND LICENSEE TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

* DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, CLICK ON AN "ACCEPT" BUTTON, OR USE THE PROGRAM;

* PROMPTLY RETURN THE UNUSED MEDIA AND DOCUMENTATION TO IBM. IF THE PROGRAM WAS DOWNLOADED, RETURN ALL COPIES OF THE PROGRAM.

1. Definitions

"Authorized Use" - the specified level at which licensee is authorized to execute or run the program. That level may be measured by number of users, millions of service units ("MSUs"), Processor Units ("PVUs"), or other level of use specified by IBM.

"Early Release" - a release of a Program for purposes of testing prior to it being made commercially available that may still be under development and therefore, potentially unreliable.

☒ I agree with the terms in the license agreement.

☐ I do not agree with the terms in the license agreement.

Figure 2-24 Setup wizard: License agreement

5. Change the password for superuser from the default as shown in Figure 2-25, then click **Apply and Next**.

System Setup

- ✓ Welcome
- ✓ License Agreement
- Change Password
- System Name
- Licensed Functions
- Date and Time
- Encryption
- System Location
- Contact
- Email Servers
- Support Assistance
- Remote Support Proxy
- Access Settings
- Summary

Change Password

The password must be reset before proceeding with system configuration.

User name: superuser

New password:

Confirm password:

Figure 2-25 Setup wizard: Change password

6. You will see password successfully changed message as shown in Figure 2-26.




Figure 2-26 Setup wizard: Password changed

7. In the System Name window, enter the system name as shown in Figure 2-27 and click **Apply and Next**.

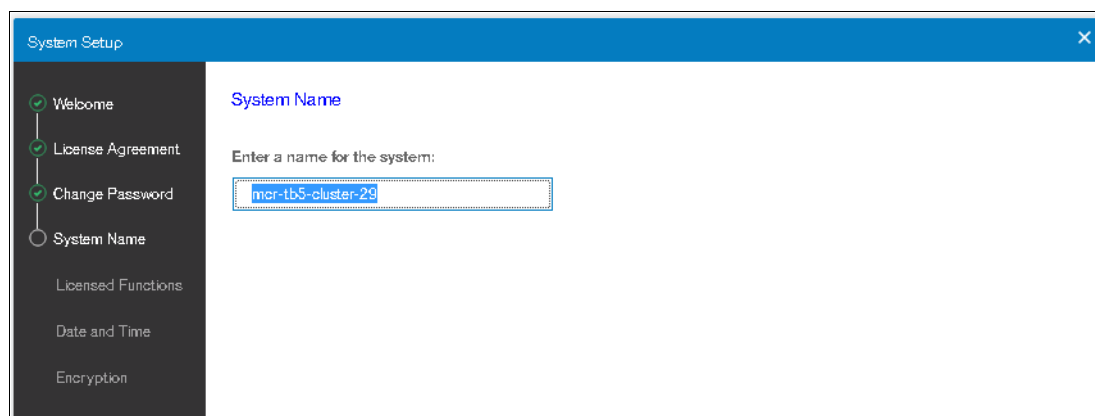


Figure 2-27 Setup wizard: System Name

Note: Use the `chsystem` command to modify the attributes of the clustered system. This command can be used any time after a system is created.

8. In the next window, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI provide help and guidance about additional licenses that are required for certain system functions. A license must be purchased for each enclosure that is attached to, or externally managed by, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. For each of the functions, enter the number of enclosures as shown in Figure 2-28 and click **Apply and Next**.

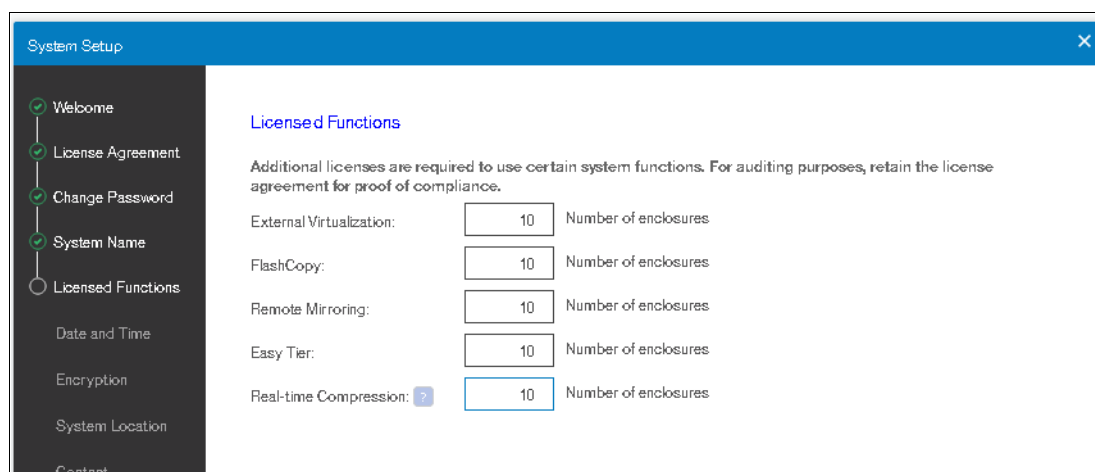


Figure 2-28 Setup wizard: Licensed Functions

The following actions are required for each of the licensed functions:

- FlashCopy: Enter the number of enclosures that are licensed to use FlashCopy function.
- Remote copy: Enter the number of Remote Mirroring licenses. This license setting enables the use of Metro Mirror and Global Mirror functions. This value must be equal to the number of enclosures that are licensed for external virtualization, plus the number of attached internal enclosures.
- Easy Tier: Enter the number of enclosures that are licensed to use Easy Tier function.
- External Virtualization: Enter the number of external enclosures that you are virtualizing. An external virtualization license is required for each physical enclosure that is attached to your system.
- Real-time Compression (RtC): Enter the number of enclosures that are licensed to use RtC.

Encryption license: The encryption feature that is available on the Lenovo Storage V3700 V2 XP and Lenovo Storage V5030 systems use a special licensing system that differs from the licensing system for the other features. Encryption requires a license key that can be activated in step 10.

9. Two options are available for configuring the date and time. Select the required method and enter the date and time manually or specify a network address for a Network Time Protocol (NTP) server. After this selection, the Apply and Next option becomes active, as shown in Figure 2-29. Click **Apply and Next**.

The screenshot shows the 'System Setup' wizard with a sidebar on the left containing the following steps: Welcome, License Agreement, Change Password, System Name, Licensed Functions, Date and Time (selected), Encryption, System Location, Contact, and Email Servers. The main content area is titled 'Date and Time' and includes the instruction: 'Select time and date settings. You can enter these settings manually or specify a Network Time Protocol (NTP) server to synchronize time on the system.' There are two radio buttons: 'Manually' and 'NTP Server' (which is selected). Below the radio buttons, there is an 'IP address:' label followed by a text input field containing '9.174.128.252'. Below that, there is a 'Time Zone:' label followed by a dropdown menu showing '(GMT) Dublin, Edinburgh, London, Lisbon'. In the bottom right corner, the current date and time are displayed as 'Sep 28, 2017, 5:40:19 PM'.

Figure 2-29 Setup wizard: Date and Time

10. If you purchased an Encryption License for a Lenovo V3700 V2 XP or Lenovo Storage V5030 system, select **Yes** as shown in Figure 2-30. One license is required for each control enclosure. Therefore, in a V5030 configuration with two I/O groups, two license keys are required.

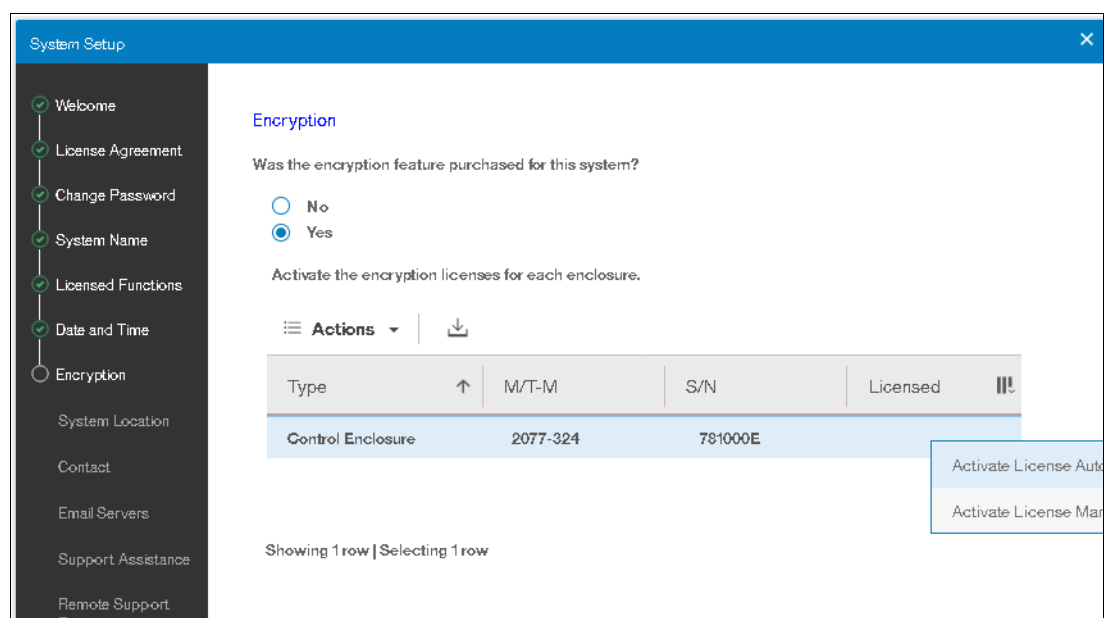


Figure 2-30 Setup wizard: Encryption feature

11. The easiest way to activate the encryption license is to highlight each enclosure that you want to activate the license for and choose **Actions** → **Activate License Automatically** and enter the authorization code that came with the purchase agreement for encryption. This action retrieves and applies a license key from ibm.com as shown in Figure 2-31.

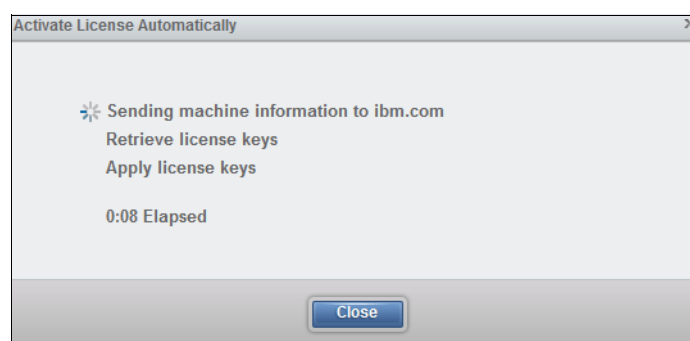


Figure 2-31 Setup wizard: Encryption license activation

12. If automatic activation cannot be performed, for example, if the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems are behind a firewall that prevents it from accessing the internet, choose **Actions** → **Activate License Manually**. Follow these steps:
- Go to this web page:
<https://www.ibm.com/storage/dsfa>
 - Select **Storwize**. Enter the machine type (6535 or 6536), serial number, and machine signature of the system. You can obtain this information by clicking **Need Help**.
 - Enter the authorization codes that were sent with your purchase agreement for the encryption function.

- d. Copy or download the key and paste it into the management GUI to activate the license.

13. When all licenses are active, click **Next** to set up the system location as shown in Figure 2-32.

The screenshot shows the 'System Setup' wizard with the 'System Location' step selected. The left sidebar lists steps: Welcome, License Agreement, Change Password, System Name, Licensed Functions, Date and Time, Encryption, System Location (selected), Contact, Email Servers, Support Assistance, Remote Support Proxy, Access Settings, and Summary. The main area is titled 'System Location' and contains the instruction: 'Service parts should be shipped to the same physical location as the system.' Below this are input fields for: Company name (IBM), System address (Maybrook House), City (Manchester), State or province (LAN), Postal code (M3 2EG), Country or region (United Kingdom), and Comment (third floor lab). At the bottom are 'Cancel', 'Back', and 'Next' buttons.

Figure 2-32 Setup wizard: system location

14. After entering the system location, click **Next** to set up the contact person for the system as shown in Figure 2-33, then click **Apply and Next**.

The screenshot shows the 'System Setup' wizard with the 'Contact' step selected. The left sidebar lists steps: Welcome, License Agreement, Change Password, System Name, Licensed Functions, Date and Time, Encryption, System Location, Contact (selected), and Email Servers. The main area is titled 'Contact' and contains the instruction: 'The support center contacts this person to resolve issues on the system.' Below this are input fields for: Name (James Whitaker), Email (james.whitaker@uk.ibm.com), Phone (primary) (+44-1234-5678), and Phone (alternate) (empty). At the bottom are 'Cancel', 'Back', and 'Next' buttons.

Figure 2-33 Setup wizard: contact person

15. You can configure your system to send email reports to field support if an issue is detected that requires hardware replacement. This function is called *Call Home*. When this email is received, field support automatically opens a problem report and contacts you to verify whether replacements parts are required.

Call Home: When Call Home is configured, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 automatically create a Support Contact with one of the following email addresses, depending on the country or region of installation:

- ▶ US, Canada, Latin America, and Caribbean Islands: `callhome1@de.ibm.com`
- ▶ All other countries or regions: `callhome0@de.ibm.com`

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can use Simple Network Management Protocol (SNMP) traps, syslog messages, and Call Home email to notify you and the Field Support Center when significant events are detected. Any combination of these notification methods can be used simultaneously.

To set up Call Home, you need the location details of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, Storage Administrator details, and at least one valid SMTP server IP address as shown in Figure 2-34.

Figure 2-34 Setup wizard: Email server details

Note: If you do not want to configure Call Home now, you can defer it using the check-box in the GUI and come back to it later via **Settings** → **Notifications**.

If your system is under warranty or you bought a hardware maintenance agreement, we advise you to configure Call Home to enable proactive support of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

To enter more than one email server, click the plus sign (+) icon, as shown in Figure 2-34. Then, click **Apply and Next** to commit.

16. The next window is for setting up support assistance if desired as shown in Figure 2-35.

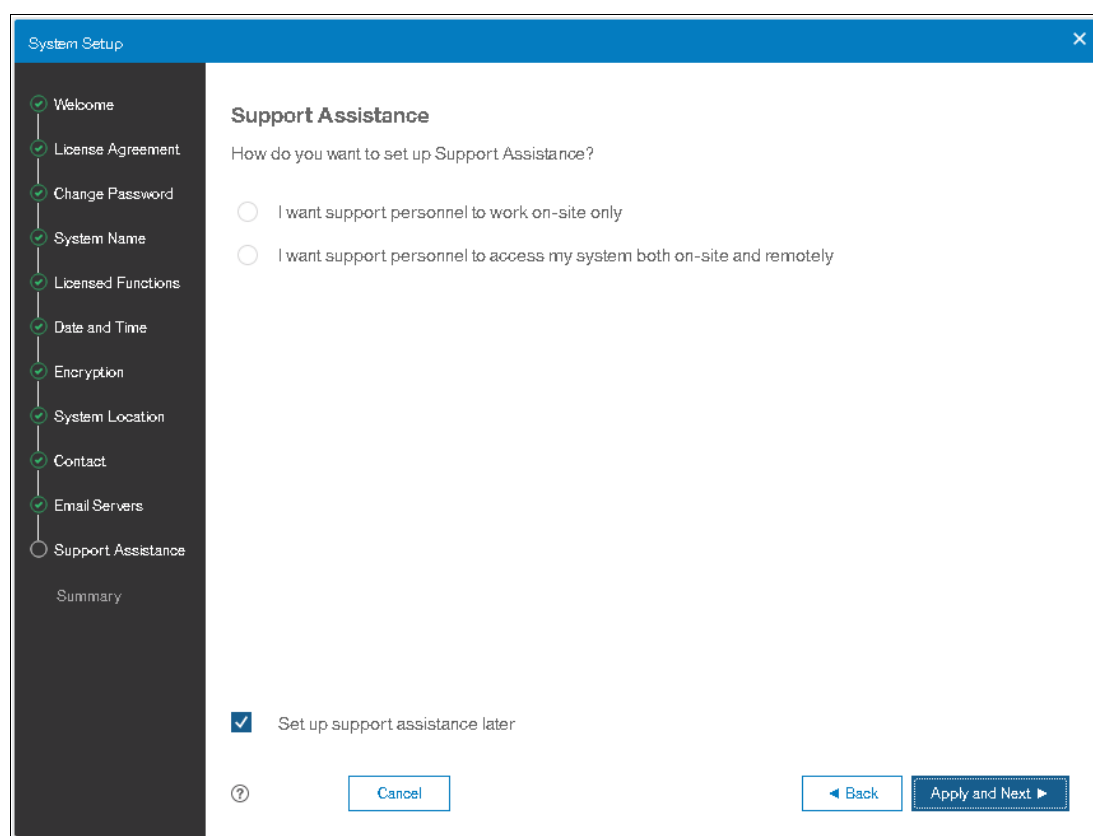


Figure 2-35 Initial setup: Support Assistance

In our setup, we chose to set up the support assistance later as it is covered extensively in Chapter 12, “RAS, monitoring, and troubleshooting” on page 625.

17. The Summary window for the contact details, system location, email server, Call Home, and email notification options is shown in Figure 2-36 on page 65.

2.10.1 Adding enclosures after the initial configuration

When the initial installation of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are complete, all expansion enclosures and control enclosures that were purchased at that time must be installed as part of the initial configuration. This process enables the system to make the best use of the enclosures and drives that are available.

Adding a control enclosure

If you are expanding the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 after the initial installation by adding a second I/O group (a second control enclosure), you must install it in the rack and connect it to the SAN. Ensure that you rezone your Fibre Channel switches so that the new control enclosure and the existing control enclosure are connected. For more information about zoning the node canisters, see 2.2, “SAN configuration planning” on page 40.

Note: Adding a second I/O group (via second controller enclosure) is only supported on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models.

After the hardware is installed, cabled, zoned, and powered on, a second control enclosure is visible from the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI, as shown in Figure 2-38.

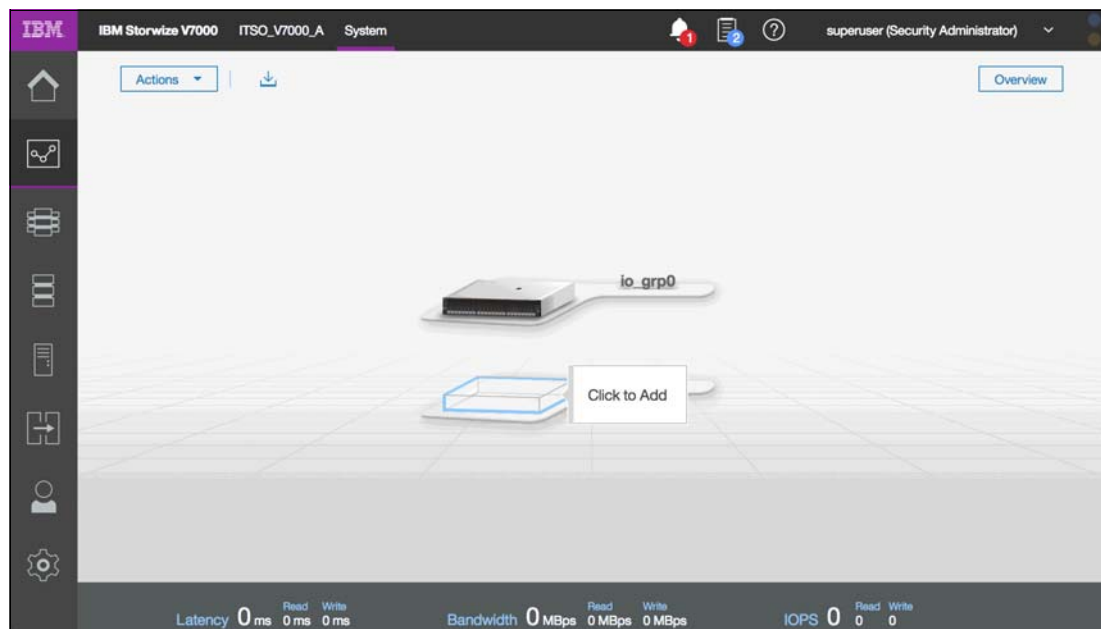


Figure 2-38 Second control enclosure

Complete the following steps to use the management GUI to configure the new enclosure:

1. In the main window, click **Actions** in the upper-left corner and select **Add Enclosures**.
Alternatively, you can click the available control enclosure as shown in Figure 2-39.

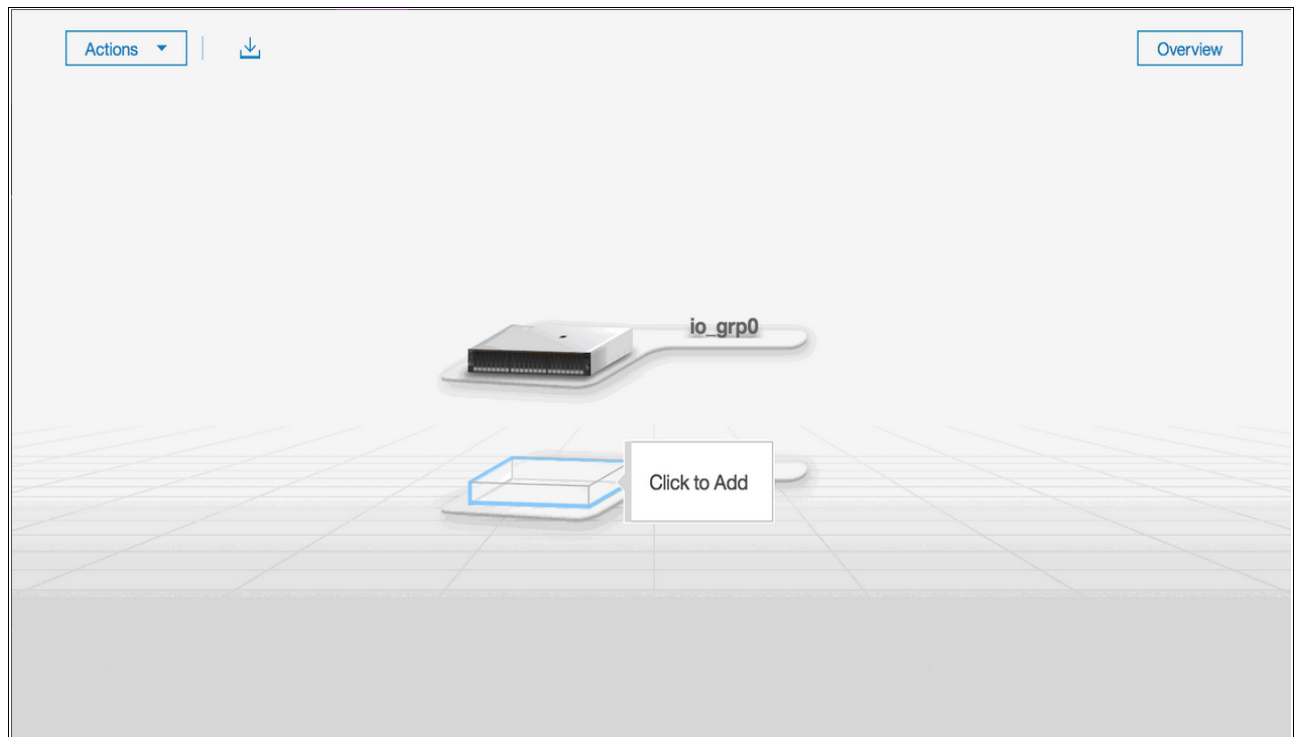


Figure 2-39 Option to add a control enclosure

2. If the control enclosure is configured correctly, the new control enclosure is identified in the next window, as shown in Figure 2-40.

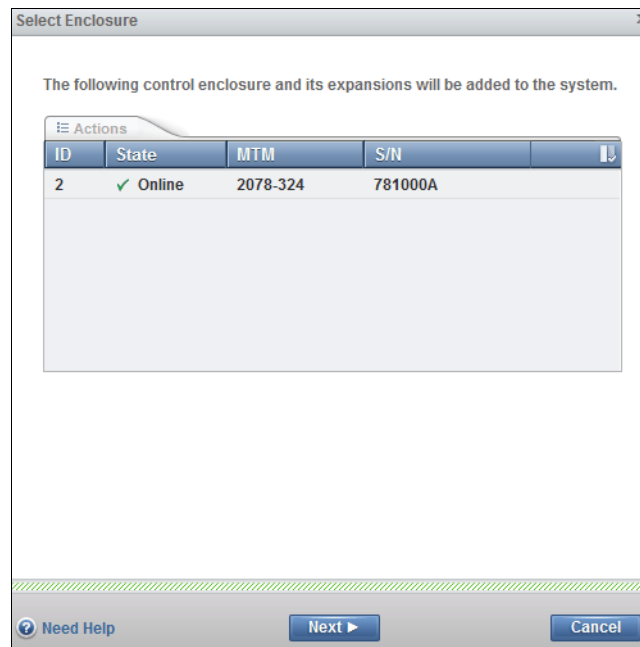


Figure 2-40 New control enclosure identification

3. Select the control enclosure and click **Actions** → **Identify** to turn on the identify LEDs of the new enclosure, if required. Otherwise, click **Next**.
4. The new control enclosure is added to the system as shown in Figure 2-41. Click **Finish** to complete the operation.

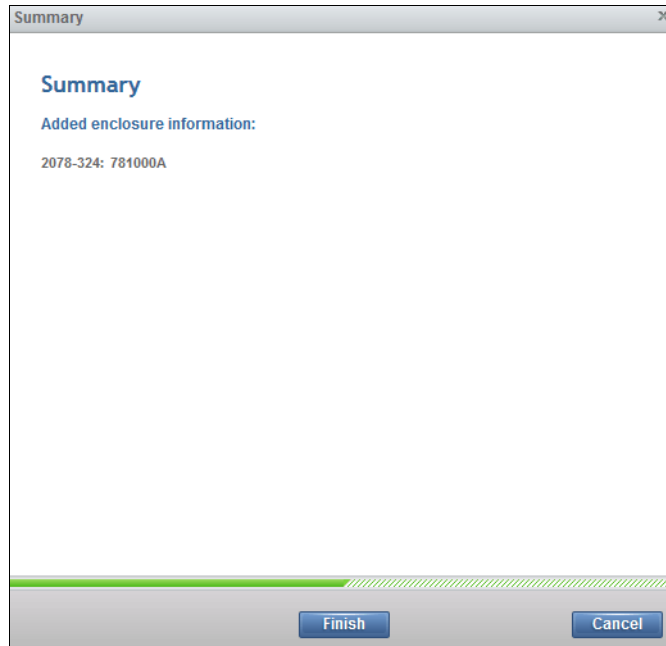


Figure 2-41 Added enclosure summary

5. When the new enclosure is added, the storage that is provided by the internal drives is available to use as shown in Figure 2-42.




Figure 2-42 Adding storage completed

6. After the wizard completes the addition of the new control enclosure, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 show the management GUI that contains two I/O groups, as shown in Figure 2-43.

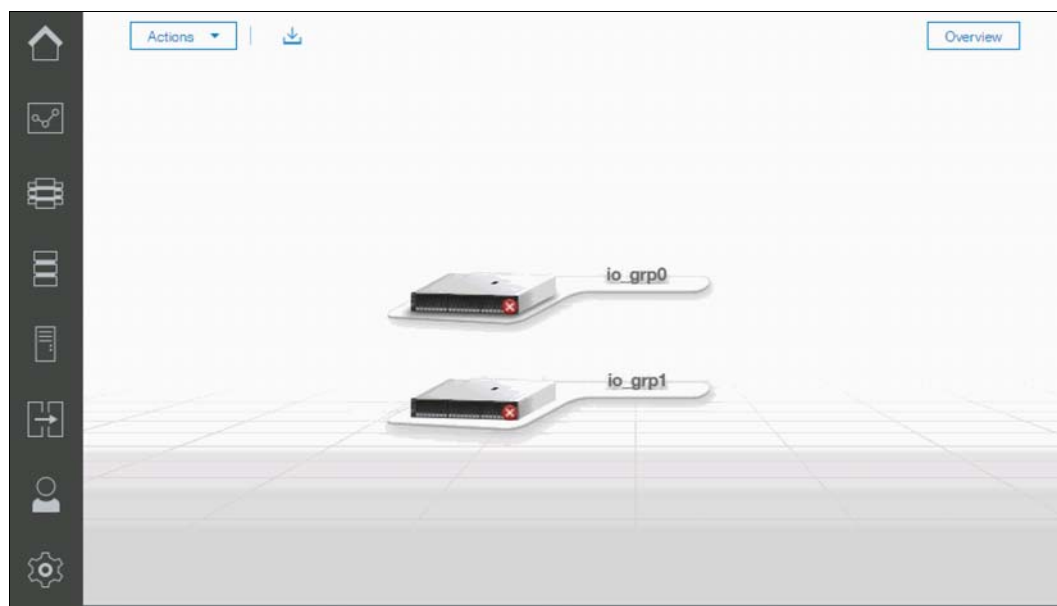


Figure 2-43 Lenovo Storage V Series GUI with two I/O groups

Adding an expansion enclosure

Complete the following steps to add an expansion controller:

1. To add an expansion enclosure, change to the **Monitoring** tab and select **System**. If no new hardware is shown, check your cabling to ensure that the new expansion enclosure is connected correctly and refresh the window.

In the main window, click **Actions** in the upper-left corner and select **Add Enclosures**. Alternatively, you can click the available expansion enclosure as shown in Figure 2-44.




Figure 2-44 Adding an expansion enclosure

2. If the enclosure is cabled correctly, the wizard identifies the candidate expansion enclosure. Select the expansion enclosure and click **Next**, as shown in Figure 2-45.

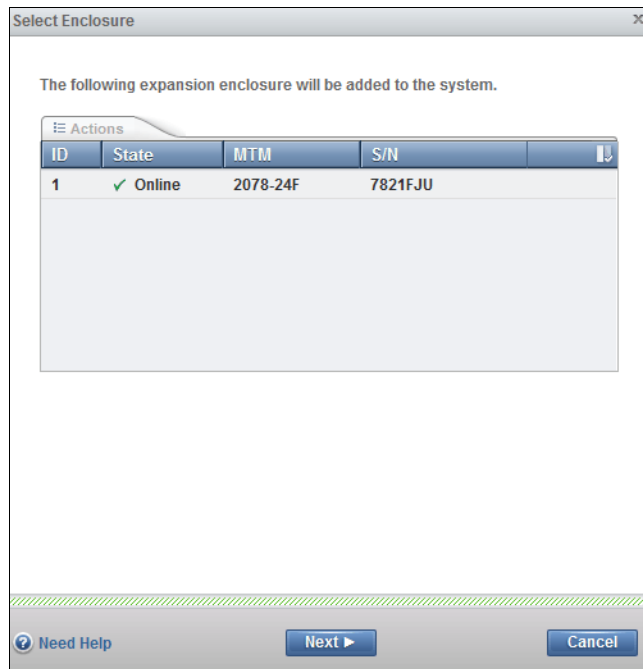


Figure 2-45 Expansion enclosure cable check

3. Select the expansion enclosure and click **Actions** → **Identify** to turn on the identify LEDs of the new enclosure, if required. Otherwise, click **Next**.
4. The new expansion enclosure is added to the system as shown in Figure 2-46. Click **Finish** to complete the operation.

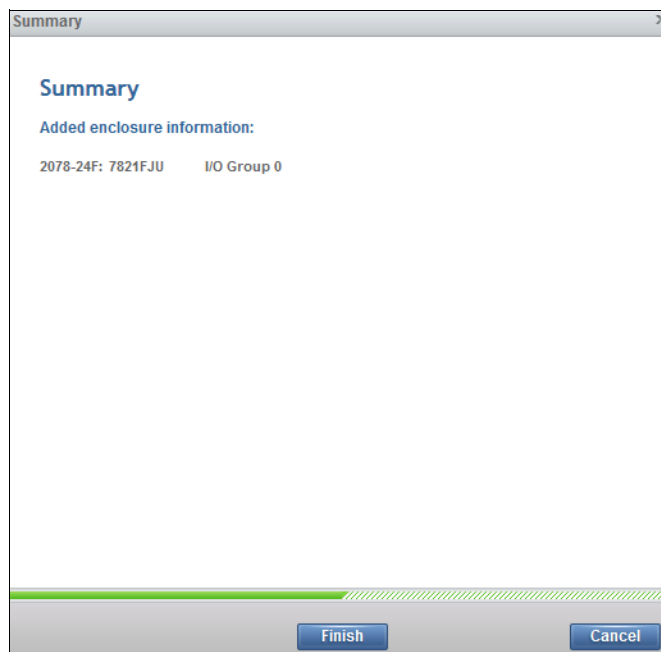


Figure 2-46 Added enclosure summary

5. After the expansion enclosure is added, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 show the management GUI that contains two enclosures, as shown in Figure 2-47.



Figure 2-47 Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI with two enclosures in a single I/O group

2.10.2 Service Assistant Tool

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, as a single I/O group, is configured initially with three IP addresses, one service IP address for each node canister, and a management IP address, which is set when the cluster is started.

The management IP and service IP addresses can be changed within the GUI as shown in Chapter 3, “Graphical user interface overview” on page 75.

IBM Service Assistant (SA) Tool is a web-based GUI that is used to service individual node canisters, primarily when a node has a fault and it is in a service state. A node cannot be active as part of a clustered system while the node is in a service state. The SA Tool is available even when the management GUI is not accessible. The following information and tasks are included:

- ▶ Status information about the connections and the node canister
- ▶ Basic configuration information, such as configuring IP addresses
- ▶ Service tasks, such as restarting the Common Information Model object manager (CIMOM) and updating the worldwide node name (WWNN)
- ▶ Details about node error codes and hints about how to fix the node error

Important: Service Assistant Tool can be accessed by using the superuser account only. You must access Service Assistant Tool under the direction of Lenovo Support only.

The Service Assistance GUI is available by using a service assistant IP address on each node. The SA GUI is accessed through the cluster IP addresses by appending service to the cluster management URL. If the system is down, the only other method of communicating with the node canisters is through the SA IP address directly. Each node can have a single

SA IP address on Ethernet port 1. We advise that these IP addresses are configured on all of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 node canisters.

The default IP address of canister 1 is 192.168.70.121 with a subnet mask of 255.255.255.0.

The default IP address of canister 2 is 192.168.70.122 with a subnet mask of 255.255.255.0.

To open the SA GUI, enter one of the following URLs into any web browser:

- ▶ `http(s)://cluster IP address of your cluster/service`
- ▶ `http(s)://service IP address of a node/service`

The following examples open the SA GUI:

- ▶ Management address: <http://1.2.3.4/service>
- ▶ SA access address: <http://1.2.3.5/service>

When you access SA by using the `<cluster address>/service`, the configuration node canister SA GUI login window opens, as shown in Figure 2-48.

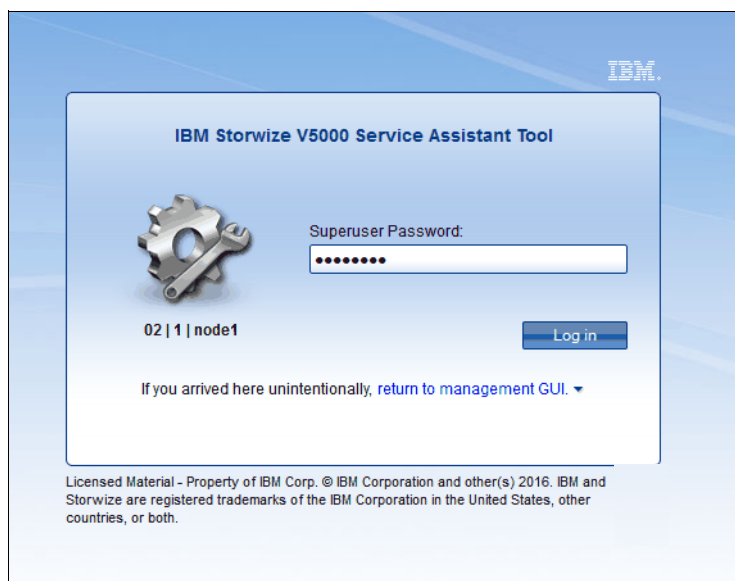


Figure 2-48 Service Assistant Tool login

The SA interface can view status and run service actions on other nodes and the node where the user is connected.

After you are logged in, you see the Service Assistant Tool Home window, as shown in Figure 2-49.

IBM Storwize V5000 Service Assistant Tool

Connected to: 02 | 1 | node1

Log out

IBM

Current: 02 | 1 | node1

Status: Active

Identify

Home

Collect Logs

Manage System

Recover System

Re-install Software

Update Manually

Configure Enclosure

Change Service IP

Configure CLI Access

Restart Service

Home

You can view detailed status and error summary, and manage service actions for the current node. The current node is the node on which service-related actions are performed. The connected node displays the service assistant and provides the interface for working with other nodes on the system. To manage a different node, select a node from the following table.

Attention:

Only perform service actions on nodes when directed by service procedures. If used inappropriately, service actions can cause a loss of access to data, or even data loss. If the node status is active, select Monitoring-->Events in the management GUI to fix any errors that are related to the active node.

Actions:

Enter Service State

GO

Change Node

Node Name	Node Status	Error	Panel	System	Site	Relationship
<input checked="" type="radio"/> node1	Active		02-1	ITSO_V5000		Local
<input type="radio"/> node2	Active		02-2	ITSO_V5000		Partner

Refresh

Node Errors

Node Detail

Node	Hardware	Access	Location	Ports
Node ID:	1			
Node Name:	node1			
Node Status:	Active			
Part Identity:	11SZZTB5C2Y000005AX00W			
Node FRU:	ZZTB5CG			
Configuration Node:	Yes			
Model:	T5M			
System:	ITSO_V5000			
Site Name:				
System Software Build:	124.4.1602111021000			
Software Version:	7.6.1.0			
Software Build:	124.4.1602111021000			
Console IP:	9.174.157.121:443			
Has File Module Key:	No			

Figure 2-49 Service Assistant Tool Home window

The current canister node is displayed in the upper-left corner of the GUI. As shown in Figure 2-49, the current canister node is node 1. To change the canister, select the relevant node in the Change Node section of the window. You see that the details in the upper-left corner change to reflect the new canister.

The SA GUI provides access to service procedures and shows the status of the node canisters. We advise that you perform these procedures only if you are directed to use them by IBM Support.

For more information about how to use the SA Tool, see this web page:

<https://ibm.biz/BdjSJq>

Chapter 2. Initial configuration 73

Graphical user interface overview

This chapter provides an overview of the graphical user interface (GUI) on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and shows you how to use the navigation tools.

Specifically, this chapter provides information about the following topics:

- ▶ 3.1, “Overview of management software” on page 76
- ▶ 3.2, “Overview pane” on page 85
- ▶ 3.3, “Monitoring menu” on page 85
- ▶ 3.4, “Pools menu” on page 93
- ▶ 3.5, “Volumes menu” on page 106
- ▶ 3.6, “Hosts menu” on page 111
- ▶ 3.7, “Copy services” on page 117
- ▶ 3.8, “Access menu” on page 125
- ▶ 3.9, “Settings menu” on page 128

3.1 Overview of management software

A GUI can simplify storage management and provide a fast and more efficient management tool. V8.1 GUI has significant changes from previous versions, such as the icons, color palette, object locations, and more. However, usability is a priority, and usability is maintained in the GUI.

JavaScript: You must enable JavaScript in your browser. For Mozilla Firefox, JavaScript is enabled by default and requires no additional configuration. For more information about configuring your web browser, go to this web page:

<https://ibm.biz/BdjS9Z>

3.1.1 Access to the storage management software

To access the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, complete the following steps:

1. To log on to the management software, type the IP address that was set during the initial setup process into the address line of your web browser. You can connect from any workstation that can communicate with the system. The login window opens (Figure 3-1).




Figure 3-1 Login window

We suggest that each user has an account that is not shared with someone else. The default user accounts need to be unavailable for remote access, or the passwords need to be changed from the default password and known only to the system owner or kept secured for emergency purposes only. This approach helps to identify the personnel who are working on the device and to track all of the important changes in the systems. The *Superuser* account must be used for initial configuration only.

- After a successful login, the System pane displays the Dashboard with all relevant details of your system. Shown in Figure 3-2.

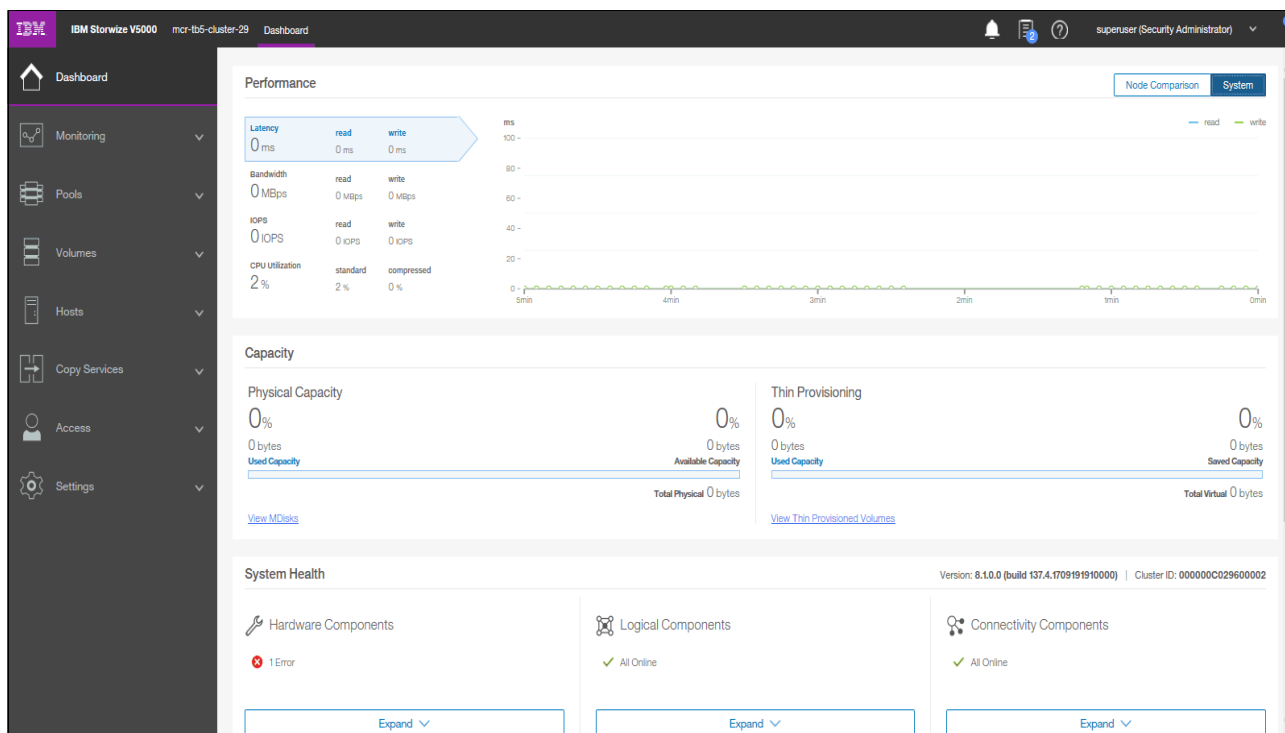


Figure 3-2 A first view

The System pane is an important user interface. In the remaining chapters, we do not explain how to access it each time.

3.1.2 System pane layout

The System pane has four main sections for navigating through the management tool:

- Top

The top menu shows the *navigation path* so that the user knows the exact path.

The Actions option can be used at any time to add more enclosures, modify existing hardware, or rename the system.

The top menu also has a *system overview* (upper right corner) so that you can see global system topology.

- Left

The *system menu* (or set of *function icons*) is the main object that the users access. By using the system menu, the user can change or view any setting or parameter in the system.

- Center

Within the *system view* object, users can view or change parameters that mainly relate to the hardware of global system settings.

- Bottom

The *informational pane* consists of running tasks, capacity information, basic performance indicator (new since version 7.5), system health status, and status alerts.

Figure 3-3 shows these main areas.

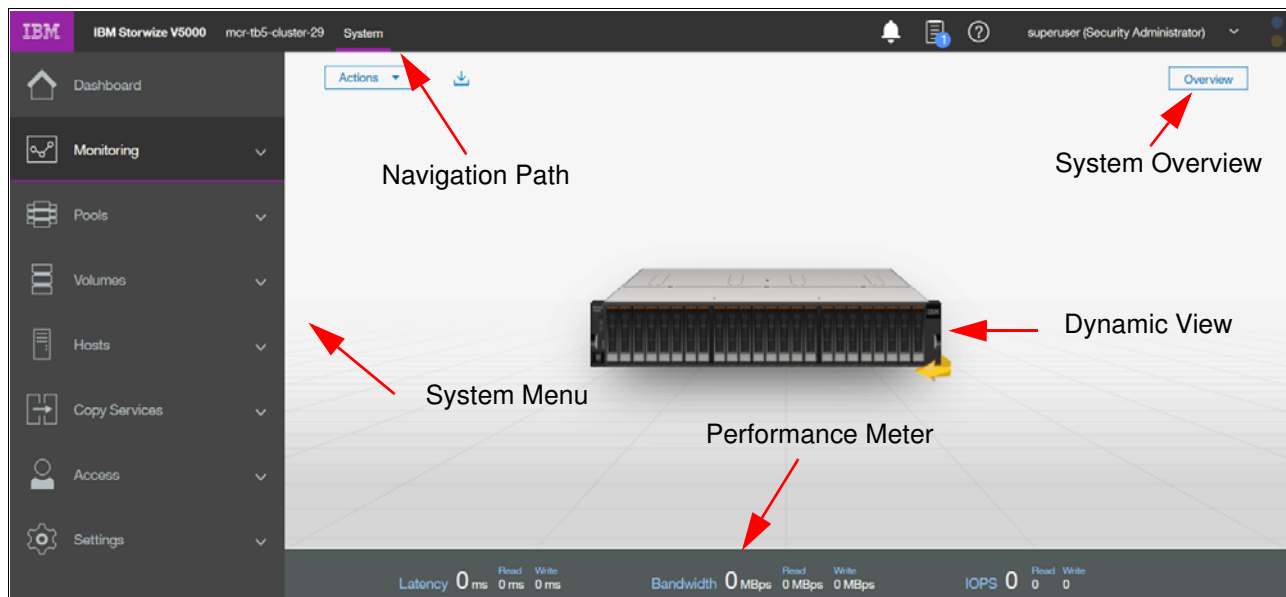


Figure 3-3 Main areas

The main areas are described:

- The left side of the window shows eight *function icons*. We refer to them collectively as a *dynamic menu*. The dynamic menu includes these function icons:
 - Dashboard
 - Monitoring menu
 - Pools menu
 - Volumes menu
 - Hosts menu
 - Copy Services menu
 - Access menu
 - Settings menu

Figure 3-4 on page 79 shows the dynamic menu.

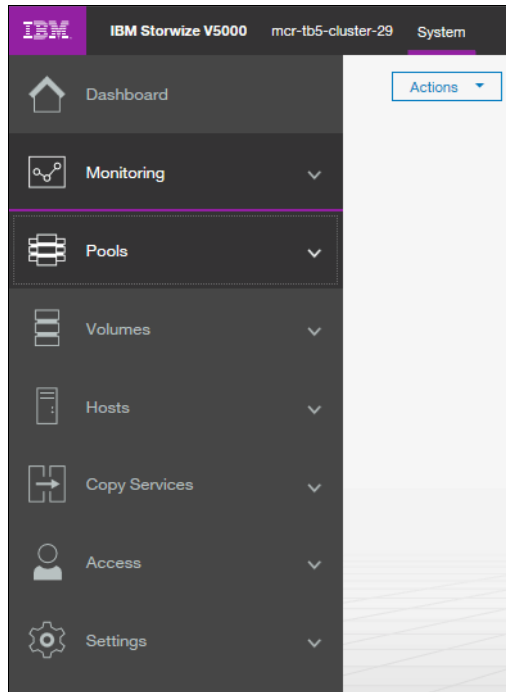


Figure 3-4 Settings menu

The middle of the window shows a component model of the existing configuration. Hovering the mouse cursor over each component and its part highlights that part and provides a pop-up menu with a description that identifies important parameters and functions of this element. To see the rear of the component, you can dynamically rotate them as in a typical 360° view. Right-clicking a component or its part opens its menu with actions, which are normally available from the dynamic menu on the left or from the Actions Button in the upper left corner. Figure 3-5 shows the component model.



Figure 3-5 Component model

The bottom of the window shows the performance indicator. It gives you information how your machine performs right now. Information covers only the external tasks for attached hosts. Figure 3-6 shows the performance indicator. How the performance internally is, shows the System statistic page as in Figure 3-7 on page 80.




Figure 3-6 Performance indicator

Figure 3-7 shows the System statistic page.

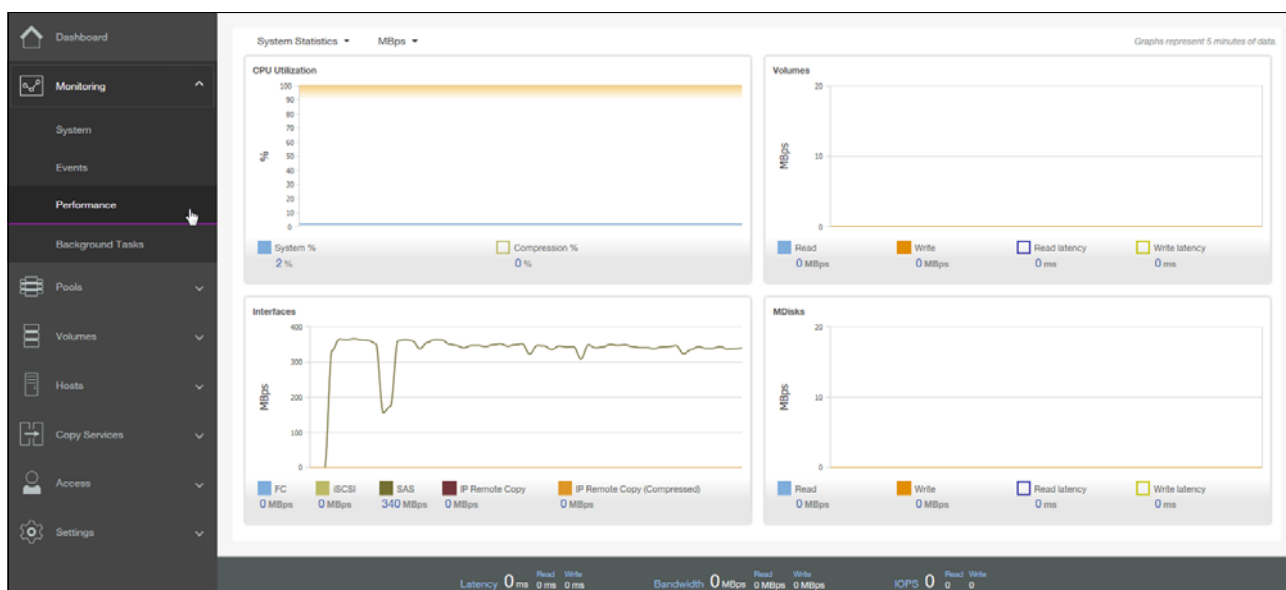


Figure 3-7 System statistic page

On the right upper area you get more information about the health of your system, as the Event button shows it in Figure 3-8,

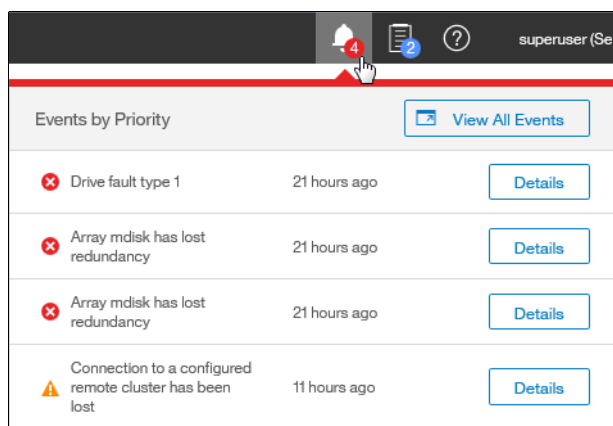


Figure 3-8 Event button

Or the Suggested tasks button shown in Figure 3-9 on page 81.

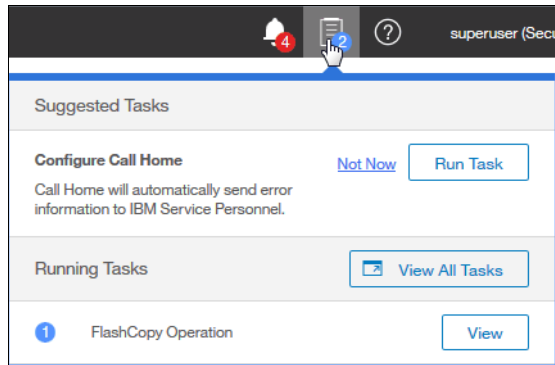


Figure 3-9 Suggested Tasks

A help menu is shown in the next picture. See Figure 3-10.




Figure 3-10 Help Menu

The bottom of the window shows five status indicators. Clicking any of them provides more detailed information about the existing configuration, situation or status of the solution. Click any of these function icons to expand them and minimize them as required, or to switch between different types of information, for example, virtual or allocated storage. In an error or warning situation, those indicators are extended by the status alerts icon in the upper-right corner See Figure 3-8 on page 80.

3.1.3 Navigation

Navigating in the management tool is simple. You can click with the cursor on one of the eight function icons and display a sub menu of options. You can move the cursor to an option and select it. Figure 3-11 on page 82 shows how to access, for example, the **Pools** option.

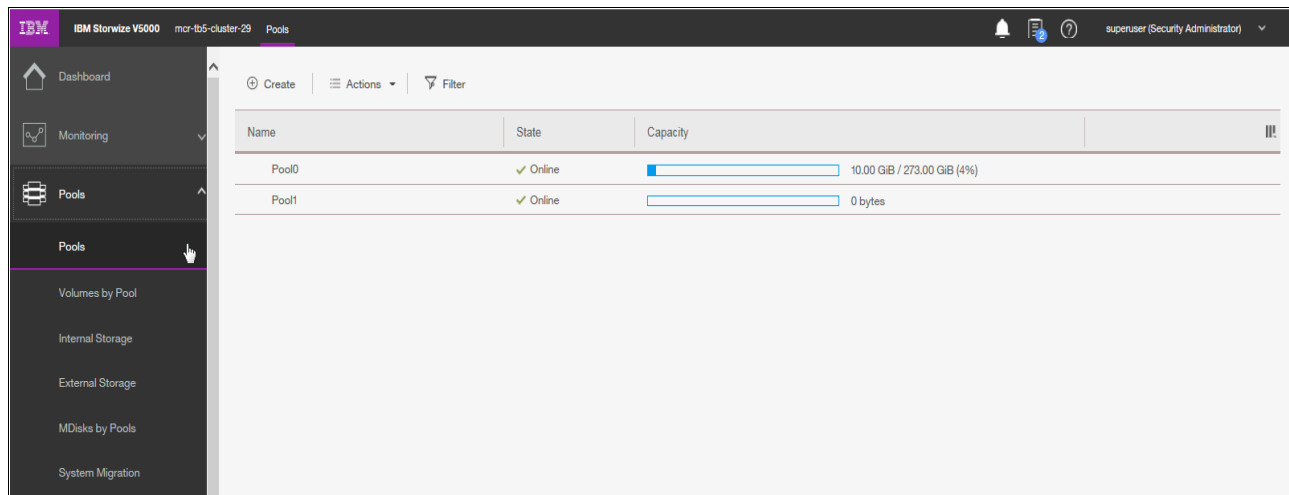


Figure 3-11 Navigate by using the menu options

3.1.4 Multiple selection

With the improved management tool, you can select multiple items by using the Shift and Ctrl keys. To select multiple items in a display, click the first item, press and hold the Shift key, and click the last item that you require in the list. All rows between those two items are selected and highlighted in light blue (Figure 3-12).

<div> + Create Volumes ≡ Actions ▼ All Volumes ▼ Filter </div> <div>Showing 21 volumes Selecting 14 volumes (14.00 GiB)</div>							
Name	State	Synchronized	Pool	UID	Host Mappings	Capacity	
V50000	✓ Online		Pool0	6005076300A580000800000000000001	No	1.00 GiB	
V50001	✓ Online		Pool0	6005076300A580000800000000000002	No	1.00 GiB	
V50002	✓ Online		Pool0	6005076300A580000800000000000003	No	1.00 GiB	
V50003	✓ Online		Pool0	6005076300A580000800000000000004	No	1.00 GiB	
V50004	✓ Online		Pool0	6005076300A580000800000000000005	No	1.00 GiB	
V50005	✓ Online		Pool0	6005076300A580000800000000000006	No	1.00 GiB	
V50006	✓ Online		Pool0	6005076300A580000800000000000007	No	1.00 GiB	
V50007	✓ Online		Pool0	6005076300A580000800000000000008	No	1.00 GiB	
V50008	✓ Online		Pool0	6005076300A580000800000000000009	No	1.00 GiB	
V50009	✓ Online		Pool0	6005076300A58000080000000000000A	No	1.00 GiB	
V50010	✓ Online		Pool0	6005076300A58000080000000000000B	No	1.00 GiB	
V50011	✓ Online		Pool0	6005076300A58000080000000000000C	No	1.00 GiB	
V50012	✓ Online		Pool0	6005076300A58000080000000000000D	No	1.00 GiB	
V50013	✓ Online		Pool0	6005076300A58000080000000000000E	No	1.00 GiB	
V50014	✓ Online		Pool0	6005076300A58000080000000000000F	No	1.00 GiB	
V50015	✓ Online		Pool0	6005076300A580000800000000000010	No	1.00 GiB	
V50016	✓ Online		Pool0	6005076300A580000800000000000011	No	1.00 GiB	
V50017	✓ Online		Pool0	6005076300A580000800000000000012	No	1.00 GiB	

Figure 3-12 Multiple selections by using the Shift key

Similarly, if you want to select multiple items that are not in sequential order, click the first item, press and hold the Ctrl key, and click the other items that you need (Figure 3-13).

Name	State	Synchronized	Pool	UID	Host Mappings	Capacity	
V50000	✓ Online		Pool0	6005076300A580000800000000000001	No	1.00 GiB	⌵
V50001	✓ Online		Pool0	6005076300A580000800000000000002	No	1.00 GiB	⌵
V50002	✓ Online		Pool0	6005076300A580000800000000000003	No	1.00 GiB	⌵
V50003	✓ Online		Pool0	6005076300A580000800000000000004	No	1.00 GiB	⌵
V50004	✓ Online		Pool0	6005076300A580000800000000000005	No	1.00 GiB	⌵
V50005	✓ Online		Pool0	6005076300A580000800000000000006	No	1.00 GiB	⌵
V50006	✓ Online		Pool0	6005076300A580000800000000000007	No	1.00 GiB	⌵
V50007	✓ Online		Pool0	6005076300A580000800000000000008	No	1.00 GiB	⌵
V50008	✓ Online		Pool0	6005076300A580000800000000000009	No	1.00 GiB	⌵
V50009	✓ Online		Pool0	6005076300A58000080000000000000A	No	1.00 GiB	⌵
V50010	✓ Online		Pool0	6005076300A58000080000000000000B	No	1.00 GiB	⌵
V50011	✓ Online		Pool0	6005076300A58000080000000000000C	No	1.00 GiB	⌵
V50012	✓ Online		Pool0	6005076300A58000080000000000000D	No	1.00 GiB	⌵
V50013	✓ Online		Pool0	6005076300A58000080000000000000E	No	1.00 GiB	⌵
V50014	✓ Online		Pool0	6005076300A58000080000000000000F	No	1.00 GiB	⌵
V50015	✓ Online		Pool0	6005076300A580000800000000000010	No	1.00 GiB	⌵
V50016	✓ Online		Pool0	6005076300A580000800000000000011	No	1.00 GiB	⌵
V50017	✓ Online		Pool0	6005076300A580000800000000000012	No	1.00 GiB	⌵

Figure 3-13 Multiple selections by using the Ctrl key

Another option for selecting volumes is selecting by mask. To select all volumes that have “V5000” in their name, click **Filter**, type V5000, and press Enter. All volumes with “V5000” in their name display. After you filter the volumes, you can easily select all of the displayed volumes or a subset by using the Ctrl key or Shift key technique that was explained previously (Figure 3-14).

Name	State	Synchronized	Pool	UID	Host Mappings	Capacity	
V50000	✓ Online		Pool0	6005076300A580000800000000000001	No	1.00 GiB	⌵
V50001	✓ Online		Pool0	6005076300A580000800000000000002	No	1.00 GiB	⌵
V50002	✓ Online		Pool0	6005076300A580000800000000000003	No	1.00 GiB	⌵
V50003	✓ Online		Pool0	6005076300A580000800000000000004	No	1.00 GiB	⌵
V50004	✓ Online		Pool0	6005076300A580000800000000000005	No	1.00 GiB	⌵
V50005	✓ Online		Pool0	6005076300A580000800000000000006	No	1.00 GiB	⌵
V50006	✓ Online		Pool0	6005076300A580000800000000000007	No	1.00 GiB	⌵
V50007	✓ Online		Pool0	6005076300A580000800000000000008	No	1.00 GiB	⌵
V50008	✓ Online		Pool0	6005076300A580000800000000000009	No	1.00 GiB	⌵
V50009	✓ Online		Pool0	6005076300A58000080000000000000A	No	1.00 GiB	⌵
V50010	✓ Online		Pool0	6005076300A58000080000000000000B	No	1.00 GiB	⌵
V50011	✓ Online		Pool0	6005076300A58000080000000000000C	No	1.00 GiB	⌵
V50012	✓ Online		Pool0	6005076300A58000080000000000000D	No	1.00 GiB	⌵
V50013	✓ Online		Pool0	6005076300A58000080000000000000E	No	1.00 GiB	⌵
V50014	✓ Online		Pool0	6005076300A58000080000000000000F	No	1.00 GiB	⌵
V50015	✓ Online		Pool0	6005076300A580000800000000000010	No	1.00 GiB	⌵
V50016	✓ Online		Pool0	6005076300A580000800000000000011	No	1.00 GiB	⌵
V50017	✓ Online		Pool0	6005076300A580000800000000000012	No	1.00 GiB	⌵

Figure 3-14 Filtering volumes

3.1.5 Status indicators area

The status indicators area at the bottom of the System pane (Figure 3-15) shows now only a high-level status of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems. Information shown there covers only the performance of the attached Host Systems.



Figure 3-15 Status Indicators

Help

Another useful interface feature is the integrated help function. You can access help for certain fields and objects by moving the mouse cursor over the question mark (?) icon (Figure 3-16) next to the field. Panel-specific help is available by clicking **Need Help** or by using the **Help** link in the upper-right corner of the GUI.




Figure 3-16 Access to panel-specific help

3.2 Overview pane

The welcome pane of the GUI changed from the well-known former Overview pane to the new System pane, as shown in Figure 3-17. Clicking **Overview** in the upper-right corner of the System pane opens a modified Overview pane with similar functionality as in previous versions of the software.

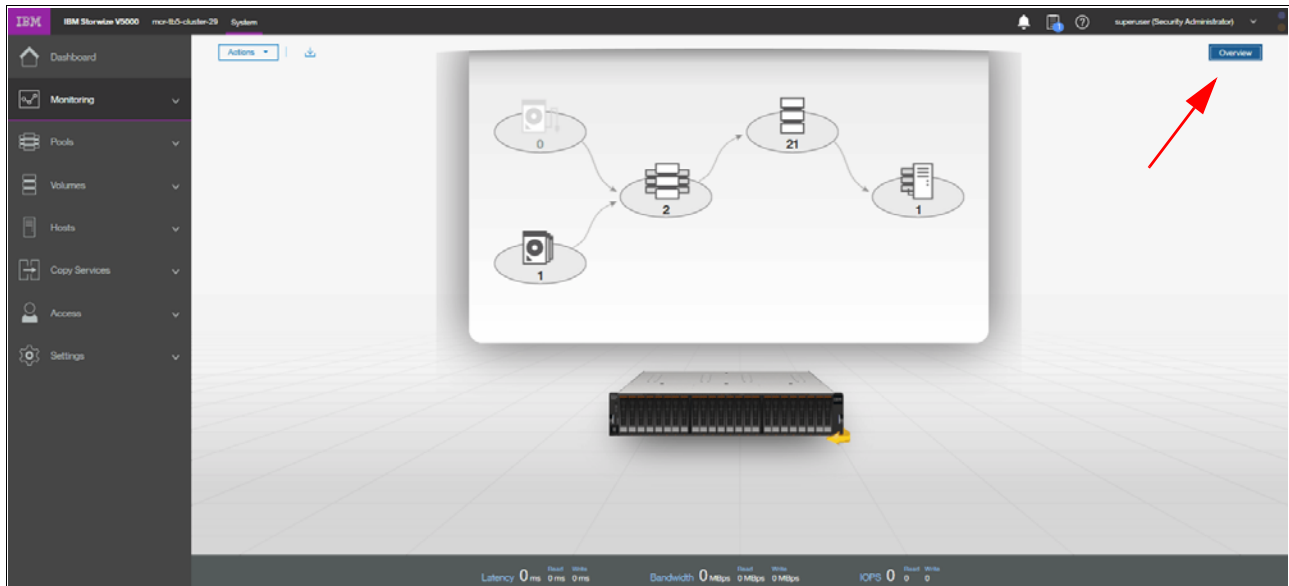


Figure 3-17 Opening the Overview pane

See 3.1.2, “System pane layout” on page 77 to understand the structure of the pane and how to navigate to various system components and manage them more efficient and quickly.

3.3 Monitoring menu

Place the cursor over the Monitoring function icon to open the Monitoring menu (Figure 3-18 on page 86). The Monitoring menu offers the following navigation directions:

► System

From the System panel, you can view details on control and expansion enclosures and various hardware components of the system. The system uses base-2 (binary numeral) as capacity indicators for volumes, drives, and other system objects. The management GUI and the command-line interface (CLI) use different abbreviations to indicate capacity, but the value for these capacity indicators is the same

See 3.3.1, “System overview” on page 87.

► Events

The Events panel displays two types of events: messages and alerts, and indicates the cause of any log entry. You can use this panel to filter messages that are related to events that occurred on the system. Some alerts have a four-digit error code and a fix procedure that helps you fix the problem. Other alerts also require action but do not have a fix procedure. Messages are fixed when you acknowledge reading them and mark them as fixed. Each event has a time stamp that indicates when the action occurred or the command was submitted on the system. When logs are displayed in the command-line interface, the time stamps for the logs in CLI are the system time. However, when logs are

displayed in the management GUI, the time stamps are translated to the local time where the web browser is running.

See 3.3.3, “Events” on page 91.

► **Performance**

Use real-time statistics to monitor CPU utilization, volume, interface, and MDisk bandwidth of your system and nodes. Each graph represents 5 minutes of collected statistics and provides a means of assessing the overall performance of your system. 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 then can 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 for Fibre Channel or SAS interfaces, to determine whether the host data-transfer rate is different from the expected rate. You can also select node-level statistics, which can help you determine the performance impact of a specific node. As with system statistics, node statistics help you to evaluate whether the node is operating within normal performance metrics.

See 3.3.4, “Performance” on page 92.

► **Background Tasks**

Use the Background Tasks page to view and manage current tasks that are running on the system. The Background Tasks page displays all long running tasks that are currently in progress on the system. Tasks, such as volume synchronization, array initialization, and volume formatting, can take some time to complete. The Background Tasks page displays these tasks and their progress. After the task completes, the task is automatically deleted from the display. If a task fails with an error, select Monitoring → Events to determine the problem.

See 3.3.5, “Background Task” on page 93

The option that was known as System Details is integrated into the device overview on the general System pane, which is available after login or when you click **System** from the Monitoring menu. The details are shown in 3.3.2, “System details” on page 89.

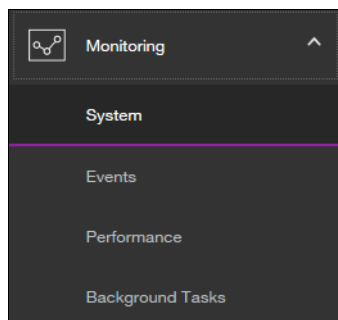


Figure 3-18 Accessing the Monitoring menu

In the following sections, we describe each option on the Monitoring menu.

3.3.1 System overview

The System option on the Monitoring menu provides a general overview about your Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, including the depiction of all devices in a rack directly connected to it. See Figure 3-19.



Figure 3-19 System overview

When you hover a mouse pointer over a specific component in an enclosure, a pop-up window indicates the details of disk drives in the unit. See Figure 3-20 for the details of Drive 0 in an enclosure.

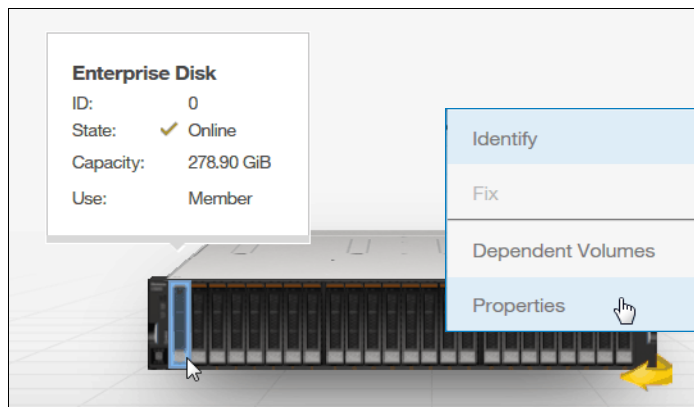


Figure 3-20 Component

By right-clicking and selecting **Properties**, you see detailed technical parameters, such as capacity, interface speed, rotation speed, and the drive status (online or offline). Click **View more details** as shown in Figure 3-21 on the properties frame.




Figure 3-21 Properties and View more details option

See Figure 3-22 for a detailed object properties view.

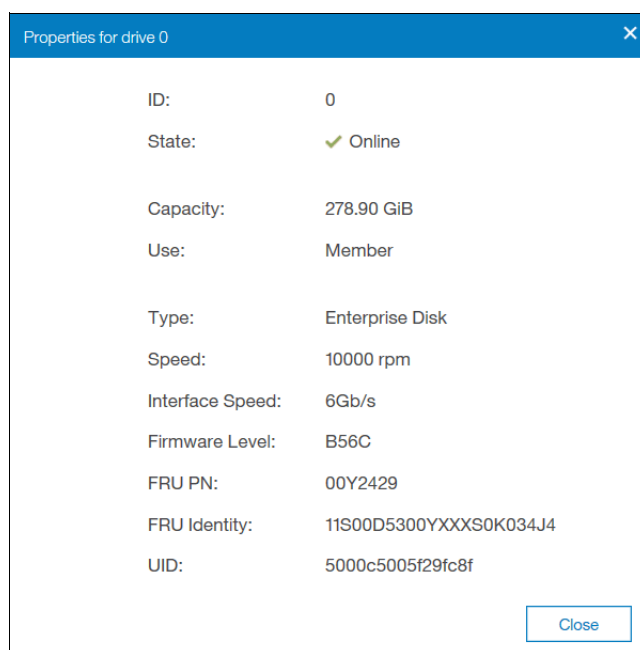


Figure 3-22 Component details

In an environment with multiple Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, you can easily navigate the onsite personnel or technician to the correct device by enabling the identification LED on the front pane. First, right-click the enclosure or drive that you want to identify. Then, click **Identify** in the pop-up window that is shown in Figure 3-23 and wait for the confirmation from the technician that the device in the data center was identified correctly.

After the confirmation, click **Turn LED Off** (Figure 3-23).

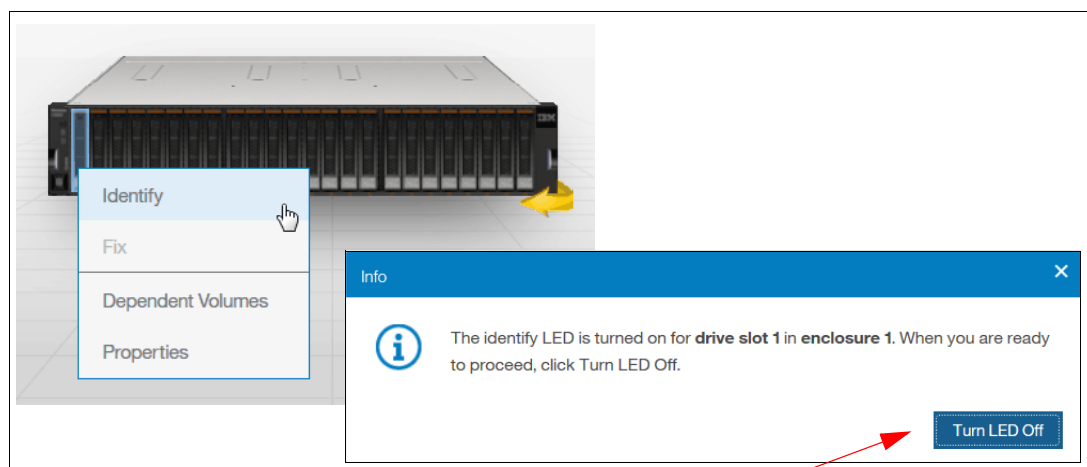


Figure 3-23 Using the identification LED

Alternatively, you can use the command-line interface (CLI) to obtain the same results. Type the following sequence of commands:

- ▶ **svctask chenclosure -identify yes 1** (or just **chenclosure -identify yes 1**)
- ▶ **svctask chenclosure -identify no 1** (or just **chenclosure -identify no 1**)

You can use the same CLI to obtain results for a specific controller or drive.

Each system that is shown in the dynamic system view in the middle of a System pane can be rotated by 180° to see its rear side. Click the rotation arrow in the lower-right corner of the device, as illustrated in Figure 3-24. In case of a malfunction, the arrow appears in red. Affected areas are marked in yellow. See arrow. Hover over it with your mouse to get more detailed information.

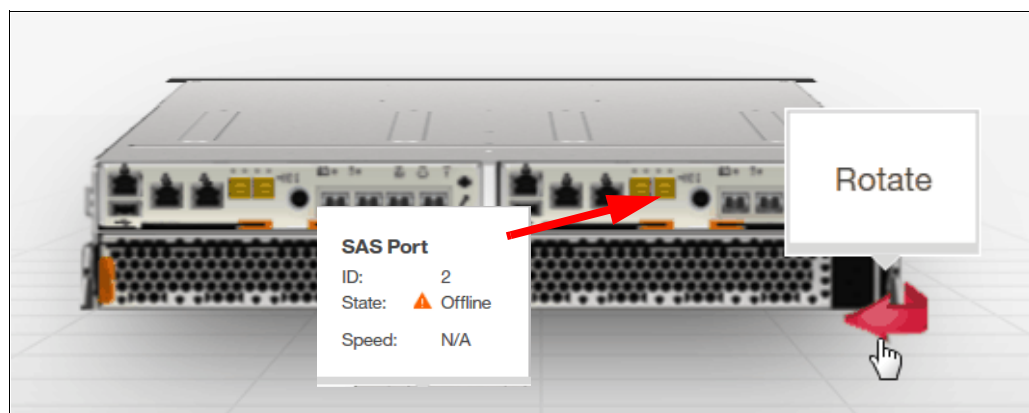


Figure 3-24 Rotating the enclosure

3.3.2 System details

The System Details option was removed from the Monitoring menu. However, its modified information is still available directly from the System pane. The Properties option provides the extended level of parameters and technical details that relate to the system, including the integration of each element with an overall system configuration. Right-click the enclosure that you want and navigate to the **Properties** option to obtain the detailed information (Figure 3-25).

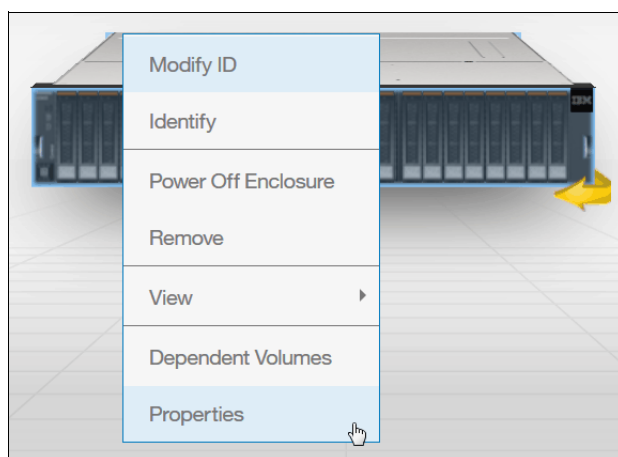


Figure 3-25 System details

The output is shown in Figure 3-26 on page 90. By using this menu, you can also power off the machine.

Warning: Powering the machine on remotely is not possible.

Remove the node or enclosure from the system, or list all volumes that are associated with the system (Show Dependent Volumes).

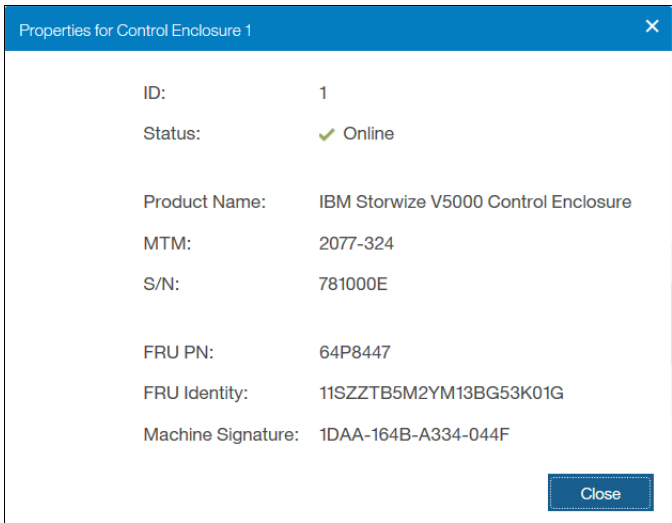


Figure 3-26 Enclosure technical details under the Properties option

Additionally, from the System pane, you can get an overview (View) of the hardware, available ports and status for Fibre Channel (FC) and serial-attached SCSI (SAS) ports and Drives (Figure 3-27).




Figure 3-27 Canister details and vital product data

By selecting, for example, **Fibre Channel Ports**, you can see the list and status of available FC ports with their speed and worldwide port names (WWPNs) in Figure 3-28.

Port ID	Owning Node	Virtualized	State	Speed	WWPN
1	node1	No	Active	8Gb	500507680D241780
1	node1	Yes	Active	8Gb	500507680D741780
1	node2	No	Active	8Gb	500507680D241781
1	node2	Yes	Active	8Gb	500507680D741781
2	node1	No	Inactive unconfigured	N/A	500507680D281780
2	node1	Yes	Inactive unconfigured	N/A	500507680D781780
2	node2	No	Inactive unconfigured	N/A	500507680D281781
2	node2	Yes	Inactive unconfigured	N/A	500507680D781781
3	node1	No	Inactive unconfigured	N/A	500507680D2C1780

Figure 3-28 Status of FC ports in the control enclosure

3.3.3 Events

The Events option, which is selected from the Monitoring menu (Figure 3-18 on page 86), tracks all informational, warning, and error messages that occur in the system. You can apply various filters to sort them or export them to an external CSV file. A CSV file can be created from the information that is included in the Events list. Figure 3-29 shows the display after you click **Events** from the Monitoring menu.

Error Code	Last Time Stamp	Status	Description	Object Type
1785	10/3/17 8:30:49 AM	Alert	A problem occurred with the Key Server	key_server
1785	10/3/17 8:30:44 AM	Alert	A problem occurred with the Key Server	key_server
1785	10/3/17 8:30:44 AM	Alert	A problem occurred with the Key Server	key_server
1785	10/3/17 8:00:42 AM	Alert	A problem occurred with the Key Server	key_server
1785	10/3/17 8:00:42 AM	Alert	A problem occurred with the Key Server	key_server
	10/3/17 1:54:27 AM	Alert	Connection to a configured remote cluster has been lost	cluster
	10/2/17 6:33:54 PM	Message	FlashCopy copied	fc_map
	10/2/17 6:33:39 PM	Message	Volume copy format completed	vdisk
	10/2/17 6:00:43 PM	Message	Array mdisk rebuild finish	mdisk

Figure 3-29 Events log

The procedures for how to work with the events log and how to run various fix procedures by using the Events option are described in Chapter 12, “RAS, monitoring, and troubleshooting” on page 625.

3.3.4 Performance

The Performance pane reports the general system statistics that relate to processor (CPU) use, host and internal interfaces, volumes, and MDisks. You can switch between MBps or IOPS, or even navigate to the statistics at the node level. The Performance pane might be useful when you compare the performance of each node in the system if problems exist after a node fail over occurs. See Figure 3-30.

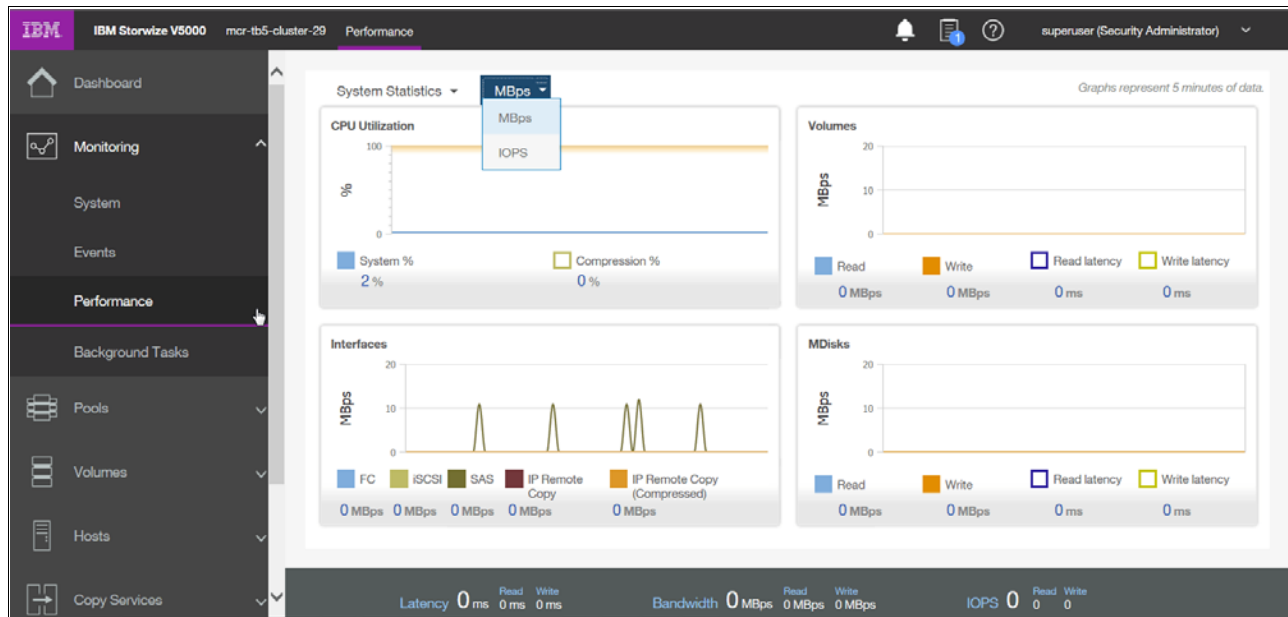


Figure 3-30 Performance statistics of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems

The performance statistics in the GUI shows, by default, the last 5 minutes of data. To see the details of each sample, click the graph and select the time stamp, as shown in Figure 3-31.

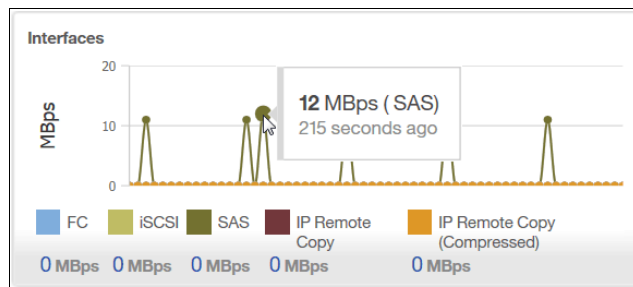


Figure 3-31 Sample details

As mentioned before, the previous charts represent 5 minutes of the data stream. For in-depth storage monitoring and performance statistics of your system with historical data, use the IBM SmartCloud Virtual Storage Center.

3.3.5 Background Task

Use the Background Tasks page to view and manage current tasks that are running on the system. The Background Tasks page displays all long running tasks that are currently in progress on the system. Tasks, such as volume synchronization, array initialization, and volume formatting, can take some time to complete. The Background Tasks page displays these tasks and their progress. After the task completes, the task is automatically deleted from the display. If a task fails with an error, select Monitoring → Events to determine the problem. Figure 3-32 shows an example of a Flash Copy Operation.

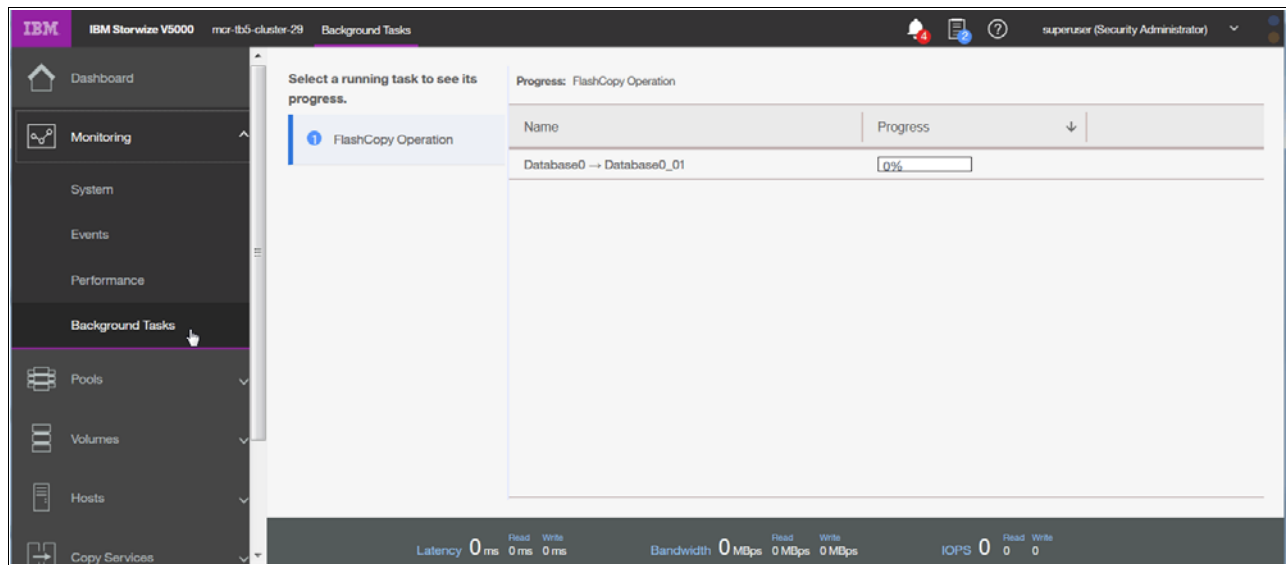


Figure 3-32 Background Task

3.4 Pools menu

Table 3-1 A *pool* or storage pool is a collection of MDisks that jointly contain all of the data for a specified set of volumes. All MDisks in a pool are split into *extents* of the same size. Volumes are created from the extents that are available in the pool. You can add MDisks to a storage pool at any time, either to increase the number of extents that are available for new volume copies or to expand existing volume copies. *Extents*

Extent size (MB)	Maximum non thin-provisioned volume capacity in GB	Maximum thin-provisioned volume capacity in GB	Maximum Compressed Volume Size	Maximum MDisk capacity in GB	Total storage capacity manageable per system*
16	2048 (2 TB)	2000		2048 (2 TB)	64 TB
32	4096 (4 TB)	4000		4096 (4 TB)	128 TB
64	8192 (8 TB)	8000		8192 (8 TB)	256 TB
128	16,384 (16 TB)	16,000		16,384 (16 TB)	512 TB
256	32,768 (32 TB)	32,000		32,768 (32 TB)	1 PB
512	65,536 (64 TB)	65,000		65,536 (64 TB)	2 PB
1024	131,072 (128 TB)	130,000	96 TB	131,072 (128 TB)	4 PB
2048	262,144 (256 TB)	260,000	96 TB	262,144 (256 TB)	8 PB

Extent size (MB)	Maximum non thin-provisioned volume capacity in GB	Maximum thin-provisioned volume capacity in GB	Maximum Compressed Volume Size	Maximum MDisk capacity in GB	Total storage capacity manageable per system*
4096	262,144 (256 TB)	262,144	96 TB	524,288 (512 TB)	16 PB
8192	262,144 (256 TB)	262,144	96 TB	1,048,576 (1024TB)	32 PB

* The total capacity values assumes that all of the storage pools in the system use the same extent size.

Volumes are created from the extents that are available in the pool. You can add MDisk to a storage pool at any time, either to increase the number of extents that are available for new volume copies or to expand existing volume copies.

Place the cursor over the Pools function icon and click to display the Pools menu options (Figure 3-33).

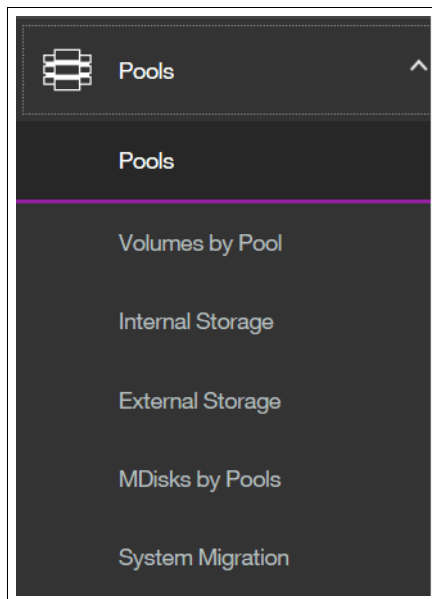


Figure 3-33 Navigate to the Pools menu

The Pools menu includes the following options:

- Pools** Shows a list of pools that are already available within the system. It provides an option to create or modify pools and *child pools*, adding additional storage.
- Volumes by Pool** Applies the high-level filter, listing all defined volumes per pool. It also provides a capacity overview, which is valid for a specific, selected pool only. This view is excellent when you plan a migration of a volume to another pool so that you have a common overview of all pools and their associated volumes. Unused volumes are not listed.
- Internal Storage** Provides an overview of all disk drives that are installed in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, including its enclosures. You can filter based on disk type and capacity and also see unused volumes that are not assigned to any pool.

- External Storage** Shows all pools and their volumes that are created from the systems that connect to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 externally and that are integrated in the system repository. It does not show any internal pools or volumes. This type of storage is also called *external virtualization*.
- MDisks by Pools** Provides the list of all managed disks (MDisks) that are either internally or externally connected and associated with one of the defined pools. It also lists all unassigned MDisks separately.
- System Migration** Offers the migration wizard to import data from image-mode MDisks to a specific pool. It is useful when you migrate data non-disruptively to the hosts from old external storage to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

3.4.1 Pools

If you plan to add storage to an existing pool, use the main Pools view. Right-click an existing pool (or create a pool in advance and add the storage) and select **Add Storage** as shown in Figure 3-34.

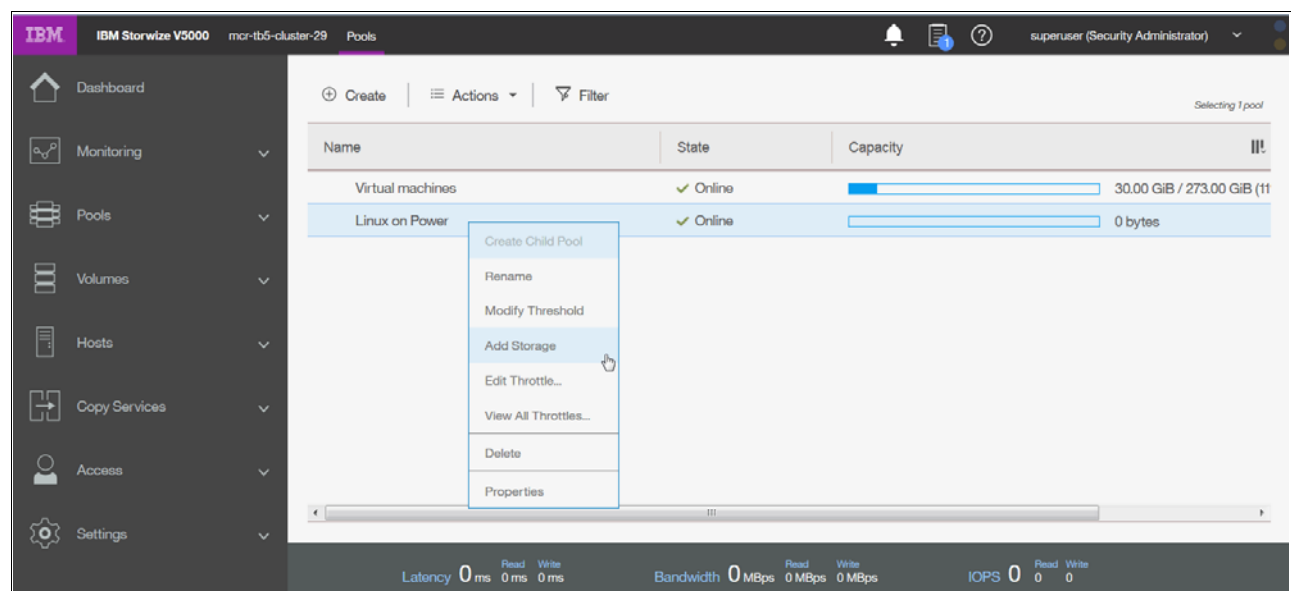


Figure 3-34 Add Storage option

In the next window, you can either select **Internal** or **Internal custom**. Internal custom gives you the chance to select the Raid type, Drive class and use of a spare drive. See Figure 3-35 on page 96.

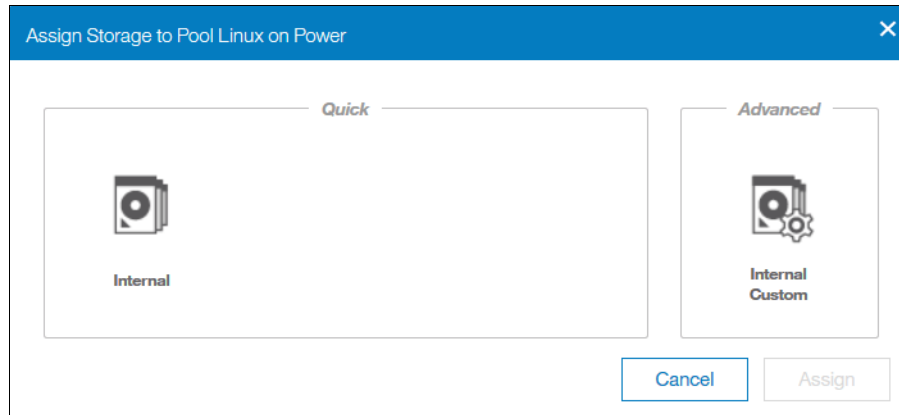


Figure 3-35 Internal storage selection

Figure 3-36 shows the window if you select Internal. Choose the default (Quick) settings to create a new MDisk for the pool.

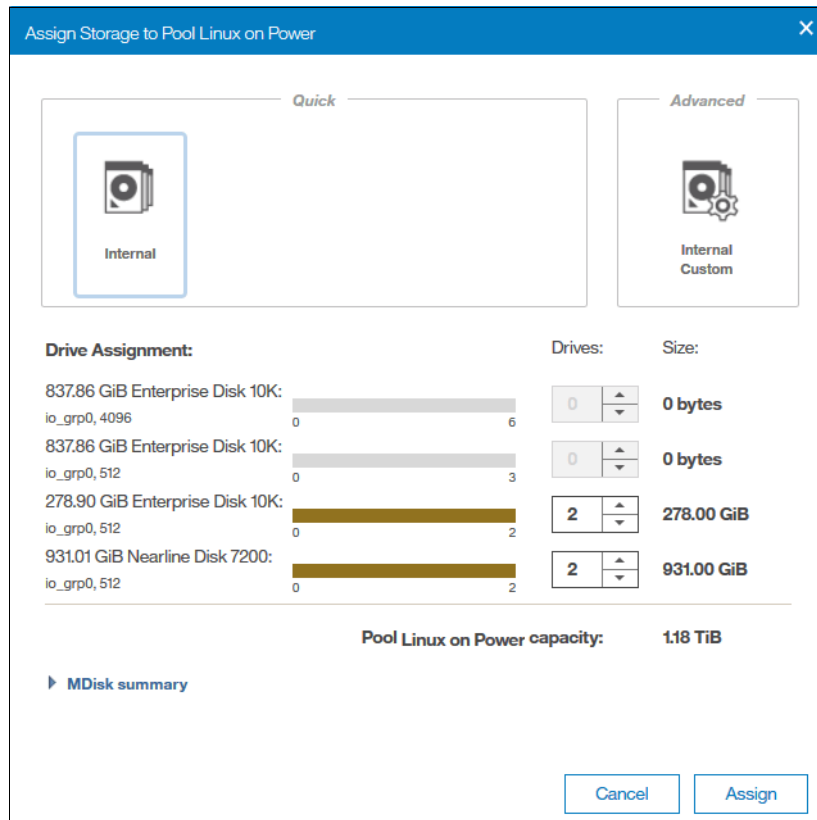


Figure 3-36 Internal storage selection

Figure 3-37 on page 97 shows you the window if you selected the Internal Custom storage panel.

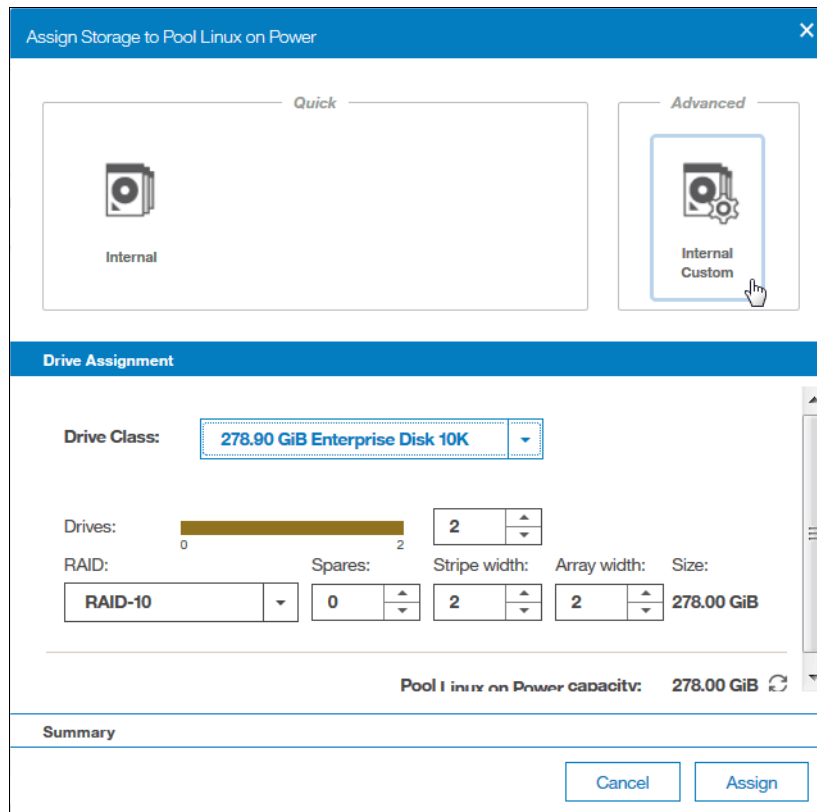


Figure 3-37 Internal Custom storage selection

You can choose from the available drive classes (depending on the installed drives) and RAID sets.

How to rename the pool is shown in Figure 3-38.

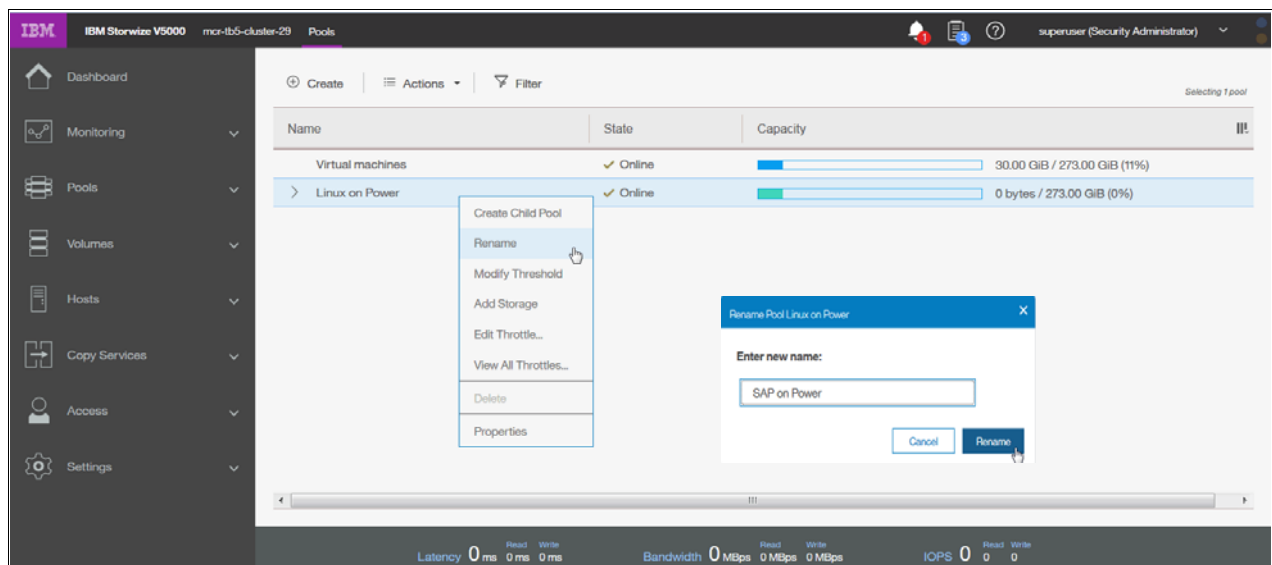


Figure 3-38 Renaming the pool

Throttles for pools

Throttling is a mechanism to control the amount of resources that are used when the system is processing I/Os on a specific pool. If a throttle is defined, the system either processes the I/O, or delays the processing of the I/O to free resources for more critical I/O. The system also supports throttles to delay processing of I/O operations for pools. If storage systems provide storage to a wide variety of applications, then production pools with more critical I/O can be competing with pools that have lower priority operations. For example, pools that are used for backup or archive operations can have I/O intensive workloads, potentially taking bandwidth from production pools. Pool throttles can be used to limit I/Os for these types of pools so that I/O operations for production pools are not affected. Figure 3-39 shows an example of pool throttling.

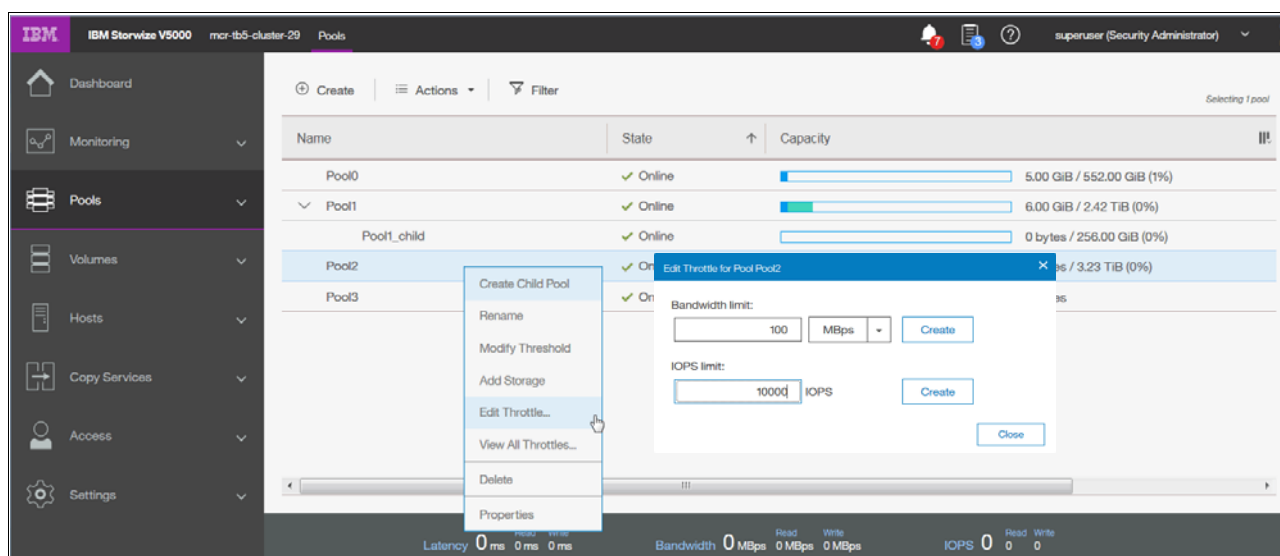


Figure 3-39 Pool throttling

Limits of the throttles are shown in Figure 3-40.

<p>Specify the maximum amount bandwidth. Valid values are 1 - 256TBps, the value with Unit GBps or TBps can be up to two fractional digits.</p>	<p>Specify the maximum amount of I/O operations per second (IOPS). Valid values are 1 - 33554432.</p>
Bandwidth limit	IOPS limit

Figure 3-40 Throttling limits

3.4.2 Child pools

Before V7.4, the disk space of a storage pool was provided from MDisks so the capacity of a storage pool depended on the MDisks' capacity. Creating or splitting a storage pool is impossible. A user cannot create a storage pool and specify the capacity that they want. A *child pool* is a new logical object that is created from a physical storage pool. A child pool provides most of the functions that pools offer (for example, volume creation), but the user can specify the capacity of the child pool at creation. Administrators can use child pools to control capacity allocation for volumes that are used for specific purposes.

A number of administration tasks benefit from being able to define and work with a part of a pool. For example, the system supports VMware vSphere Virtual Volumes, sometimes

referred to as VVols, that are used in VMware vCenter and VASA applications. Before a child pool can be used for Virtual Volumes for these applications, the system must be enabled for Virtual Volumes.

A child pool is an object that is similar to a storage pool, and a child pool can be used interchangeably with a storage pool. A child pool supports volume copy and migration. However, limitations and restrictions exist for child pools:

- ▶ The maximum capacity cannot exceed the parent pool's size.
- ▶ The capacity can be allocated at creation (thick) or flexible (thin).
- ▶ You must always specify the parent storage pool. The child pool does not own any MDisks.
- ▶ Child pools can also be created by using the GUI.
- ▶ The maximum number of child pools in one parent pool is 127.
- ▶ You are restricted to migrating image-mode volumes to a child pool.
- ▶ Volume extents cannot be migrated out of the child pool.
- ▶ You cannot shrink capacity smaller than the real capacity.

You can view the list of child pools from the Pools menu option by clicking the plus sign (+) of a parent pool, as shown in Figure 3-41.

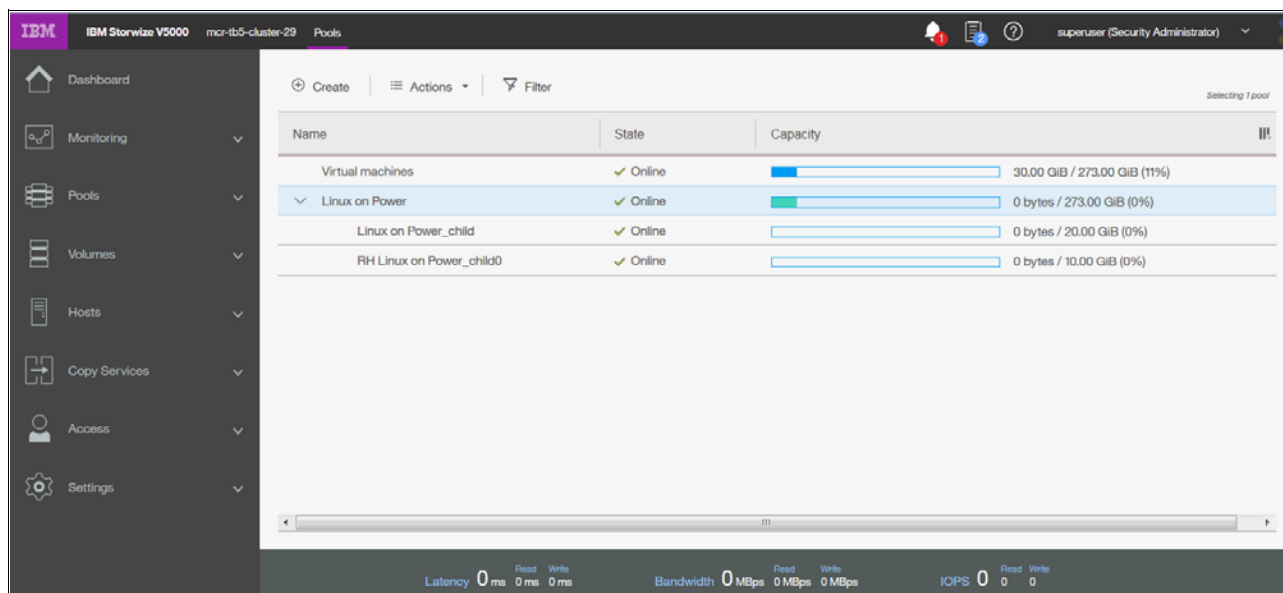


Figure 3-41 Working with child pools

3.4.3 Volumes by pool

The Volumes by Pool menu option lists all defined volumes, which are sorted by their pool assignment (Figure 3-42 on page 100). Unassigned volumes are not visible in this window. By using this window, you can, for example, create volumes, map or unmap volumes to and from hosts, migrate volumes to another pool, and rename, shrink, or expand volumes.

In addition, you can choose a different icon (Figure 3-42 on page 100) that represents this pool.

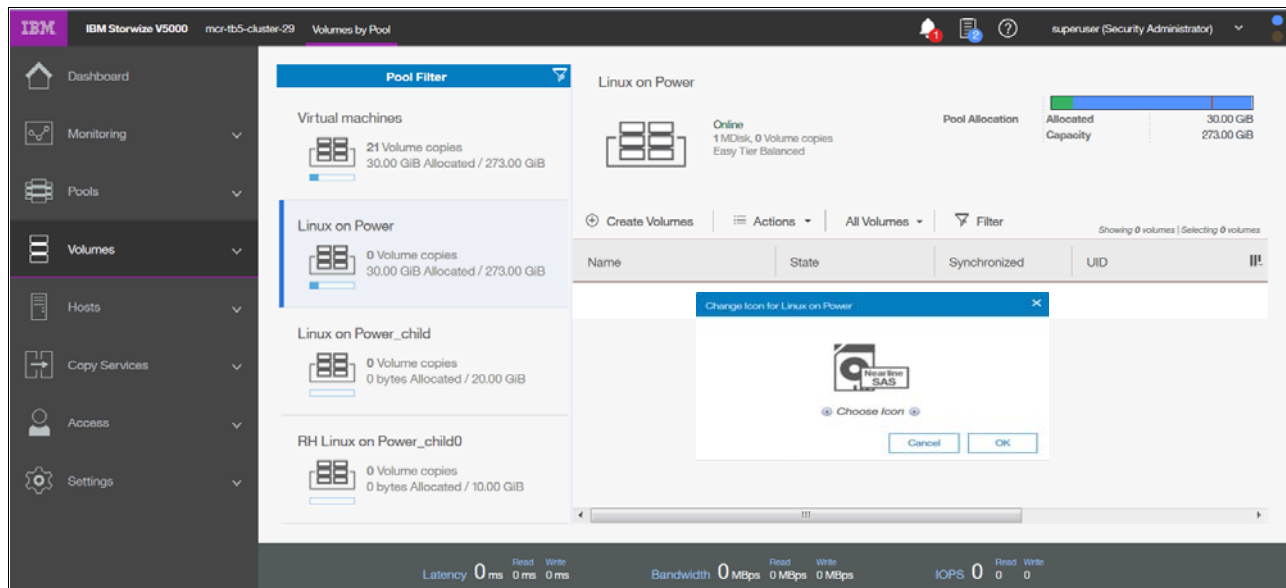


Figure 3-42 Volume by Pool option and changing the icon

To change the icon, use the pen sign as shown in Figure 3-43.




Figure 3-43 Pen sign to change icon

When the pools are defined and the volumes are assigned, the pool shows one of the following operational states:

- Online** The storage pool is online and available. All of the MDisks in the storage pool are available.
- Degraded path** One or more nodes in the clustered system cannot access all of the MDisks in the pool. A degraded path state is most likely the result of the incorrect configuration of either the storage system or the FC fabric. However, hardware failures in the storage system, FC fabric, or node can also be a contributing factor to this state.
- Degraded ports** One or more 1220 errors were logged against the MDisks in the storage pool. The 1220 error indicates that the remote FC port was excluded from the MDisk. This error might cause reduced performance on the storage system and usually indicates a hardware problem with the storage system.

To fix this problem, you must resolve any hardware problems on the storage system and fix the 1220 errors in the event log. To resolve these errors in the log, select **Troubleshooting** → **Recommended Actions** in the management GUI.

This action displays a list of unfixed errors that are in the event log. For these unfixed errors, select the error name to begin a guided maintenance procedure to resolve the errors. Errors are listed in

descending order with the highest priority error listed first. Resolve the highest priority errors first.

Offline

The storage pool is offline and unavailable. No nodes in the system can access the MDisks. The most likely cause is that one or more MDisks are offline or excluded.

Important: In this view, volumes from child pools are shown the same way as volumes from standard pools. The relationships between the child and parent pools are not visible.

3.4.4 Internal storage

Click the **Internal Storage** option in the Pools menu to open a window similar to Figure 3-44. From this window, you can allocate Redundant Array of Independent Disks (RAID) arrays of internal disk drives into storage pools. This window also offers the option to display internal drives, based on their capacity and speed.

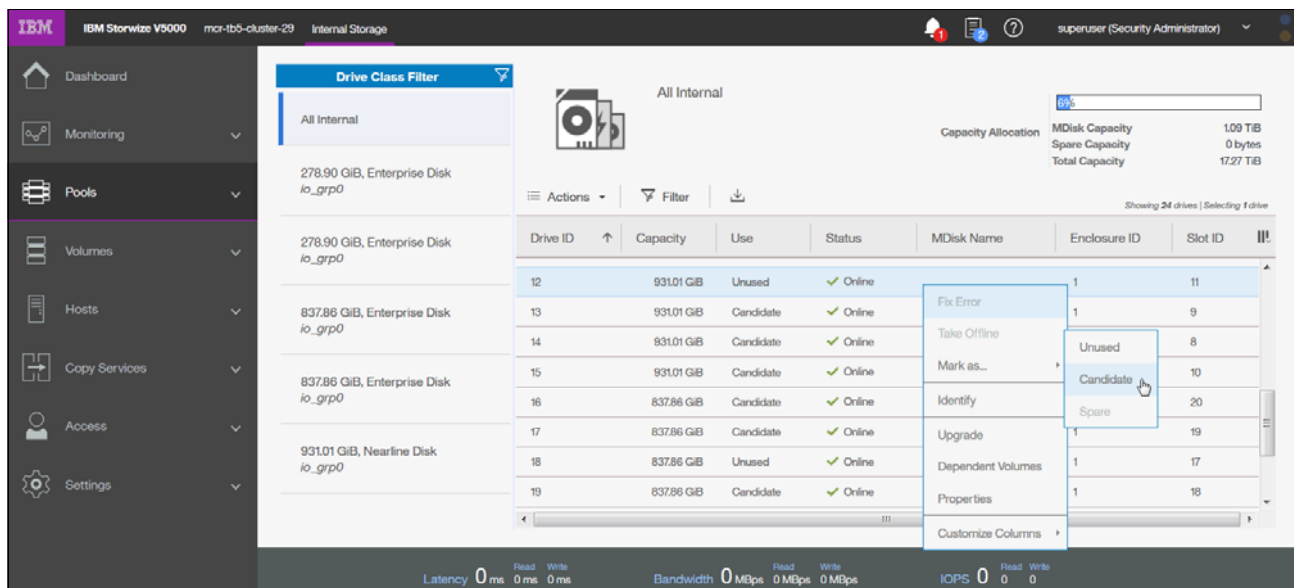


Figure 3-44 Internal storage window

Click **Actions** in the table header, or right-click a specific drive to take the drive offline, show its properties, update firmware of a single drive (or all under Actions) or mark the drive as Unused, Candidate, or Spare.

3.4.5 External storage

Before you can add external storage your system has to be in the **Replication Layer**.

This can be either done via CLI:

► **chsystem layer replication -cacheprefetch yes/no**

If you need to switch it back to storage layer use following command:

► **chsystem layer storage -cacheprefetch yes/no**

If you search for external storage and you are not in the replication layer you will get the following warning in the GUI as shown in Figure 3-45 on page 102.

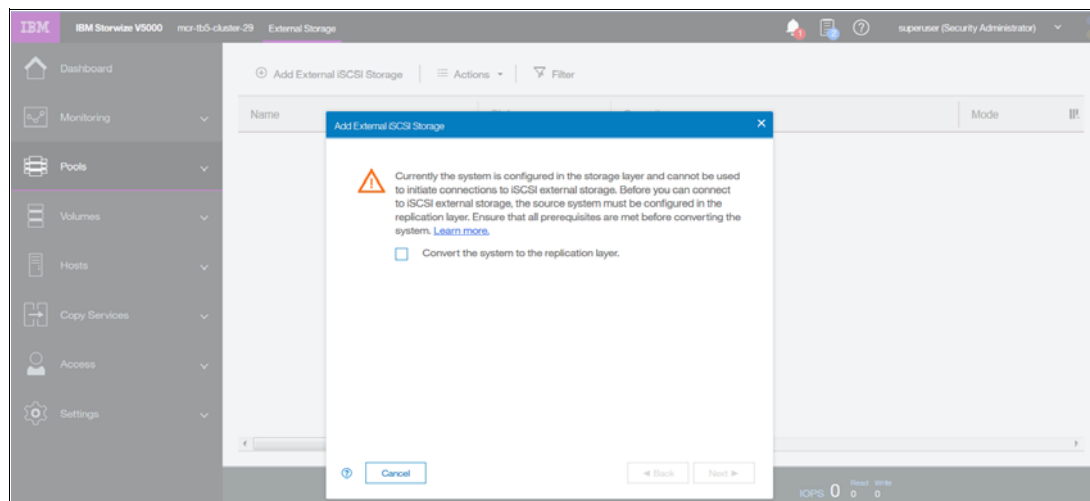


Figure 3-45 Convert System to replication layer

Clicking the **External Storage** option opens the window as shown in Figure 3-46. It provides the list of all externally connected (storage area network (SAN)- and iSCSI attached) disk systems to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. The system supports also iSCSI connections to systems that are used as external storage. Unlike Fibre Channel connections, you need to manually configure the connections between the source system and these target external storage systems. Direct attachment between the system and external storage systems is not supported and requires Ethernet switches between the system and the external storage. To avoid a single point of failure, a dual switch configuration is recommended. For full redundancy, a minimum of two paths between each initiator node and target node must be configured with each path on a separate switch. In addition, extra paths can be configured to increase throughput if both initiator and target nodes support more ports. The system supports a maximum of four paths per node.

When the new external storage system is zoned correctly to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, run the **Discover storage** procedure either from the Actions menu in the table header or by right-clicking any of the existing MDisk in the list (Figure 3-46).

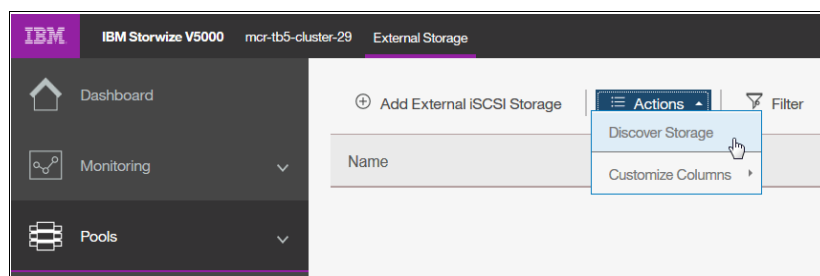


Figure 3-46 Detecting external storage systems

A new storage controller (external storage system) is listed automatically when the SAN zoning is configured, but typically without detecting disk drives (Figure 3-47 on page 103).




Figure 3-47 Figure 3-47Automatically detected new external storage system

By right-clicking a newly detected storage system, you can rename the controller's default name, in our case, the IBM System Storage FlashSystem 840, to reflect the real type of the storage device. We suggest that you use a simple naming convention, which in our case is FlashSystem 840 (Figure 3-49).

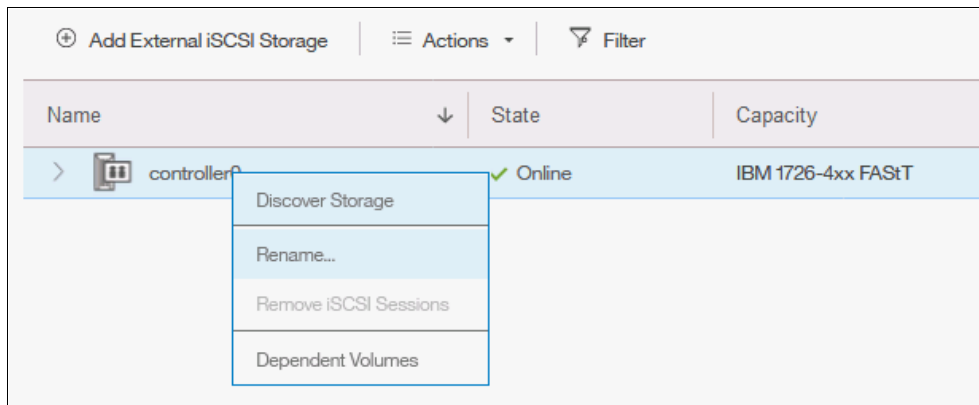


Figure 3-48 Select Rename of the detected system

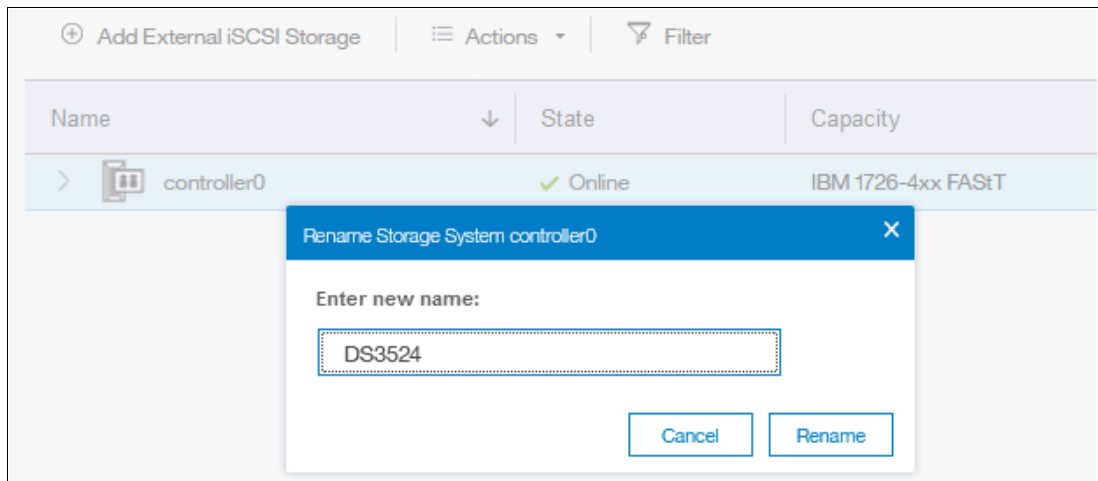


Figure 3-49 Renaming the detected Controller

After the new external storage system is named correctly, detect all disks that are configured on that external storage, in our case, System Storage FlashSystem 840. You can also discover storage from the CLI by using the **svctask detectmdisk** command or **detectmdisk**. Figure 3-50 on page 104 shows details about detected managed disks.


⊕ Add External iSCSI Storage		☰ Actions ▾		🔍 Filter	
Name	↓	State	Capacity	Mode	Storage System
⌵  DS3524		✓ Online	IBM 1726-4xx FASTT		
mdisk0		✓ Online	64.00 GiB	Unmanaged	DS3524
mdisk1		✓ Online	64.00 GiB	Unmanaged	DS3524
mdisk2		✓ Online	64.00 GiB	Unmanaged	DS3524
mdisk3		✓ Online	64.00 GiB	Unmanaged	DS3524
mdisk4		✓ Online	10.00 GiB	Unmanaged	DS3524
mdisk5		✓ Online	32.00 GiB	Unmanaged	DS3524
mdisk6		✓ Online	32.00 GiB	Unmanaged	DS3524
mdisk7		✓ Online	32.00 GiB	Unmanaged	DS3524
mdisk8		✓ Online	32.00 GiB	Unmanaged	DS3524

Figure 3-50 Newly discovered managed disks

All newly discovered disks are always interpreted in an *unmanaged* mode. You must assign them to the specific pool to be able to operate them.

Important: The MDisks are not physical disk drives, but storage arrays that are configured on external systems.

If you add a managed disk that contains existing data to a managed disk group, you lose the data that it contains. The *image* mode is the only mode that preserves its data.

3.4.6 MDisks by pools

This option on the Pools menu provides the list of all managed disks and arrays of disks, either internally or externally connected, and associated with one of the defined pools. It also lists all unassigned MDisks (which are only provided by external storage systems) separately, see Figure 3-51 on page 105.

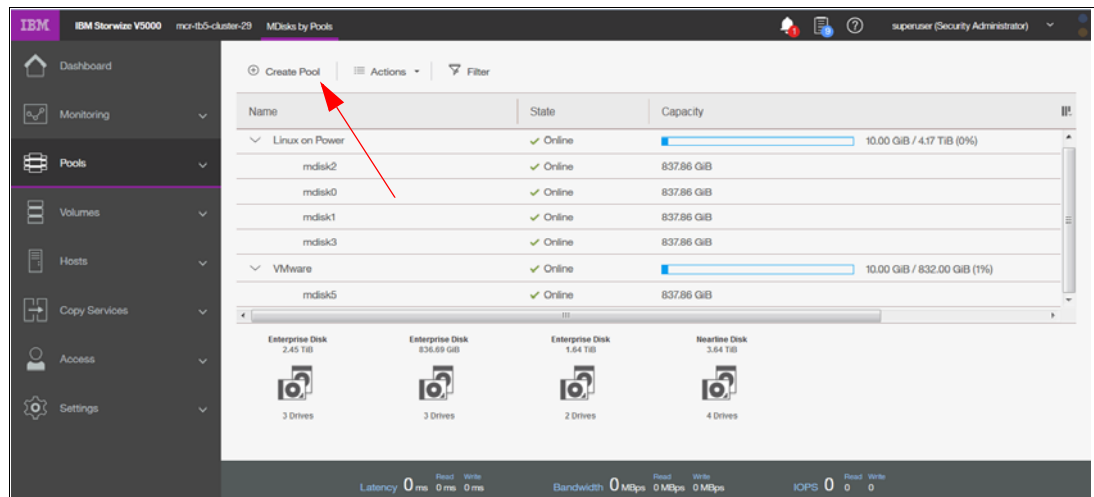


Figure 3-51 List of managed disks that are sorted within pools

All disks that are not yet assigned to any pool are listed in the Unassigned MDisks section. This section is always at the top of the list, even if you sort the list by pool name (clicking the **Name** header of the table). Right-click a specific disk to open a window where you can assign selected unmanaged disks to the pool.

From the same pane, you can define a new storage pool by clicking **Create Pool** in the upper-left corner of the table (highlighted in Figure 3-51). The wizard window opens and you need to specify pool parameters, such as Pool Name, Extent Size, and Warning Threshold. You can directly select Unmanaged MDisks that you want to include in the pool, or skip this task and add MDisks later.

Note: All sort functions in the header of the table apply to MDisks *within* pools. You cannot sort volumes based on specific criteria *across* all pools.

3.4.7 System migration

Migrating data from older storage systems to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage system enables applications to benefit from the new features, such as IBM Easy Tier, Space Efficient volumes, an intuitive management GUI, and advanced storage replication functions that better support applications.

To migrate existing data, use the storage migration wizard to guide you through the procedure. This wizard is available by selecting **Pools** → **System Migration** as shown in Figure 3-52 on page 106.

The migration of external volumes to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems is one of the key benefits and features of external storage virtualization that are provided by this product. Therefore, we dedicate a whole chapter to this topic. See Chapter 7, “Storage migration” on page 323 for detailed steps of the migration process.

Administrators can migrate data from the external storage system to the system that uses either iSCSI connections, serial-attached SCSI connections, and Fibre Channel or Fibre Channel over Ethernet connections. To use Fibre Channel connections, the system must have the optional Fibre Channel host interface adapter installed.

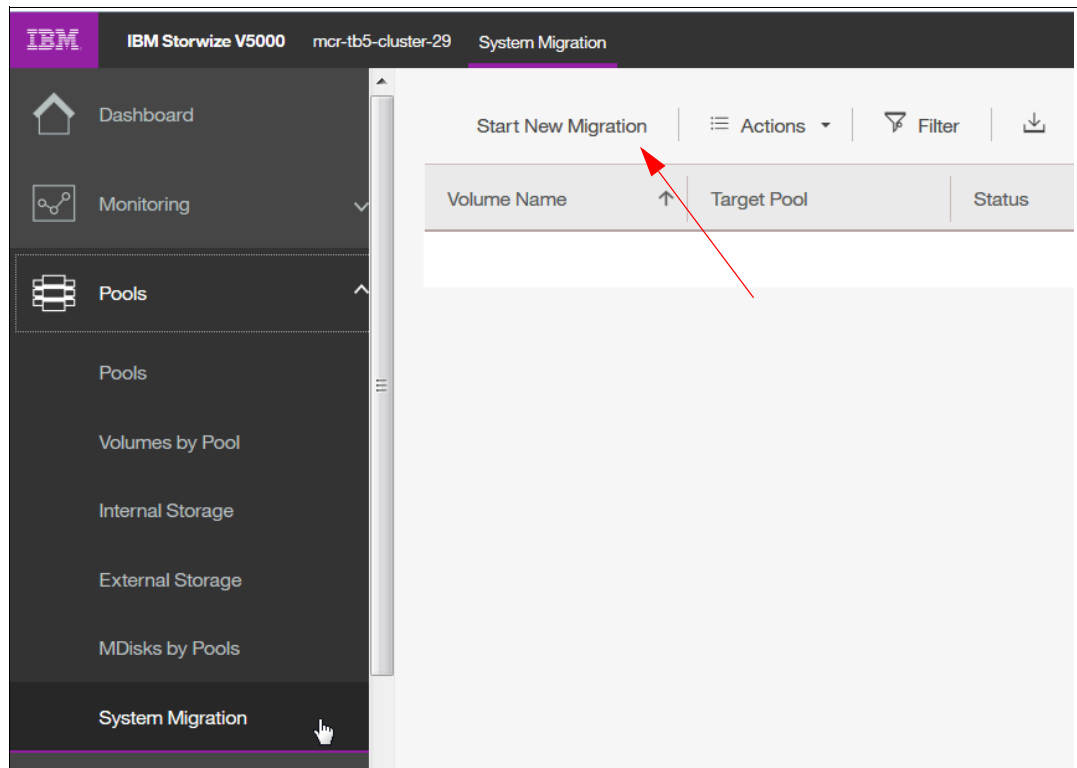


Figure 3-52 System migration

Note: Before migrating storage, ensure that all host operations are stopped, all the appropriate changes are made to the environment based on the connection type, and the storage that is being migrated is configured to use the device.

At any time, you can pause the running migration processes or create a new one. No license for External Virtualization is required to migrate from old storage to your new Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

3.5 Volumes menu

A *volume* is a logical disk that the system presents to the attached host. Application servers access volumes, not MDisk or drives. Volumes have additional characteristics. They can be automatically expanded, mirrored, or pre-allocated. The following list provides a description of the volume characteristics and differences:

- Basic** A basic (*fully allocated*) volume is the traditional data store method when any host input/output (I/O) is destaged to the drives. Even zeros are destaged. All of the zeros that exist are written.
- Mirrored** By using volume mirroring, a volume can have two physical copies. Each volume copy can belong to a different pool, and each copy has the same virtual capacity as the volume. In the management GUI, an asterisk (*) indicates the primary copy of the mirrored volume. The primary copy indicates the preferred volume for read requests.
- HyperSwap** HyperSwap volumes create copies on separate sites for systems that are configured with HyperSwap topology. Data that is written to a

HyperSwap volume is automatically sent to both copies so that either site can provide access to the volume if the other site becomes unavailable. HyperSwap volumes are supported on Lenovo storage systems (for example, Lenovo storage V5030 system) that contain more than one I/O group.

Custom	Custom volumes create volumes that are based on user-defined customization rather than taking the standard default settings for each of the options under quick volume creation.
Thin-provisioned	<p>When you create a volume, you can designate it as thin-provisioned. A thin-provisioned volume has a virtual capacity and a real capacity. <i>Virtual capacity</i> is the volume storage capacity that is available to a host. <i>Real capacity</i> is the storage capacity that is allocated to a volume copy from a storage pool.</p> <p>In a fully allocated volume, the virtual capacity and real capacity are the same. In a thin-provisioned volume, the virtual capacity can be much larger than the real capacity.</p>
Compressed	This volume is a special type of volume where data is compressed and thin-provisioned at the same time. Any compressed volume is a thin-provisioned volume by default, and no option is available to change this characteristic. Data within the compressed volume is compressed as it is written to disk. This design saves additional space on the storage drive so that you can store more data within the same storage system.
Change volumes	Change volumes are used in Global Mirror relationships where cycling mode is set to Multiple. Change volumes can also be used between HyperSwap volume copies, and other relationship types, to automatically maintain a consistent image of a secondary volume when a relationship is being re synchronized. Change volumes create periodic point-in-time-copies of the source volumes and replicate them to the secondary site. Using change volumes lowers bandwidth requirements by only addressing the average throughput and not the peak.

Important: Compression is only available in the Lenovo Storage V5030. It requires 64 GB of RAM. To use the compression function, you must obtain the Real-time Compression license.

To keep a volume accessible even when an MDisk on which it depends is unavailable, a mirrored copy can be added to a selected volume. Any volume (generic, thin-provisioned or compressed) can be mirrored with a mirror from any type, even the same one. Therefore, a volume can be either thin-provisioned with compressed copy or compressed with compressed copy. Each volume can have a maximum of two copies.

Each volume copy is created from a set of extents in a storage pool. By using volume mirroring, a volume can have two physical copies. Each volume copy can belong to a different storage pool, and each copy has the same virtual capacity as the volume. In the management GUI, an asterisk (*) indicates the primary copy of the mirrored volume. The primary copy indicates the preferred volume for read requests.

Select the **Volumes** function icon to display the Volumes menu options (Figure 3-53 on page 108).

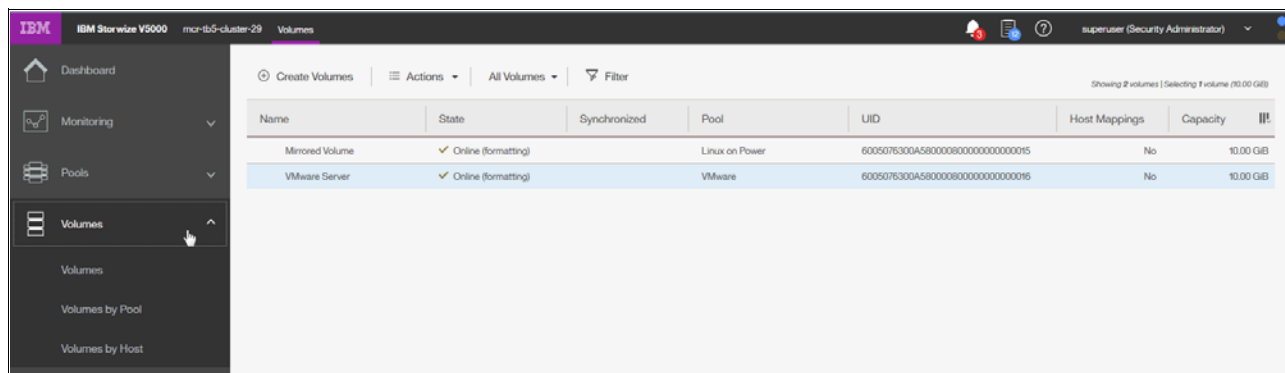


Figure 3-53 Volumes menu

3.5.1 All volumes

Select **Volumes** as shown in Figure 3-53. A list of all defined volumes, alphabetically sorted by the volume name (by default), is displayed. At any time, you can change the sort options by clicking a specific header in the table. You can directly configure a new volume by clicking **Create Volumes** (as shown by the arrow in Figure 3-54).

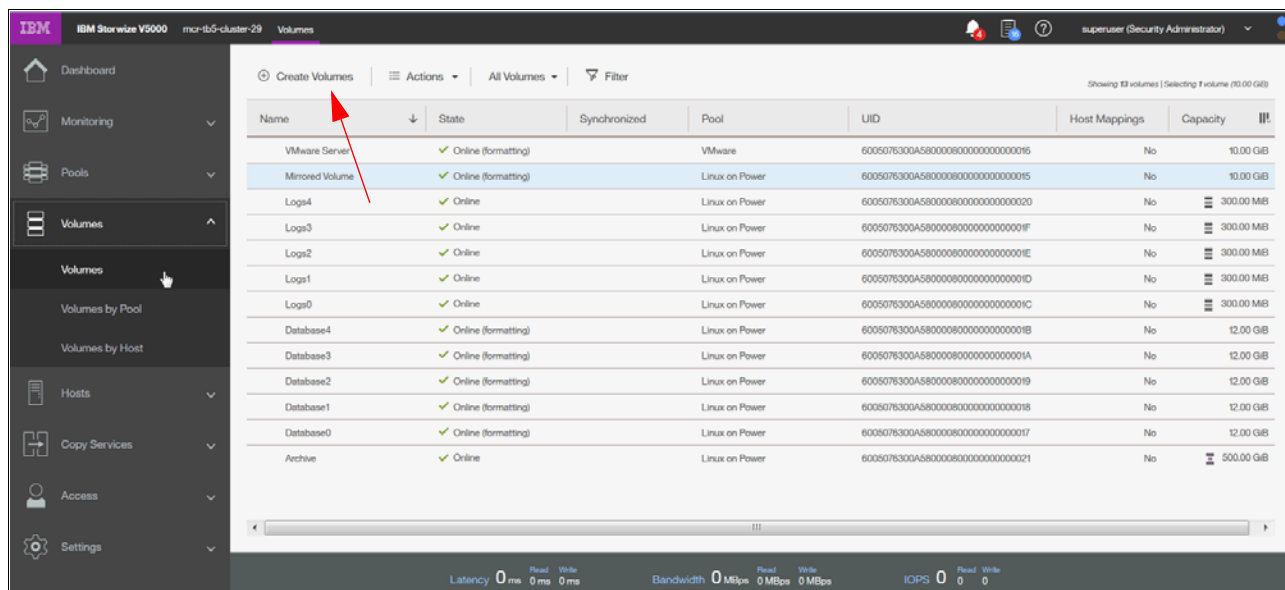


Figure 3-54 Create a volume

The wizard opens and the list of volume options is displayed (Figure 3-55 on page 109).

Figure 3-55 Create Volumes wizard

The description of each type of volume and the procedures for how to effectively create these volumes are described in Chapter 6, “Volume configuration” on page 269.

In addition to the volume creation, other direct volume functions are available:

- ▶ Mapping and unmapping volumes to hosts
- ▶ Renaming, shrinking, or expanding existing volumes
- ▶ Modify Mirror Synchronisation Rate
- ▶ Space savings → Estimate compression savings
- ▶ Migrating to a different pool
- ▶ Defining a volume copy

All of these tasks are available when you select a specific volume and click **Actions** (Figure 3-56 on page 110). Not all options are shown.

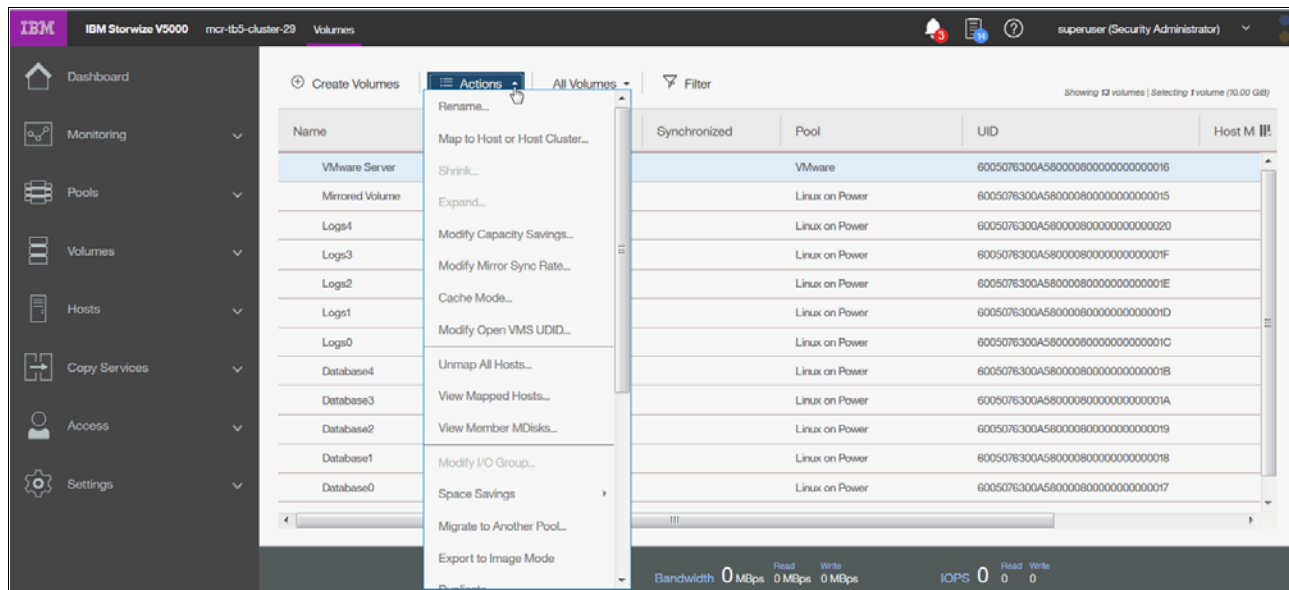


Figure 3-56 Actions menu for volumes

When you move a volume to another I/O group (a different Lenovo Storage V5030 system), be sure that the correct host zoning is in place. The target host must have access to both systems: source and target. This function is only available on the Lenovo Storage V5030.

3.5.2 Volumes by pool

This menu is identical to the one that is described in 3.4.3, “Volumes by pool” on page 99. See that section for details.

3.5.3 Volumes by host

Click **Volumes by Host** to open the window that is shown in Figure 3-57 on page 111. This window shows the volumes that are mapped to a certain host. You can perform the same actions with volumes as in all previous views, either by clicking **Actions** or by using the menu that opens after you right-click a specific volume. See Figure 3-57 on page 111 for details.

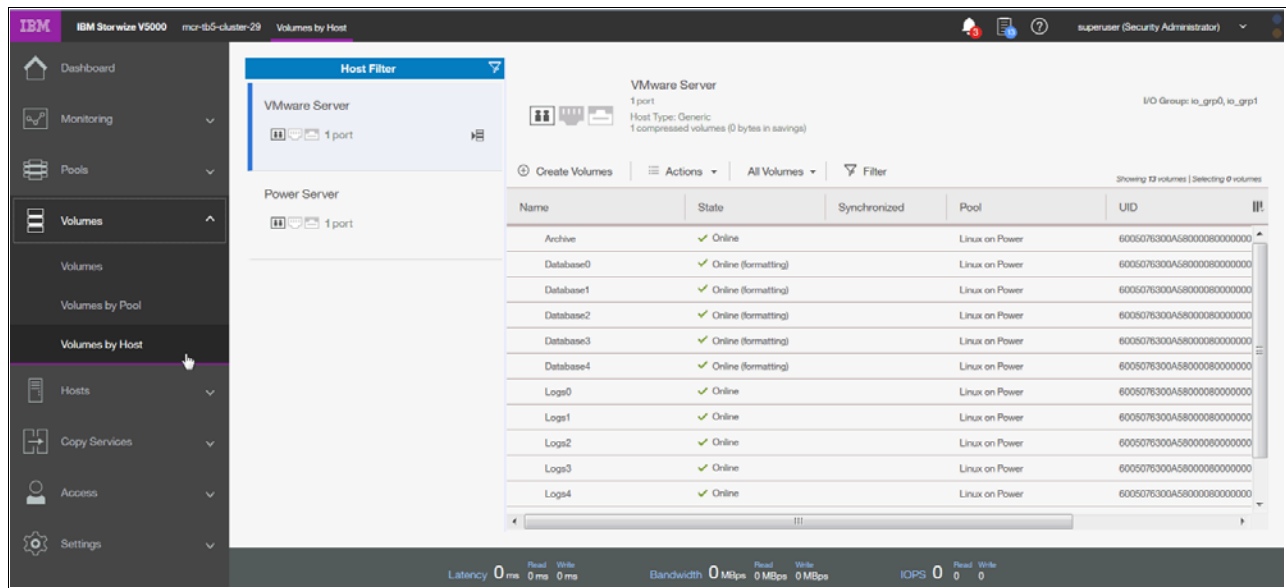


Figure 3-57 Listing volumes by host

3.6 Hosts menu

In a SAN environment, a host system is a computer that is connected to the system through one of the following: a Fibre Channel interface, serial-attached SCSI (SAS) connections, or an IP network. To use Fibre Channel or Fibre Channel over Ethernet connections to a storage area network (SAN), an optional host interface adapter must be installed. You can use several tools to manage hosts, including the management GUI, the CLI, and specialized utilities for working with host bus adapters (HBAs). To work with hosts in the management GUI, select **Hosts**. When you click the cursor on the Host function icon, the Hosts menu opens, providing the following options (Figure 3-58 on page 112):

- ▶ Hosts
- ▶ Host Clusters
- ▶ Ports by Host
- ▶ Host Mappings
- ▶ Volumes by Host

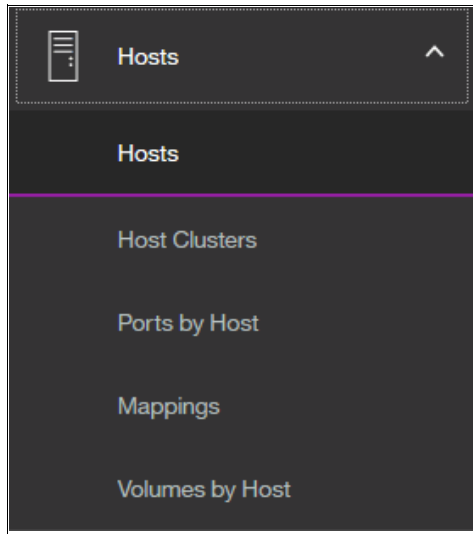


Figure 3-58 Hosts menu

3.6.1 Hosts

This option provides an overview about all hosts that are connected (zoned) to the system, detected, and configured to be ready for storage allocation. This overview shows the following information about the hosts:

- ▶ The name of the host as defined in the management GUI
- ▶ The type of the host
- ▶ Its access status
- ▶ The number of ports that is used for host mapping
- ▶ Whether host mapping is active or not

From the same pane, you can create a new host, rename a host, delete a host, or modify a host mapping. The output of the menu selection is shown in Figure 3-59.

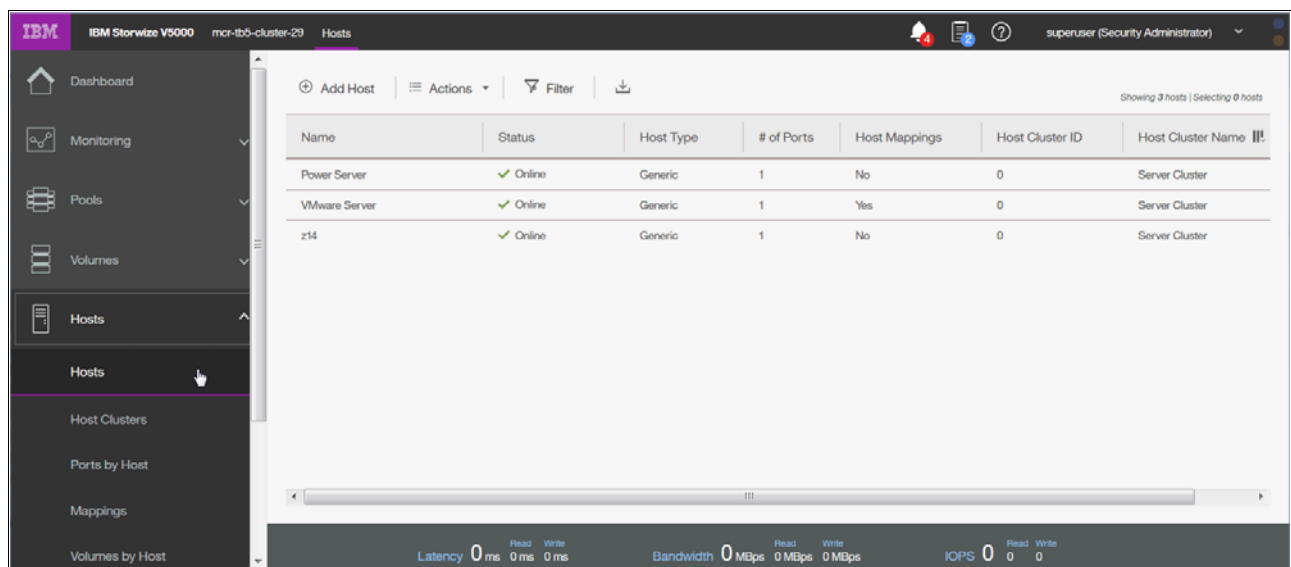


Figure 3-59 Overview of configured hosts

For example, when you click **Add Host** in a pane header, a wizard opens where you define either a Fibre Channel host or an iSCSI host (Figure 3-60).

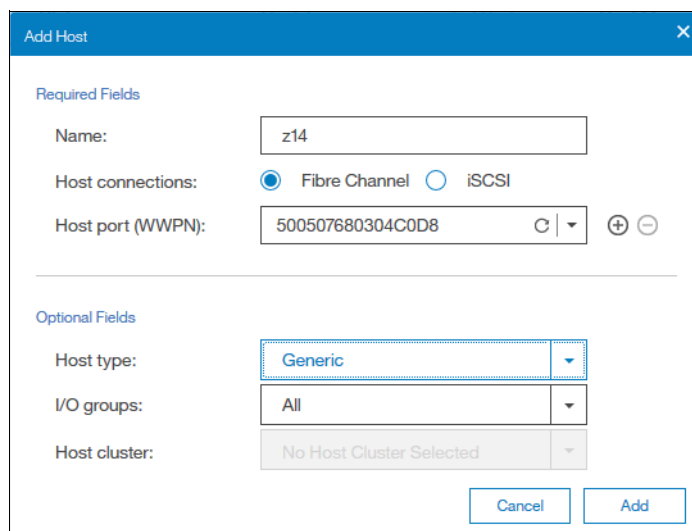
The 'Add Host' dialog box is divided into two sections: 'Required Fields' and 'Optional Fields'. In the 'Required Fields' section, the 'Name' field contains 'z14'. The 'Host connections' section has two radio buttons: 'Fibre Channel' (selected) and 'iSCSI'. The 'Host port (WWPN)' field contains '500507680304C0D8' and includes a copy icon and expand/collapse buttons. The 'Optional Fields' section contains three dropdown menus: 'Host type' set to 'Generic', 'I/O groups' set to 'All', and 'Host cluster' set to 'No Host Cluster Selected'. At the bottom right are 'Cancel' and 'Add' buttons.

Figure 3-60 Add Host wizard

To rename multiple hosts in a single step, mark all hosts that you want by using the Ctrl or Shift key, right-click, and then from the opened menu, select **Rename**. The window that is shown in Figure 3-61 opens.

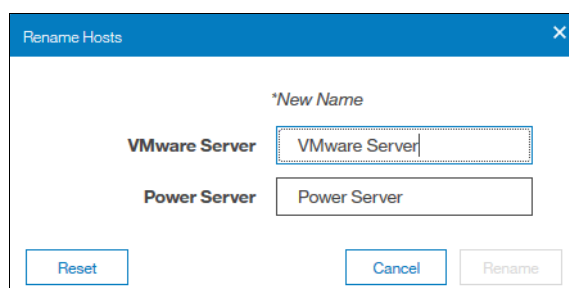
The 'Rename Hosts' dialog box shows a list of hosts with their current names and input fields for new names. The first host is 'VMware Server' with a new name of 'VMware Server'. The second host is 'Power Server' with a new name of 'Power Server'. At the bottom are 'Reset', 'Cancel', and 'Rename' buttons.

Figure 3-61 Renaming multiple hosts

Many of the actions that are described are available from different menus. For example, you can select **Volumes** and its option **Volumes by Hosts**, where you can also rename hosts. This flexibility is one of the advantages of the enhanced, redesigned management GUI.

3.6.2 Host Clusters

A host cluster is a group of logical host objects that can be managed together. For example, you can create a volume mapping that is shared by every host in the host cluster. See Figure 3-62 on page 114.

Create Host Cluster

Name:

Server Cluster

Optional: Select hosts to assign to a new host cluster. Any current volume mappings become the shared mappings for all the hosts in the host cluster.

i It is recommended that all hosts in a host cluster have access to the same I/O Groups.

Filter

Showing 3 hosts | Selecting 3 hosts

Name	Status	Host Type	Host Mappings
Power Server	✓ Online	Generic	No
VMware Server	✓ Online	Generic	Yes
z14	✓ Online	Generic	No

Need Help

Cancel

Back

Next

Figure 3-62 Create Host Cluster

The systems use internal protocols to manage access to the volumes and ensure consistency of the data. Host objects that represent hosts can be grouped in a host cluster and share access to volumes. New volumes can also be mapped to a host cluster, which simultaneously maps that volume to all hosts that are defined in the host cluster. Each host cluster is identified by a unique name and ID, the names of the individual host objects within the cluster, and the status of the cluster. A host cluster can contain up to 128 hosts. However, a host can be a member of only one host cluster. The management GUI displays the status of each host cluster. A host cluster can have one of the following statuses:

- **Online** All hosts in the host cluster are online.
- **Host degraded** All hosts in the host cluster are either online or degraded.
- **Host cluster degraded** At least one host is offline and at least one host is either online or degraded.
- **Offline** All hosts in the host cluster are offline (or the host cluster does not contain any hosts).

By default, hosts within a host cluster inherit all shared volume mappings from that host cluster, as if those volumes were mapped to each host in the host cluster individually. Hosts in a host cluster can also have their own private volume mappings that are not shared with other hosts in the host cluster. With shared mapping, volumes are mapped on a host cluster basis. The volumes are shared by all of the hosts in the host cluster, if there are no Small Computer System Interface (SCSI) LUN conflicts among the hosts. Volumes that contain data that is needed by other hosts are examples of a shared mapping. If a SCSI LUN conflict occurs, a shared mapping is not created. SCSI LUN conflicts can occur if multiple volumes are mapped with the same SCSI LUN ID or if same volume is mapped to multiple SCSI LUN IDs. The system does not allow a volume to be mapped more than once to the same host. With private mapping, individual volumes are directly mapped to individual hosts. These volumes are not

shared with any other hosts in the host cluster. A host can maintain the private mapping of some volumes and share other volumes with hosts in the host cluster. The SAN boot volume for a host would typically be a private mapping.

3.6.3 Ports by host

Click **Ports by Hosts** to open the pane that is shown in Figure 3-63. This pane lists the Fibre Channel and iSCSI ports that are assigned to a particular host.

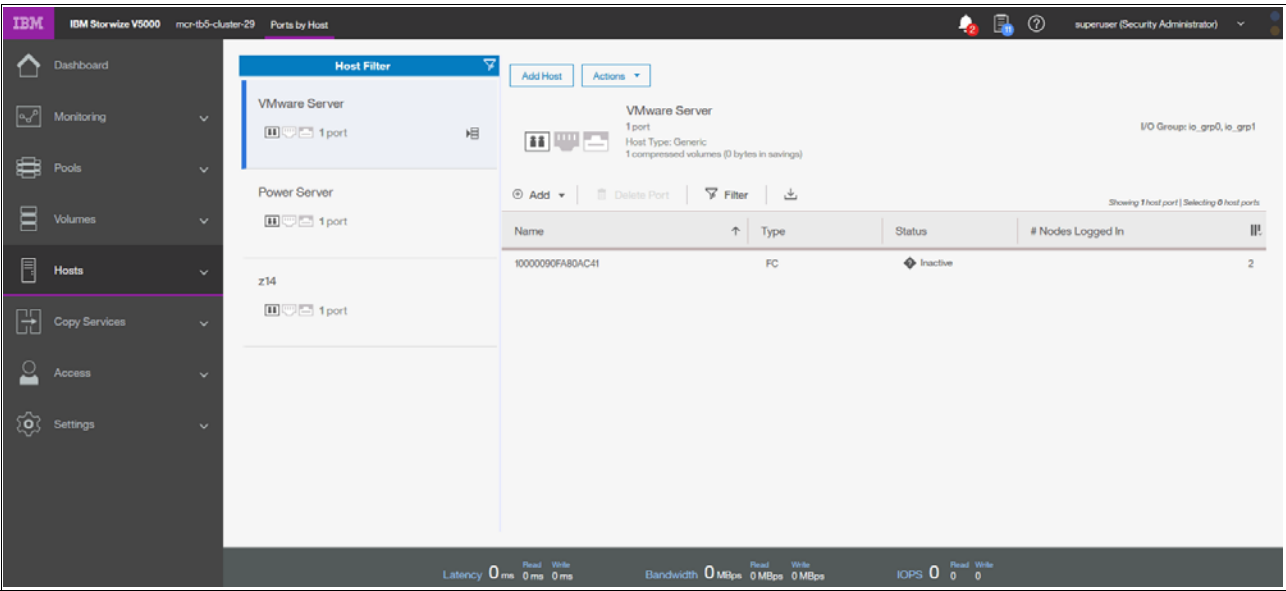


Figure 3-63 Ports by Host window

This overview shows hosts with active, inactive, or degraded ports. You can delete or add a port, or modify its characteristics. Also, in this pane, you can create a new host or rename the existing one.

To perform any of the tasks that are shown in Figure 3-64 on page 116, click **Actions** and select a menu item.

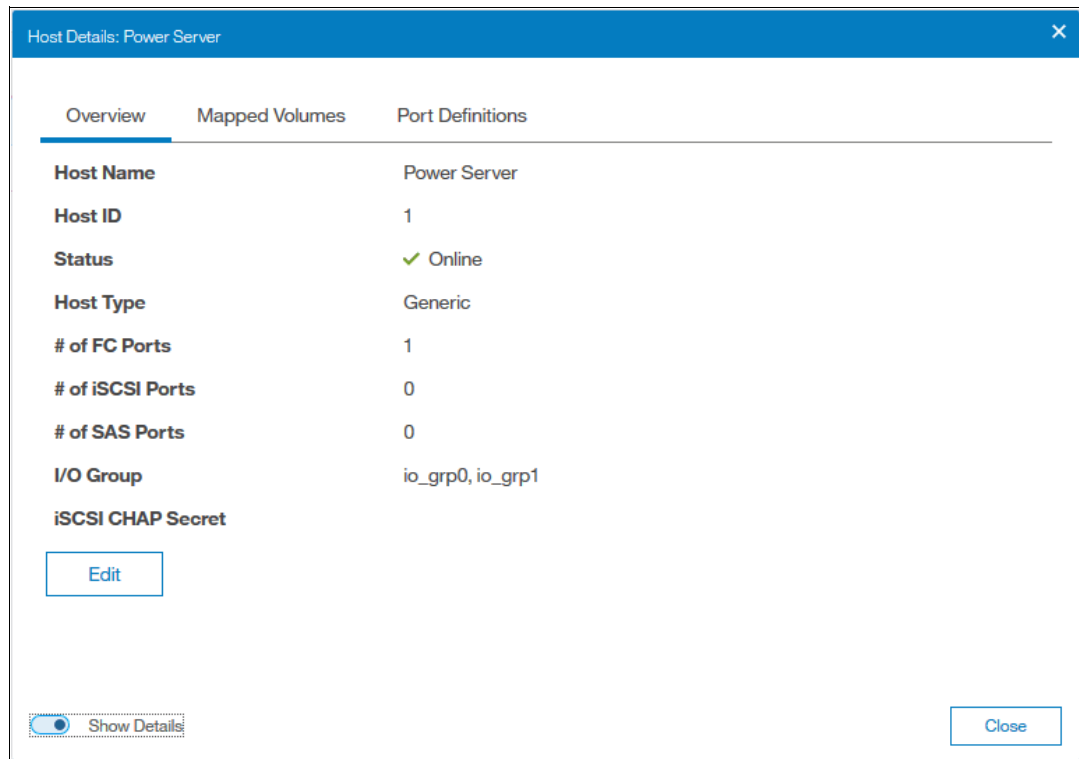


Figure 3-66 Host properties

With enabled details, you can modify host name, host type, I/O group assignment, or iSCSI Challenge Handshake Authentication Protocol (CHAP) Secret by clicking **Edit** and then **Save**, as shown in Figure 3-66.

3.6.5 Volumes by host

This option is identical to the option that is available in the dynamic menu Volumes. For a description, see 3.5.3, "Volumes by host" on page 110.

3.7 Copy services

The copy services are available in all IBM Storwize for Lenovo and Lenovo V series family products. It consists of the following functions:

- ▶ FlashCopy
- ▶ Metro Mirror and Global Mirror
- ▶ Global Mirror with Changed Volumes
- ▶ Volume Mirroring function (Volume Copy)
- ▶ HyperSwap volume mirroring

Figure 3-67 on page 118 shows the Copy Services menu functions.

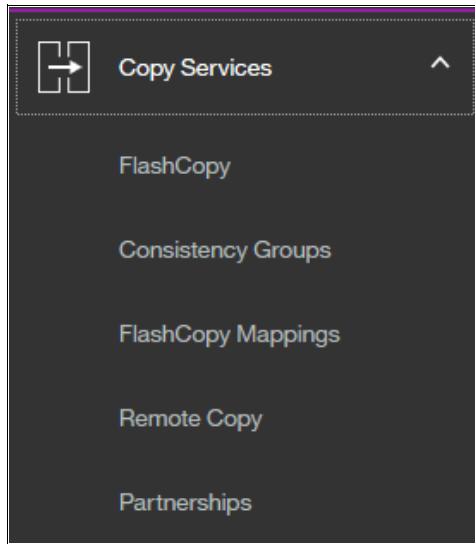


Figure 3-67 Copy Services menu

In this section, we briefly describe how to navigate in the Copy Services menu.

3.7.1 FlashCopy

IBM FlashCopy is a function that you use to create a point-in-time copy of one of your volumes. This function might be helpful when you back up a data or test applications. These copies can be cascaded one on another, read from, written to, and even reversed.

FlashCopy snapshots can conserve storage, if needed, by being space-efficient copies (rather than full copies) that record only items that changed from the originals. Select **FlashCopy** from the dynamic menu to open a pane similar to the information that is shown in Figure 3-68.

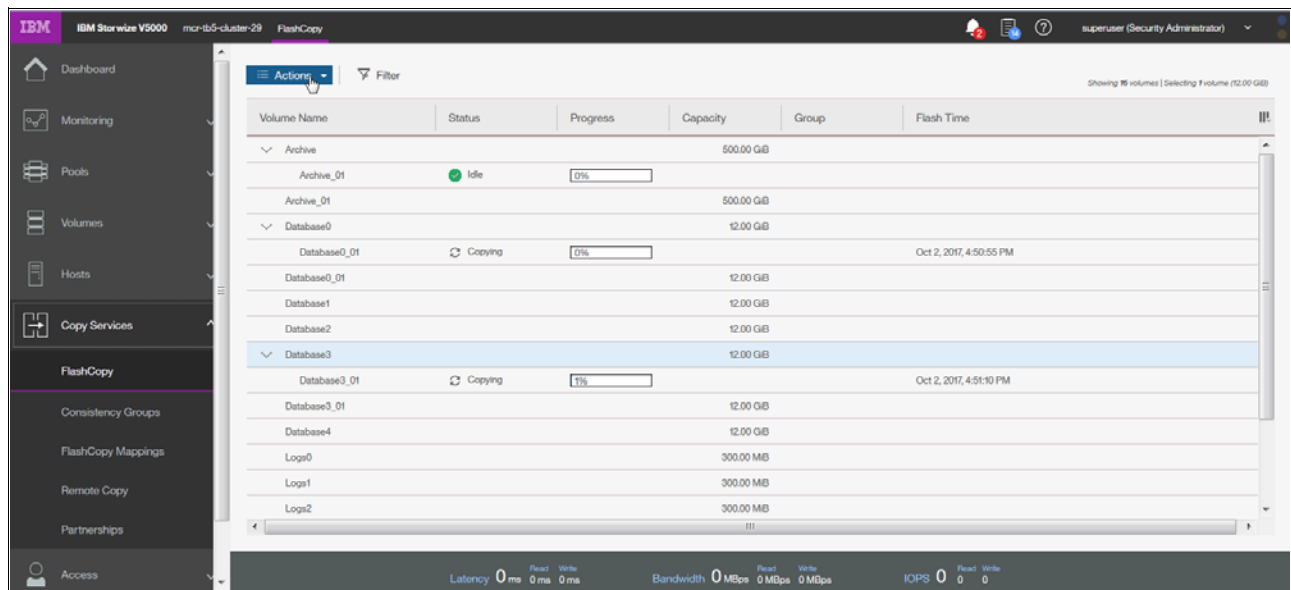


Figure 3-68 FlashCopy operations

If you need to create a FlashCopy of an additional volume, right-click the volume and the list of available functions displays. You can perform several tasks, such as initiate a new snapshot, and clone or back up a volume.

Clicking the volume name opens the window that is shown in Figure 3-69. You can click the tabs at the top of the window to display additional information, such as the hosts that the volume or FlashCopy volume is mapped to and its dependent MDisks.

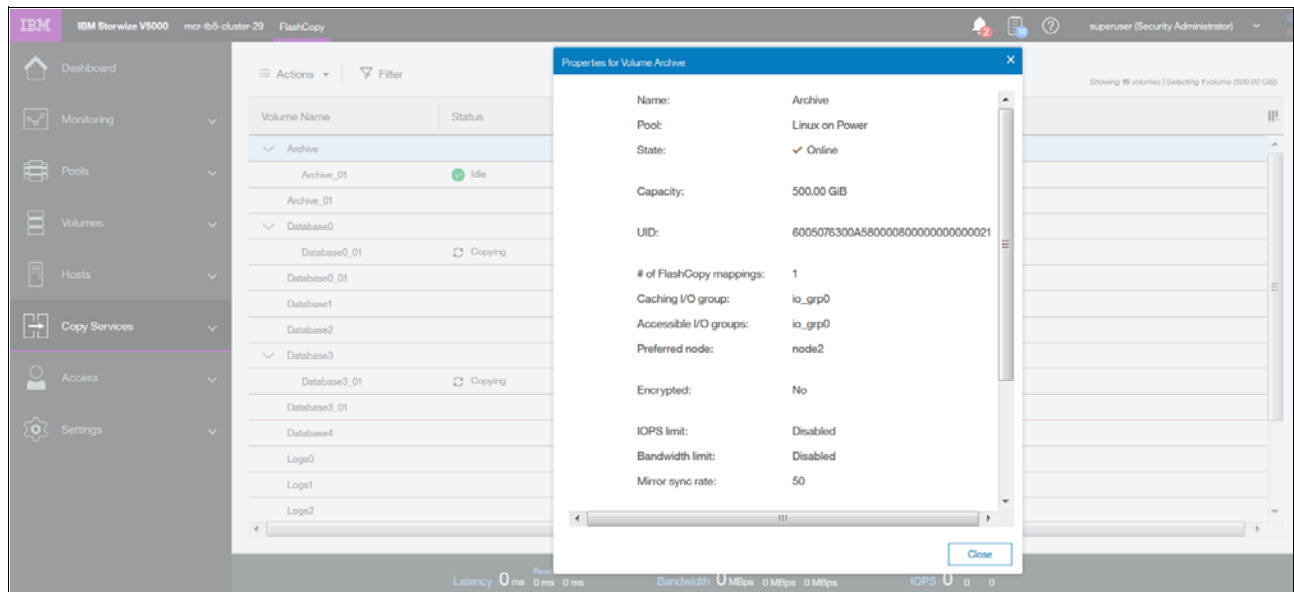


Figure 3-69 FlashCopy volume details

3.7.2 Consistency group

FlashCopy *consistency groups* can be used to create a consistent point-in-time copy across multiple volumes, and even across multiple managed storage systems, managing the consistency of dependent writes.

Click **Consistency Group** to open the window that is shown in Figure 3-70 on page 120. FlashCopy relationships can be placed into a consistency group. You can also use start and stop commands against the FlashCopy consistency group from this window by right-clicking the relationship.

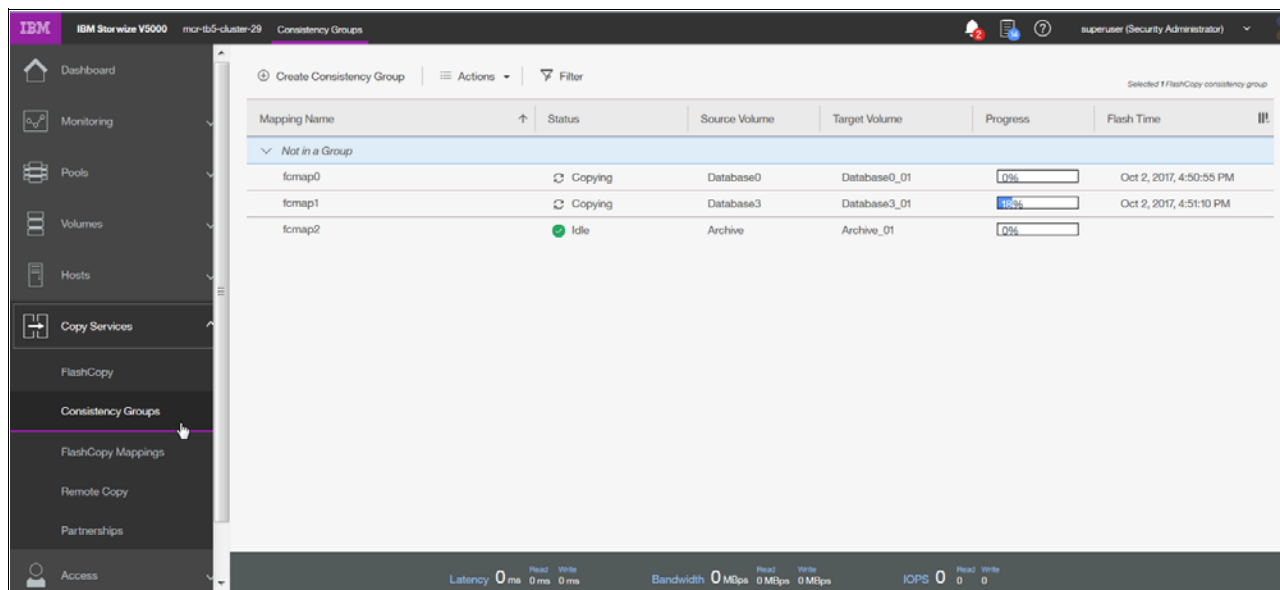


Figure 3-70 FlashCopy Consistency Groups window

When any FlashCopy consistency group is available, either empty or with existing relationships, you can move an existing relationship to that group. Right-click a relationship and select **Move to Consistency Group** as shown in Figure 3-71.

Other actions on the same menu include Remove from Consistency Group, Start (resume) or Stop that FlashCopy operation, Rename Mapping (rename a target volume or FlashCopy mapping), and Delete Mapping.

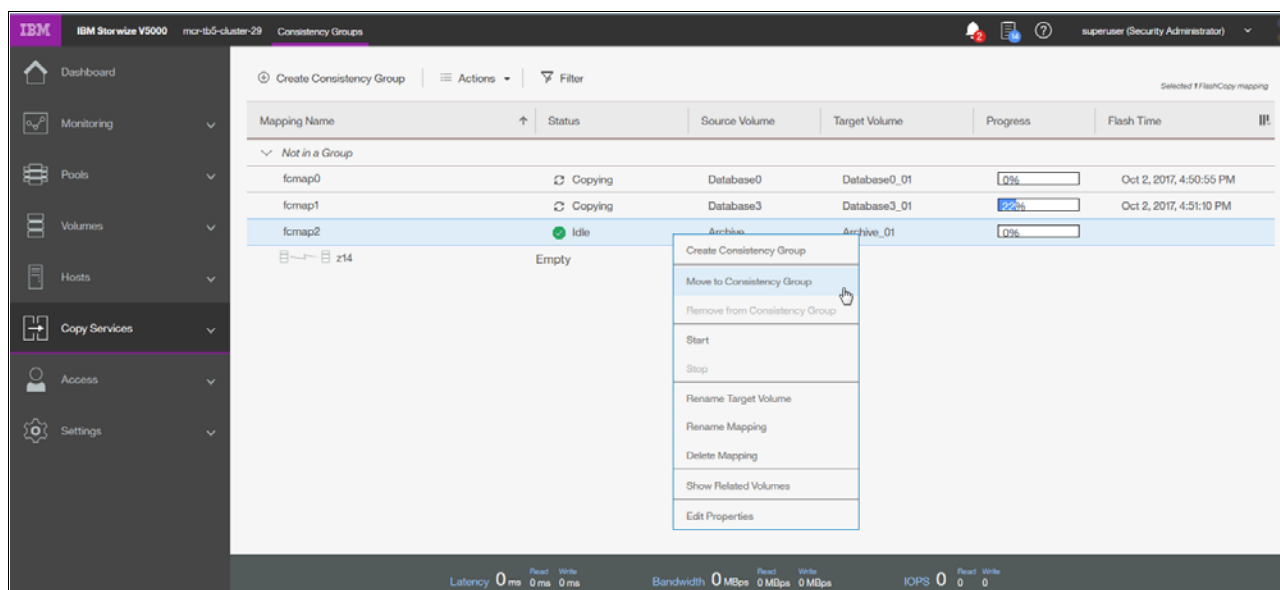


Figure 3-71 Moving a relationship to the consistency group

From the menu, select the appropriate group (in our case, the only one available) and confirm the selection (Figure 3-72 on page 121).

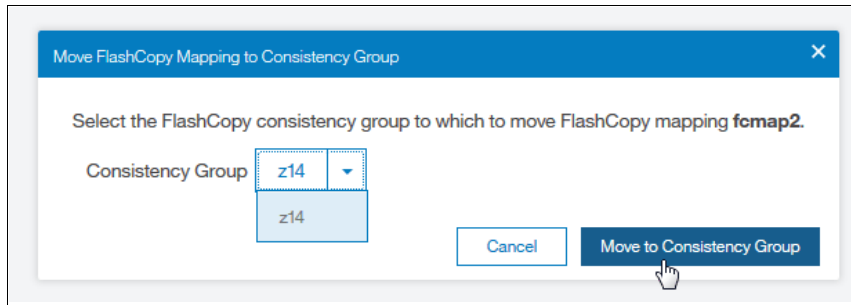


Figure 3-72 Assigning the consistency group

The result of the operation is similar to the result that is shown in Figure 3-73.

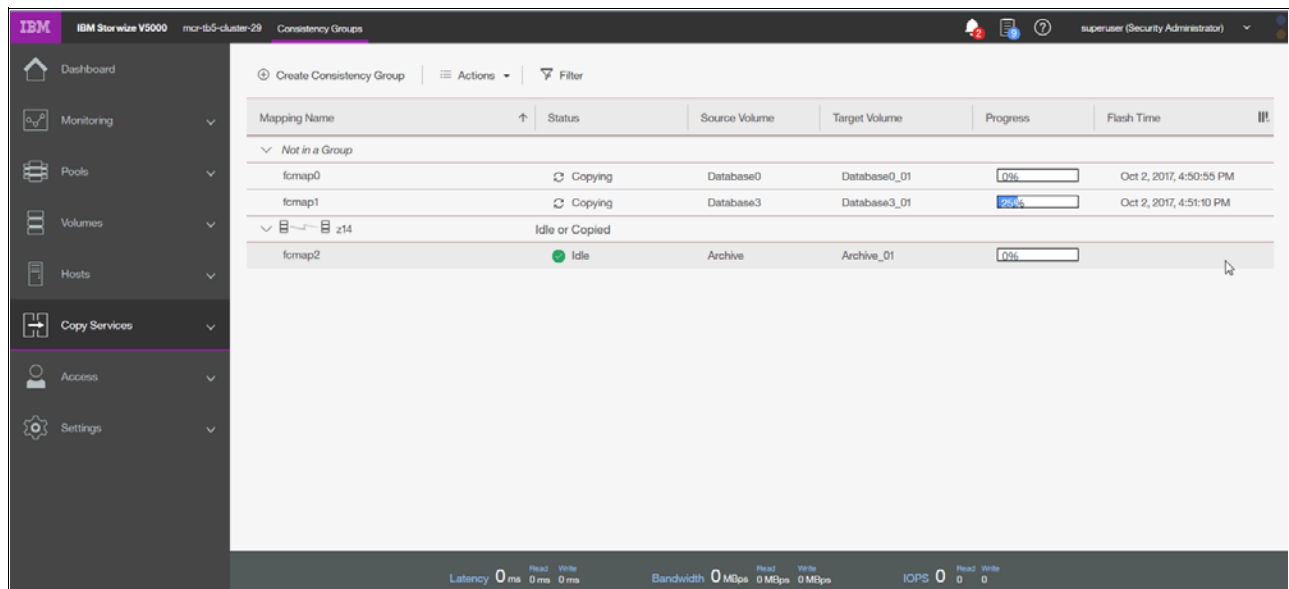


Figure 3-73 Consistency groups

3.7.3 FlashCopy mappings

To create a new FlashCopy mapping, click **Create FlashCopy Mapping** (shown in Figure 3-74 on page 122) to start a wizard. This wizard maps a source volume to a target volume to prepare for a subsequent copy. This mapping persists until it is deleted. The mapping specifies the source and destination volumes. The destination must be identical in size to the source or the mapping fails.

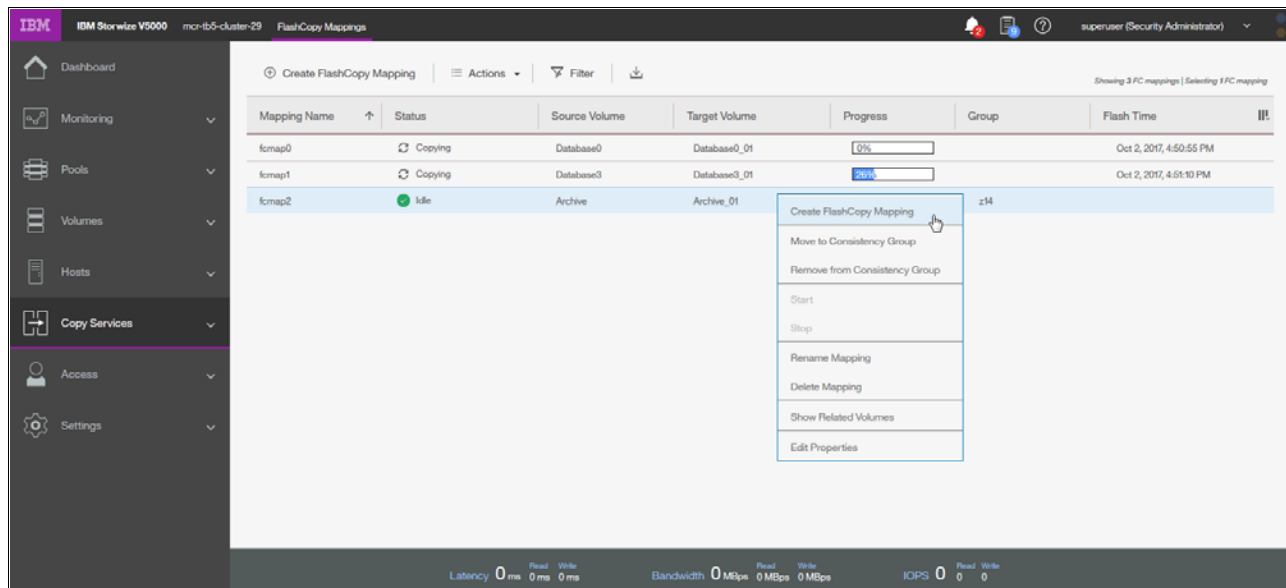


Figure 3-74 FlashCopy mappings

In a single mapping, the source and destination cannot be on the same volume. A mapping is triggered at the point in time when the copy is required. The mapping can optionally be given a name and assigned to a consistency group. These groups of mappings can be triggered at the same time, enabling multiple volumes to be copied at the same time, which creates a consistent copy of multiple disks. A consistent copy of multiple disks is required for database products in which the database and log files are on separate disks.

If a consistency group (ID or Name) is not specified, the mapping is assigned to the default *group 0*, which is a special group that cannot be started as a whole. Mappings in this group can be started only on an individual basis.

An example of the wizard for FlashCopy mapping creation is shown in Figure 3-75 on page 123. Select source and target volumes from the wizard.

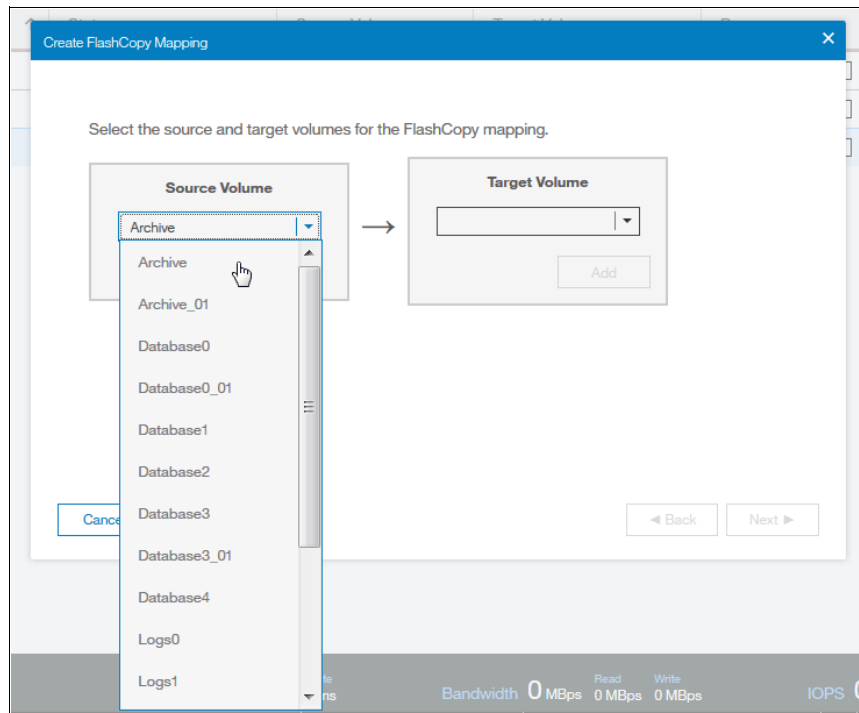


Figure 3-75 Selecting volumes for FlashCopy mappings

You can select the Snapshot (copy-on-write), Clone (replica of the volume without effect on original one), or Backup (data recovery) type of relationship. When selected, you can specify whether you also want to add the mapping to the consistency group.

3.7.4 Remote copy

Click **Remote Copy** to open the window that is shown in Figure 3-76. This window shows the existing remote copy relationships, and you can set up and modify consistency groups. From this window, you can also start and stop relationships, add relationships to a consistency group, and switch the direction of the mirror.

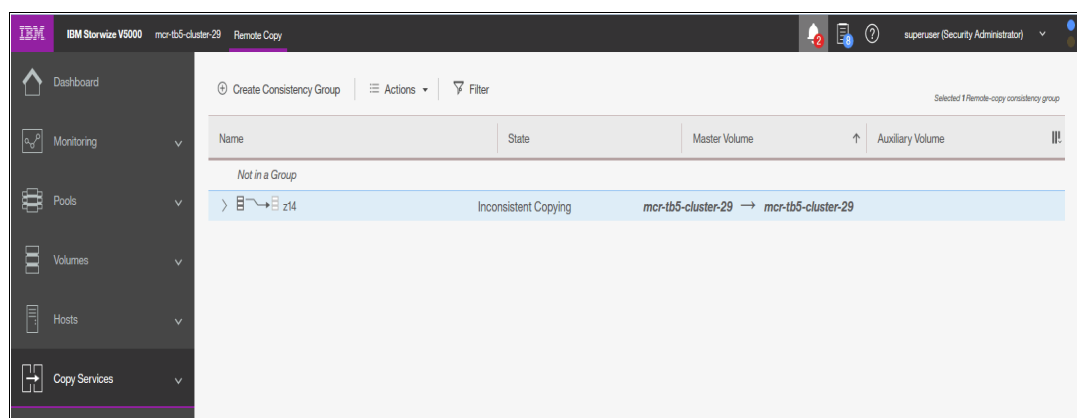


Figure 3-76 Remote Copy window

The menu provides the options to create Metro Mirror, Global Mirror, or Global Mirror with Changed Volumes:

Metro Mirror

This option makes *synchronous* copies. The original write operations are not considered complete until the write operation to the destination disk is confirmed. The distance between your two sites is determined by how much latency your applications can handle.

Global Mirror

This option makes *asynchronous* copies of your disk. The write is considered complete after it is complete at the local disk. It does not wait for the write to be confirmed at the remote cluster as Metro Mirror does. This method greatly reduces the latency that is experienced by your applications if the other cluster is far away. However, it also means that during a failure, the data on the remote copy might not contain the most recent changes that were committed to the local disk.

Global Mirror with Changed Volumes

This option is best described as “Continuous Remote FlashCopy.” If you use this feature, controller firmware essentially takes a periodic FlashCopy of a disk and writes it to your remote destination. This feature completely isolates the local copy from wide area network (WAN) issues and from sudden spikes in workload that might occur. The drawback is that your remote copy might lag behind the original by a significant amount of data, depending on how you set up the cycle time.

3.7.5 Partnerships

Click **Partnerships** to open the window that is shown in Figure 3-77. You can use this window to set up a new partnership, or delete an existing partnership for remote mirroring with another Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems. To create a partnership, click **Create Partnership**. A new window displays. When you select the partnership type, for example, Fibre Channel, the window expands to a more detailed view, as shown in Figure 3-77.

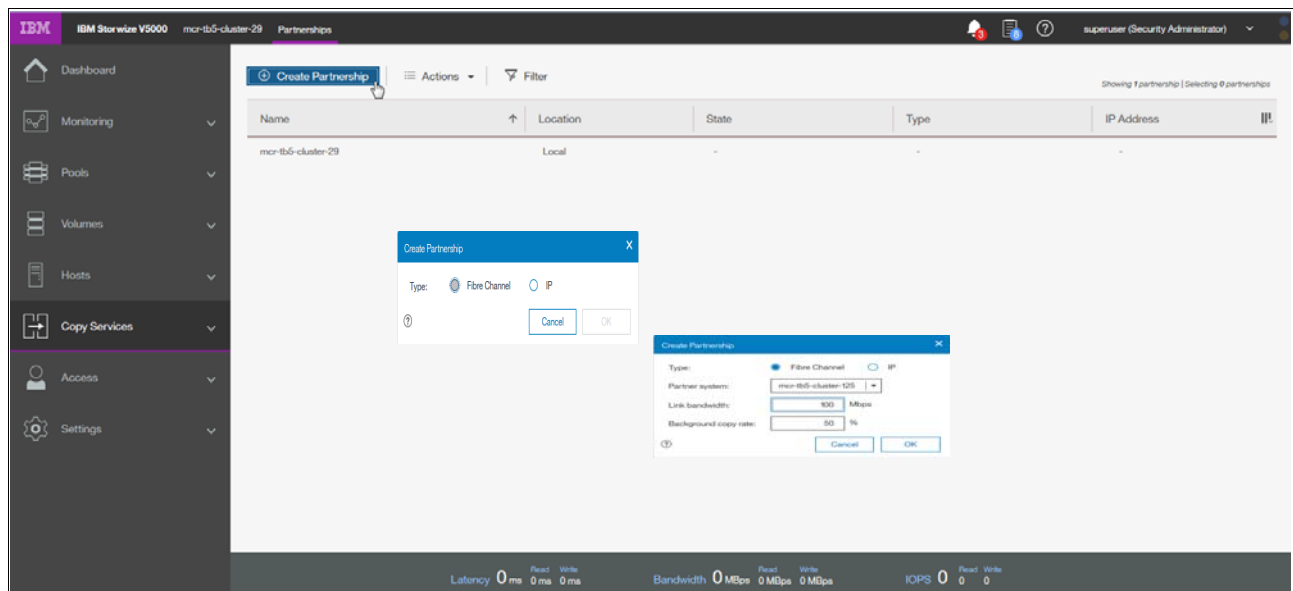


Figure 3-77 Creating a partnership

Clicking an existing partnership opens a window, as shown in Figure 3-78 on page 125. From this window, you can also set the background copy rate. This rate specifies the bandwidth, in

Mbps, that is used by the background copy process between the clusters (Figure 3-78). In our case, we configured the partnership only on one side. You can see it in the State row. It shows **Partially Configured: Local**, which is an indication that the configuration was only configured on one side. If you see this message, go to the second system, and configure the Create Partnership settings there, too.

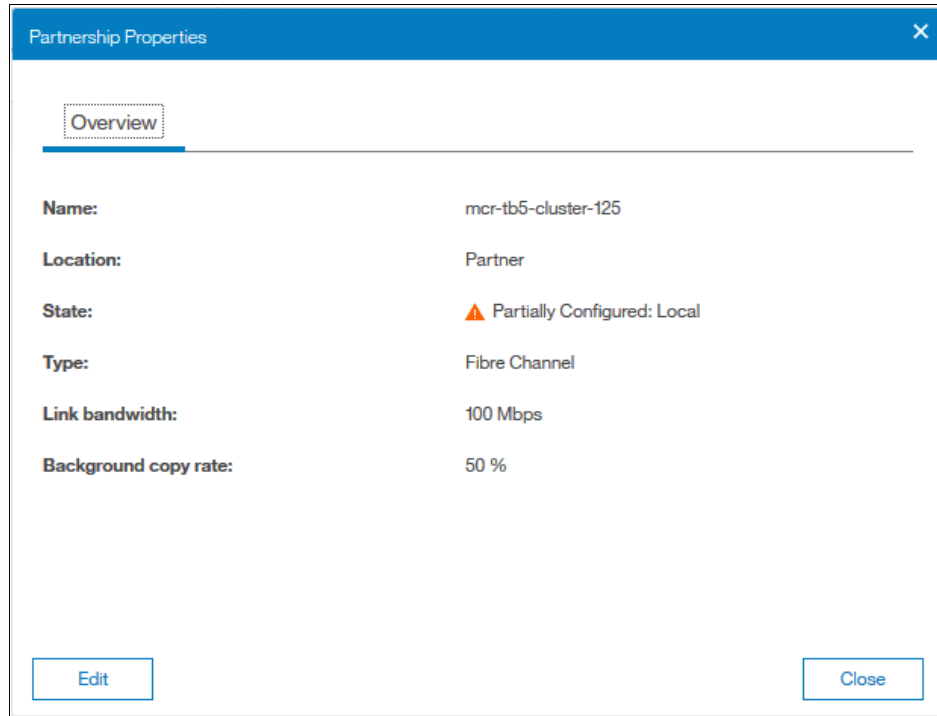


Figure 3-78 Partnership properties

3.8 Access menu

The Access menu has two options:

- ▶ Users (for user management)
- ▶ Audit Log

Figure 3-79 shows these options.

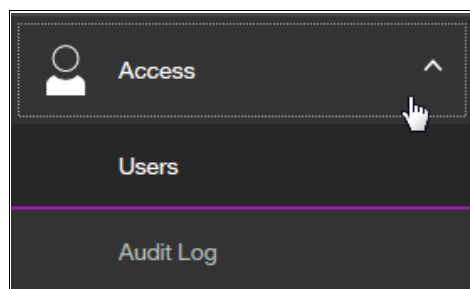


Figure 3-79 Access menu

3.8.1 Users

Figure 3-80 shows the Users pane. You can create and delete new users, change and remove passwords, and add and remove Secure Shell (SSH) keys.

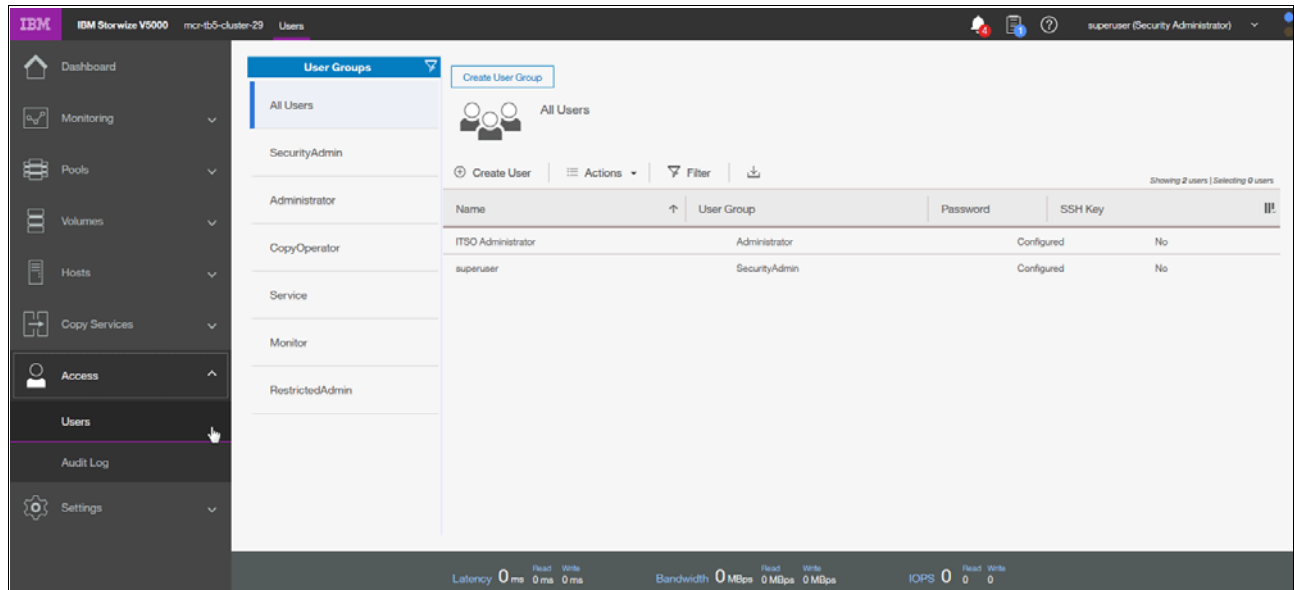


Figure 3-80 Users window

Click **Create User** to open the pane that is shown in Figure 3-81. Use this pane to specify the name and password of the user, and load the SSH key (if the SSH key was generated). SSH key is not required for CLI access, and you can choose to use either SSH or a password for CLI authentication.

The 'Create User' dialog box has a blue header with a close button. It contains a 'Name' field with a user icon and the text 'Security Admin'. Below this is the 'Authentication Mode' section with two radio buttons: 'Local' (selected) and 'Remote'. Underneath is a 'User Group' dropdown menu showing 'SecurityAdmin'. The 'Local Credentials' section has a note: 'Users must have a password, an SSH public key, or both.' It includes a 'Password' field, a 'Verify password' field, and an 'SSH Public Key' section with a 'Browse...' button and the text 'No file selected.' At the bottom are 'Cancel' and 'Create' buttons.

Figure 3-81 Adding a user

3.8.2 Audit Log option

Click **Audit Log** to open the window that is shown in Figure 3-82. The cluster maintains an audit log of successfully run commands and displays the users that performed particular actions at certain times.

Date and Time	User Name	Command	Object ID
10/3/17 11:37:36 AM	superuser	svctask mkuser -gui -name "Security Admin" -password #### -use...	2
10/3/17 11:32:56 AM	superuser	svctask mkuser -gui -name "ITSO Administrator" -password #### -...	1
10/2/17 5:33:33 PM	superuser	svctask mkfcpartnership -backgroundcopyrate 50 -gui -linkbandwi...	
10/2/17 5:28:37 PM	superuser	svctask startroconsistgrp -gui 0	
10/2/17 5:25:51 PM	superuser	svctask chrrelationship -consistgrp 0 -gui 2	
10/2/17 5:25:50 PM	superuser	svctask chrrelationship -cyclingmode multi -gui 2	
10/2/17 5:25:50 PM	superuser	svctask chrrelationship -gui -masterchange 16 2	
10/2/17 5:25:50 PM	superuser	svctask chrrelationship -auxchange 17 -gui 2	
10/2/17 5:25:49 PM	superuser	svctask mkvdisk -autoexpand -gui -iogrp io_grp0 -mdiskgrp "Linux...	17
10/2/17 5:25:49 PM	superuser	svctask mkvdisk -autoexpand -gui -iogrp io_grp0 -mdiskgrp "Linux...	16
10/2/17 5:25:46 PM	superuser	svctask mkrelationship -aux Database1 -cluster mcr-tb5-cluster-...	2
10/2/17 5:25:44 PM	superuser	svctask mkroconsistgrp -gui -name z14	0
10/2/17 5:17:48 PM	superuser	svctask chfomap -consistgrp z14 -gui 2	
10/2/17 5:13:07 PM	superuser	svctask mkfconsistgrp -gui -name z14	1
10/2/17 5:10:54 PM	superuser	svctask chenclosureslot -gui -identify no -slot 11	

Figure 3-82 Audit Log entries

You can filter audit log records by date or within a specific time frame (Figure 3-83).

Date and Time	User Name	Command	Object ID
10/3/17 11:37:36 AM	superuser	svctask mkuser -gui -name "Security Admin" -password #### -use...	2
10/3/17 11:32:56 AM	superuser	svctask mkuser -gui -name "ITSO Administrator" -password #### -...	1
10/2/17 5:33:33 PM	superuser	svctask mkfcpartnership -backgroundcopyrate 50 -gui -linkbandwi...	
10/2/17 5:28:37 PM	superuser	svctask startroconsistgrp -gui 0	
10/2/17 5:25:51 PM	superuser	svctask chrrelationship -consistgrp 0 -gui 2	
10/2/17 5:25:50 PM	superuser	svctask chrrelationship -cyclingmode multi -gui 2	
10/2/17 5:25:50 PM	superuser	svctask chrrelationship -gui -masterchange 16 2	
10/2/17 5:25:50 PM	superuser	svctask chrrelationship -auxchange 17 -gui 2	
10/2/17 5:25:49 PM	superuser	svctask mkvdisk -autoexpand -gui -iogrp io_grp0 -mdiskgrp "Linux...	17
10/2/17 5:25:49 PM	superuser	svctask mkvdisk -autoexpand -gui -iogrp io_grp0 -mdiskgrp "Linux...	16
10/2/17 5:25:46 PM	superuser	svctask mkrelationship -aux Database1 -cluster mcr-tb5-cluster-...	2
10/2/17 5:25:44 PM	superuser	svctask mkroconsistgrp -gui -name z14	0
10/2/17 5:17:48 PM	superuser	svctask chfomap -consistgrp z14 -gui 2	
10/2/17 5:13:07 PM	superuser	svctask mkfconsistgrp -gui -name z14	1
10/2/17 5:10:54 PM	superuser	svctask chenclosureslot -gui -identify no -slot 11	

Figure 3-83 Filtering the records

The following commands are *not* recorded in the audit log:

- ▶ All commands that failed
- ▶ dumpconfig
- ▶ cpdumps
- ▶ cleardumps

- ▶ finderr
- ▶ dumperrlog
- ▶ dumpinternallog
- ▶ svc servicetask dumperrlog
- ▶ svc servicetask finderr

3.9 Settings menu

The Settings menu provides various options to adjust your system parameters according to your needs. You can configure these options (Figure 3-84):

- ▶ Notifications
- ▶ Network
- ▶ Security (remote authentication)
- ▶ System
- ▶ Support
- ▶ GUI Preferences

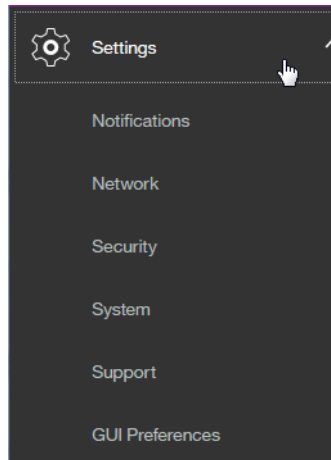


Figure 3-84 Settings menu

3.9.1 Notifications

It is important to correct any issues that are reported by your system as soon as possible. Configure your system to send automatic notifications when a new event is reported. To avoid monitoring for new events that use the management GUI, select the type of event that you want to be notified about, for example, restrict notifications to events that require immediate action.

You can use email, Simple Network Management Protocol (SNMP), or syslog types of notifications. If your system is within warranty, or if you use a hardware maintenance agreement, configure your Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems to send email events to Lenovo directly if an issue that requires hardware replacement is detected. This mechanism is called *Call Home*.

When an event is received, Lenovo automatically opens a problem report. If appropriate, Lenovo contacts you to verify whether replacement parts are required. The configuration window for e-mail notifications is shown in Figure 3-85.

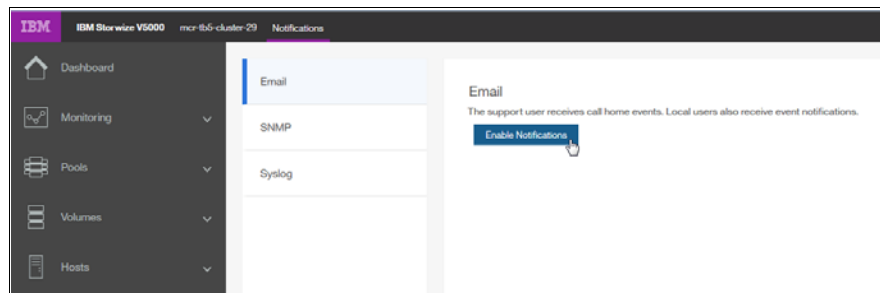


Figure 3-85 Email event notifications

The procedure for how to enable e-mail notifications is described in Chapter 12, “RAS, monitoring, and troubleshooting” on page 625.

3.9.2 Network

Click **Network** to open the window that is shown in Figure 3-86. You can update the network configuration, set up iSCSI definitions, and view information about the Fibre Channel connections.

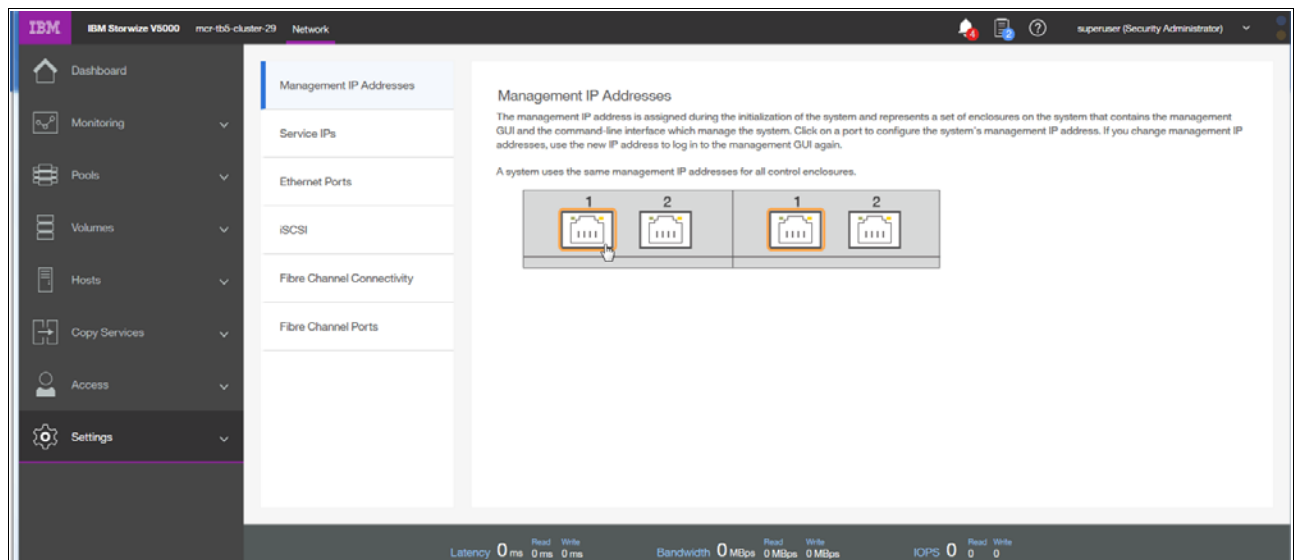


Figure 3-86 Network window

When you click **Fibre Channel Connectivity** (Figure 3-87 on page 130), useful information is displayed. In this example, we click '**All nodes, storage systems, and hosts**' from the menu and then select *Show Results* to display the details. Other options that are available from the menu include displaying Fibre Channel details for a host, clusters, nodes, or storage systems.

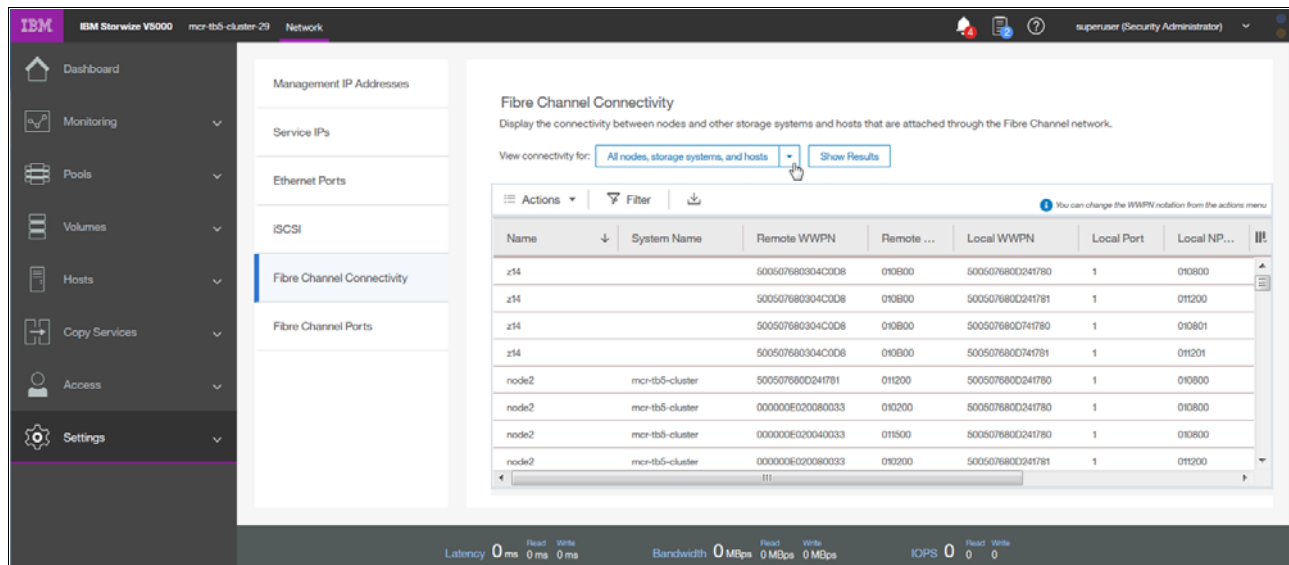


Figure 3-87 Fibre Channel connectivity

3.9.3 Security

The different security features are described below.

Remote authentication

With security and its directory services, the user can remotely authenticate to the controller firmware without the need for a local account. Therefore, when you log on, you authenticate with your domain user ID and password rather than a locally created user ID and password.

The benefits of remote authentication are listed:

- ▶ You do not need to configure every user on every machine. If multiple machines are in your environment, you can set up authentication more efficiently.
- ▶ When commands are run on the controller firmware, the audit log shows the domain user name that issued that command, rather than a local user name, or worse, just “superuser”. (In this case, determining who mapped a volume, acted as the superuser, and so on might be difficult.)
- ▶ You have central control over access. If someone leaves the company, you simply remove access at the domain controller, which means that orphan accounts do not remain on your storage equipment.

The access pane to configure remote authentication is shown in Figure 3-88 on page 131.

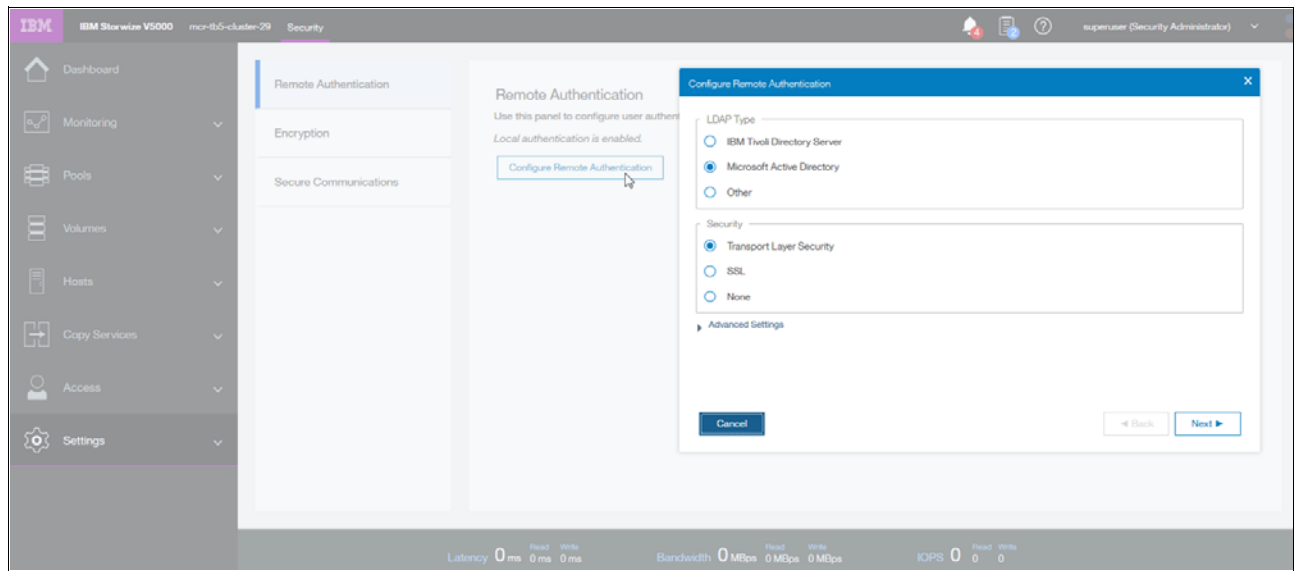


Figure 3-88 Configuring remote authentication

The detailed steps to configure remote logon are described at the following web pages:

- ▶ <http://ibm.biz/Bd4Cr5>
- ▶ <https://ibm.biz/BdjSL5>

Encryption

On the panel that is shown in Figure 3-89, you can enable or disable the encryption function on an Lenovo Storage V5030. The panel shows that no USB drives that contain encryption keys were detected. These are no longer needed if you have an external Key Management Server. Figure 3-89 shows four available external Key Management Server. If there is no external Server available, you need the USB keys to be able to en/decrypt your data on power on.

External Encryption Management server is implemented in the controller firmware 8.1.

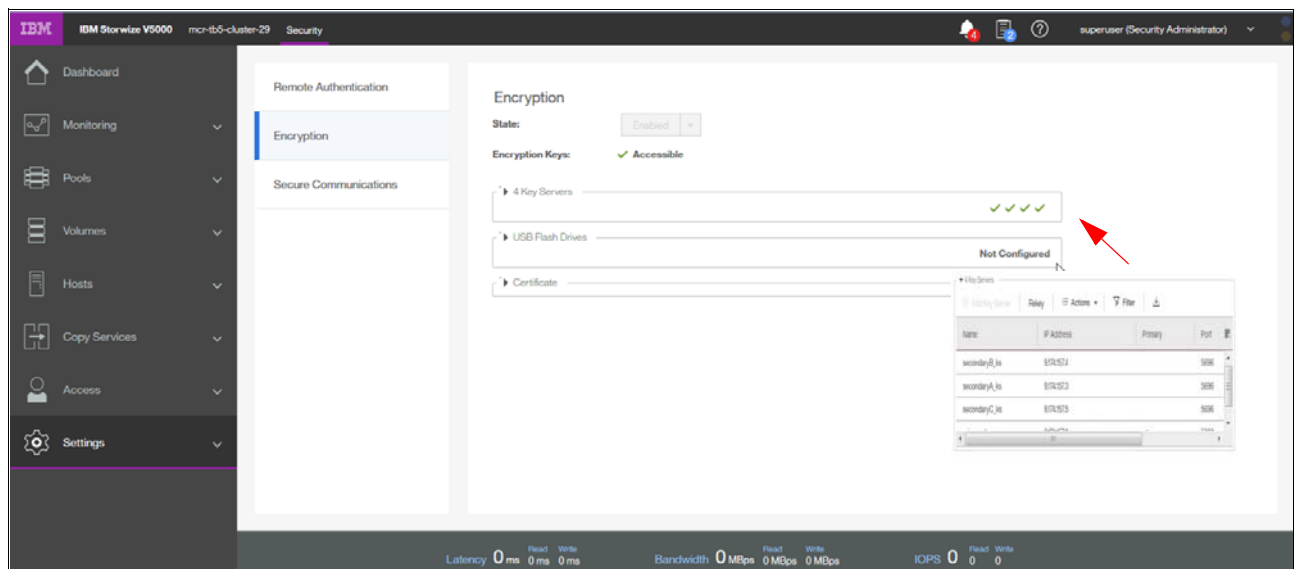


Figure 3-89 Encryption panel

Secure communications

Use the Secure Communications page to enable and manage secure connections. During system setup, an initial certificate is created to use for secure connections between web browsers. Based on the security requirements for your system, you can create either a new self-signed certificate or install a signed certificate that is created by a third-party certificate authority.

Self-signed certificates are generated automatically by the system and encrypt communications between the browser and the system. Self-signed certificates can generate web browser security warnings, and they might not comply with organizational security guidelines. Signed certificates are created by a third-party certificate authority. These certificate authorities ensure that certificates have the required security level for an organization based on purchase agreements. Signed certificates usually have higher security controls for the encryption of data and do not cause browser security warnings.

To use a signed certificate, first generate and download a request for a certificate that is based on the values that are specified on the Secure Communication page. Submit this request to the certificate authority to receive a signed certificate and then install it by using the Secure Communication page. Before you create a request for either type of certificate, ensure that your current browser does not restrict the type of keys that are used for certificates. Certain browsers limit the use of specific key types for security and compatibility.

On Figure 3-90, you can see the details about the security certificates.

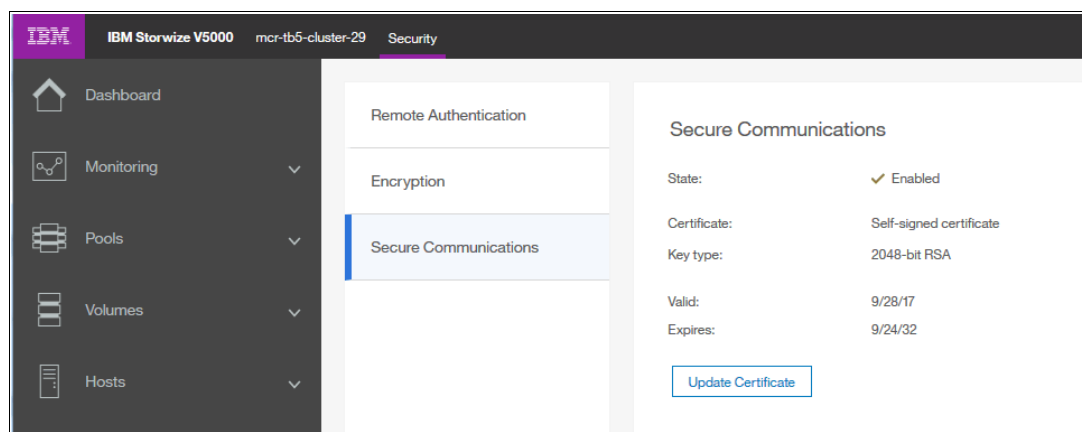


Figure 3-90 Secure communications

If you want to update or change the certificate, click **Update Certificate**.

The Update Certificate panel opens, as shown in Figure 3-91 on page 133.

Figure 3-91 Update Certificate panel

3.9.4 System

The System menu provides the following options:

- ▶ Set the system date and time
- ▶ Manage licenses
- ▶ Upgrade System
- ▶ Virtual Volumes (VVols)
- ▶ IP Quorum
- ▶ I/O Groups
- ▶ DNS

The Date and Time window opens (Figure 3-92 on page 134) when you select **Date and Time** from the System menu. You can add a Network Time Protocol (NTP) server if available.

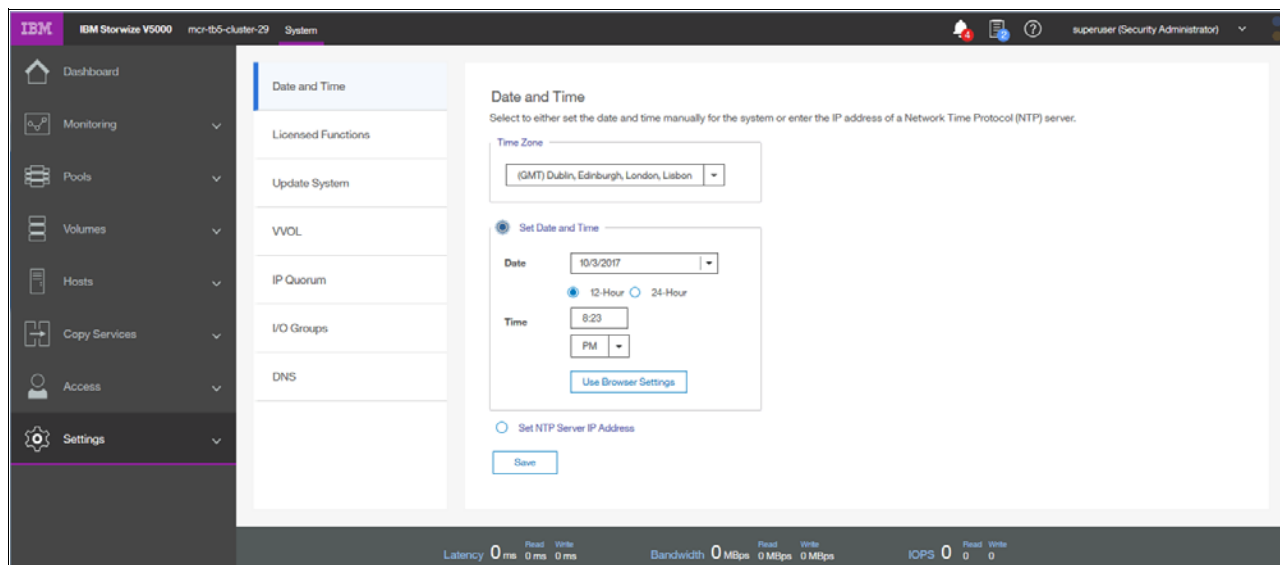


Figure 3-92 Date and Time window

You can also update the license information for specific features, as shown in Figure 3-93.

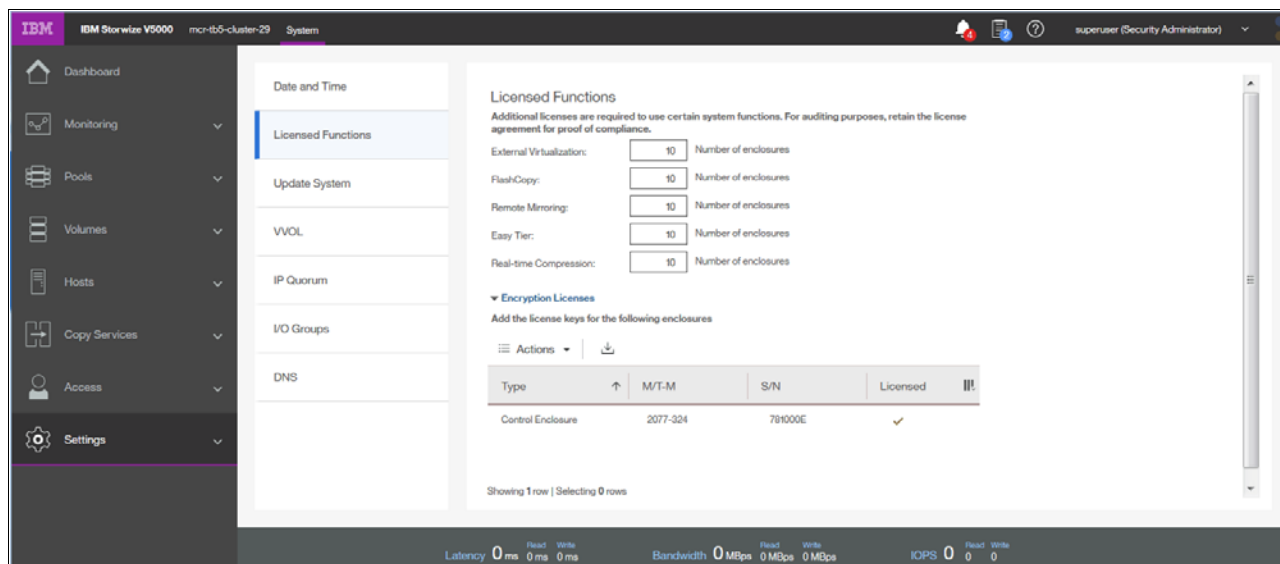


Figure 3-93 Licensing options

To upgrade your controller firmware, use the procedure that is described in Chapter 12, “RAS, monitoring, and troubleshooting” on page 625.

Virtual Volume (VVol) is a tape volume that resides in a tape volume cache of a virtual tape server (VTS). VVol is a new feature that was introduced in controller firmware 7.6. With this new functionality, users can create volumes on your system directly from a VMware vCenter server.

On the VVOL page, you can enable or disable the functionality, as shown in Figure 3-94 on page 135. Before you can enable VVol, you must set up an NTP server. See the Date and Time settings to set up the NTP server.

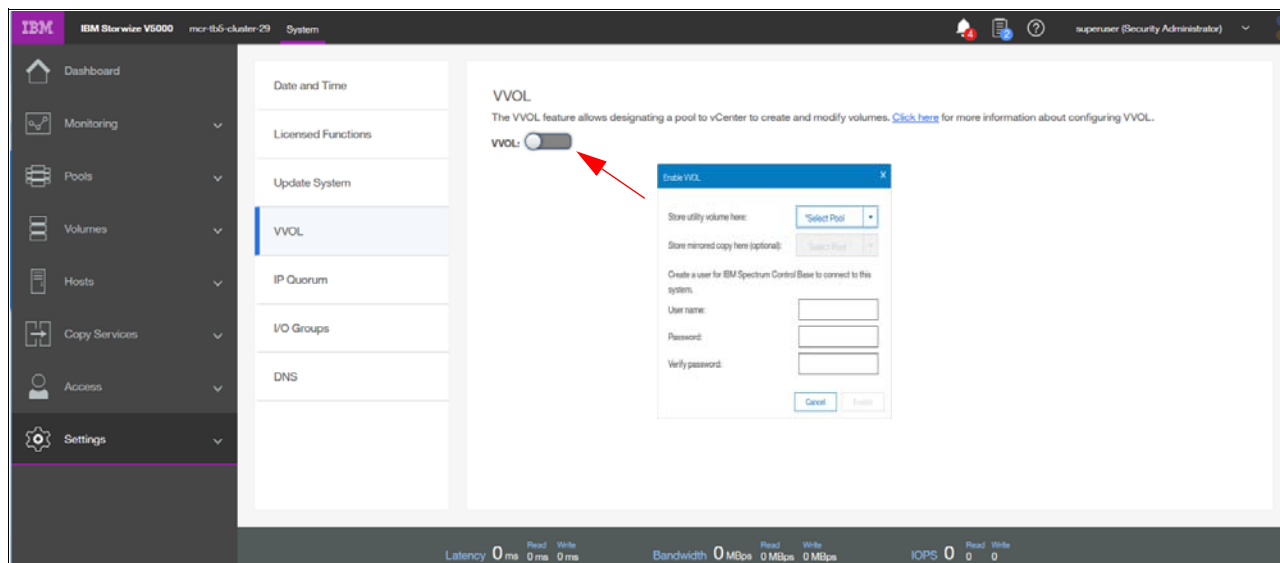


Figure 3-94 Activating VVol

In some HyperSwap configurations, IP quorum applications can be used at the third site as an alternative to third-site quorum disks. No Fibre Channel connectivity at the third site is required to use an IP quorum application as the quorum device. The IP quorum application is a Java application that runs on a host at the third site. The IP network is used for communication between the IP quorum application and node canisters in the system. If you currently have a third-site quorum disk, you must remove the third site before you use an IP quorum application. The round trip time limitations from 80 micro seconds for a IP quorum are still existent. Figure 3-95 shows where you can download the Java application.

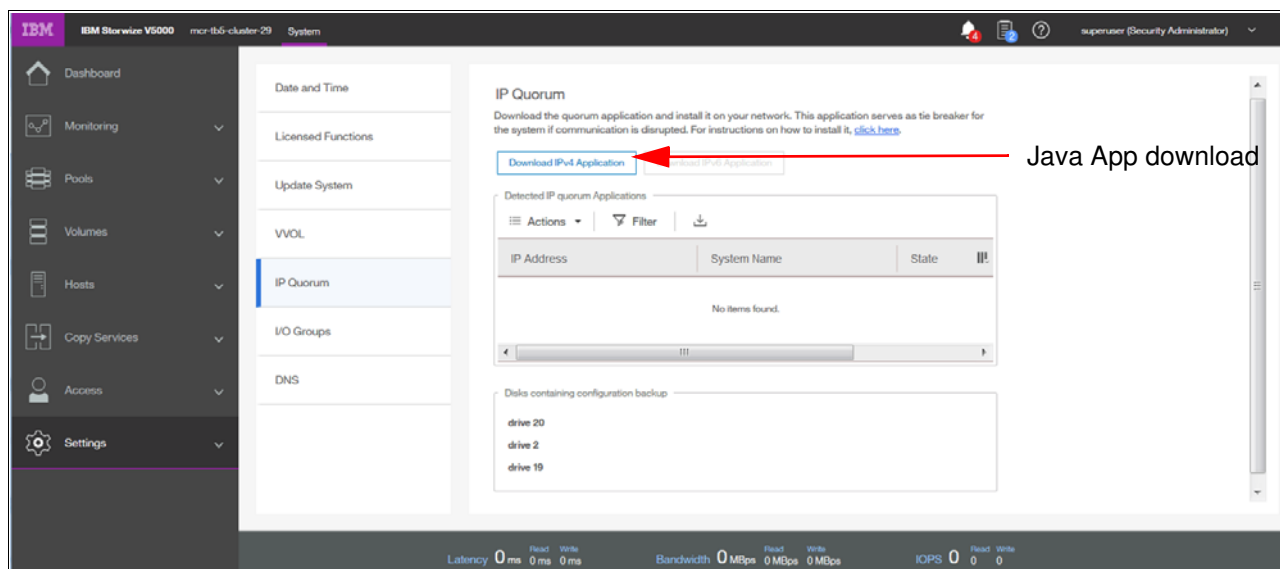


Figure 3-95 IP Quorum

For ports within an I/O group, you can enable virtualization of Fibre Channel ports that are used for host I/O operations. With N_Port ID virtualization (NPIV), the Fibre Channel port consists of both a physical port and a virtual port. When port virtualization is enabled, ports do not come up until they are ready to handle I/O, which improves host behavior around node unpends. In addition, path failures due to an offline node are masked from hosts. The target

port mode on the I/O group indicates the current state of port virtualization. Figure 3-96 shows the panel with the I/O Groups.

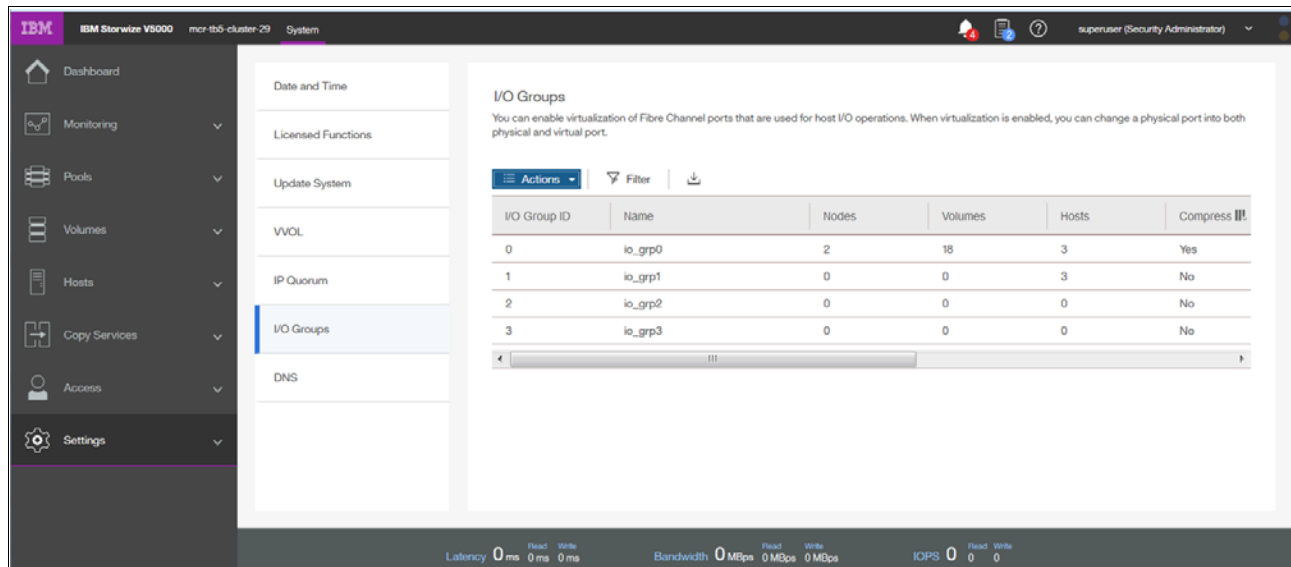


Figure 3-96 I/O Groups

Domain Name System (DNS) translates IP address to host names. You can create, delete, or change domain name servers, which manage names of resources that are located on external networks.

You can have up to two DNS servers that are configured on the system. To configure DNS for the system, enter a valid IP address and name for each server. Both IPv4 and IPv6 address formats are supported. Figure 3-97 shows the DNS setup Window.

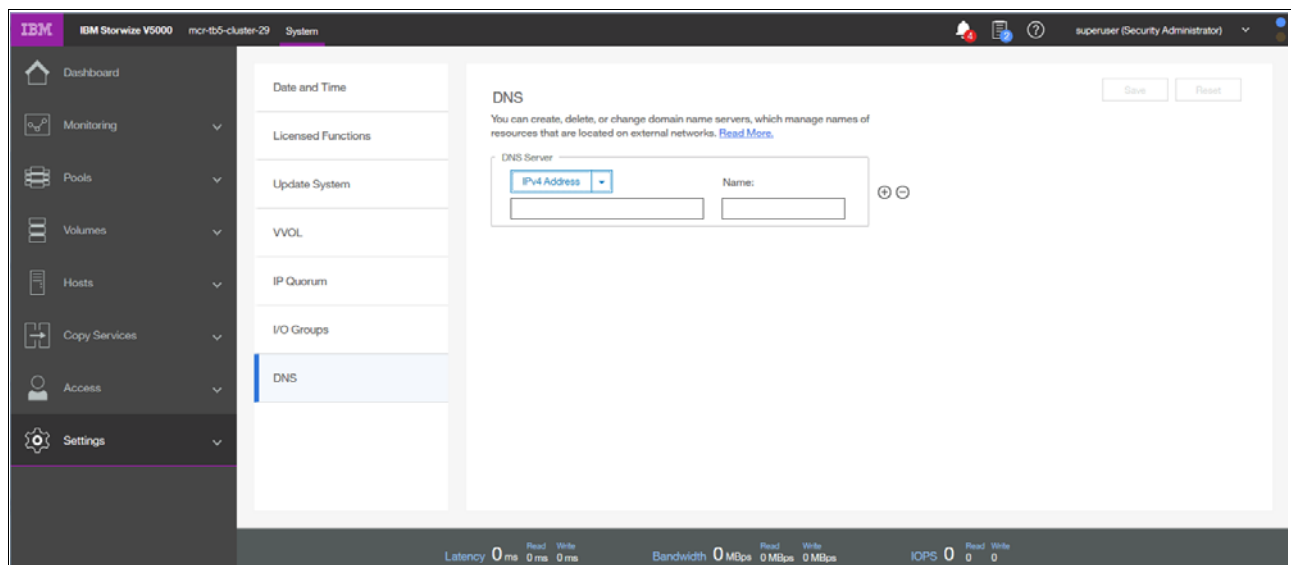


Figure 3-97 DNS

3.9.5 Support

Support assistance enables support personnel to access the system to complete troubleshooting and maintenance tasks. You can configure either local support assistance,

where support personnel visit your site to fix problems with the system, or remote support assistance. Both local and remote support assistance uses secure connections to protect data exchange between the support center and system. To enable Support assistance you need to enable an email Server. More access controls can be added by the system administrator. Figure 3-98 shows the Support assistance panel.

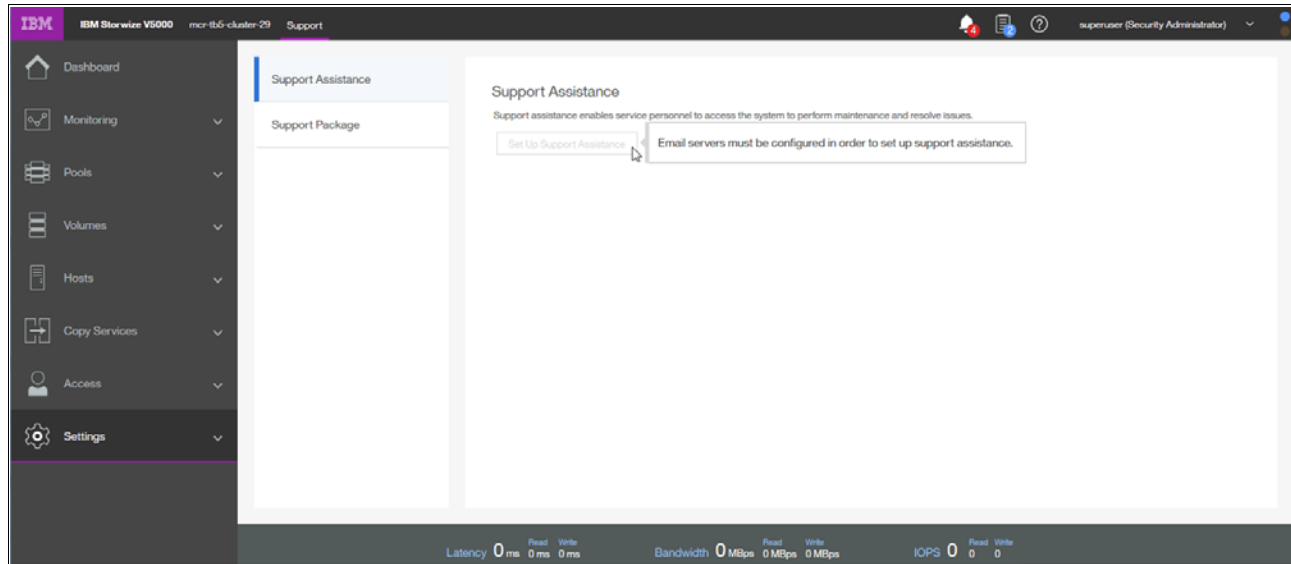


Figure 3-98 Support assistance

If support assistance is configured on your system, you can either automatically or manually upload new support packages to the support center to help analyze and resolve errors on the system. You can select individual logs to either download to review or send directly to the support center for analysis.

Figure 3-99 shows how to upload or download support logs.

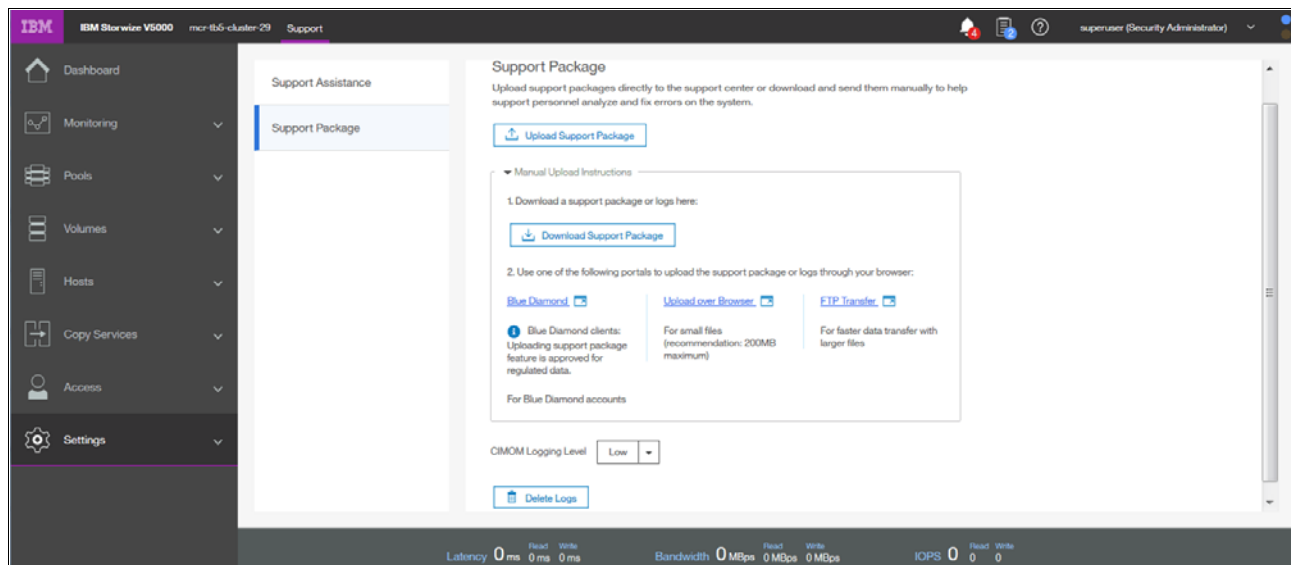


Figure 3-99 Up- or downloading support packages

For more information, see Chapter 12, “RAS, monitoring, and troubleshooting” on page 625.

3.9.6 GUI preferences

By using this menu, you can configure the appearance and behavior of the GUI. Click **GUI Preferences** in the Settings option of the Dynamic menu. To display the login message, select **Enable**. You can create a customized login message as shown in Figure 3-100.

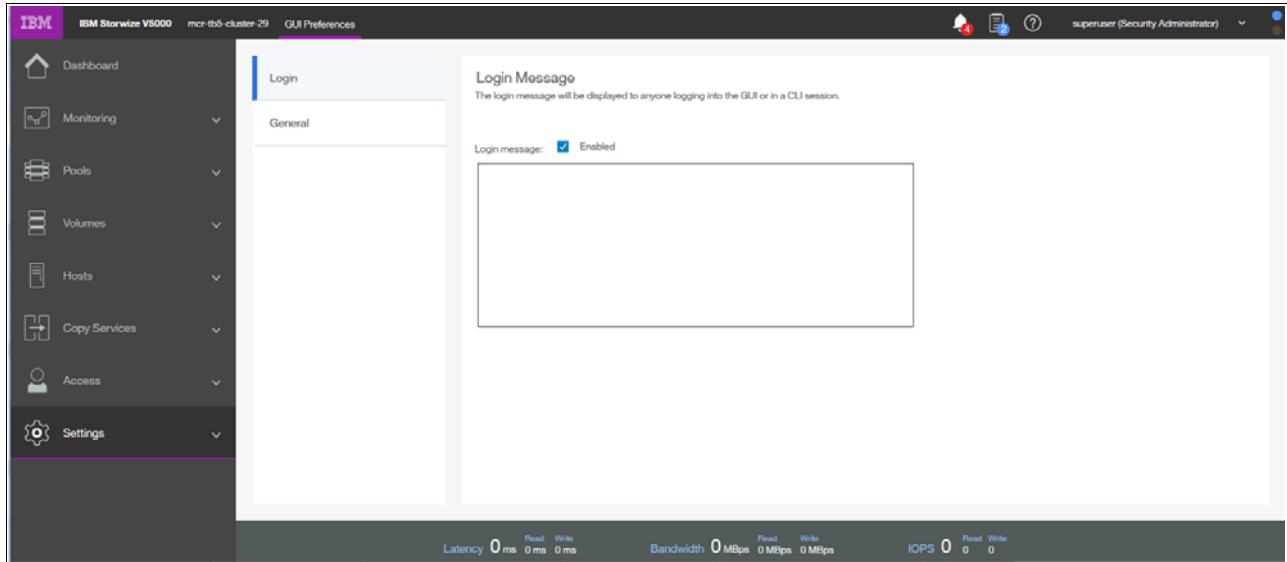


Figure 3-100 GUI preferences

Select **General** to adjust the browser settings as shown in Figure 3-101.

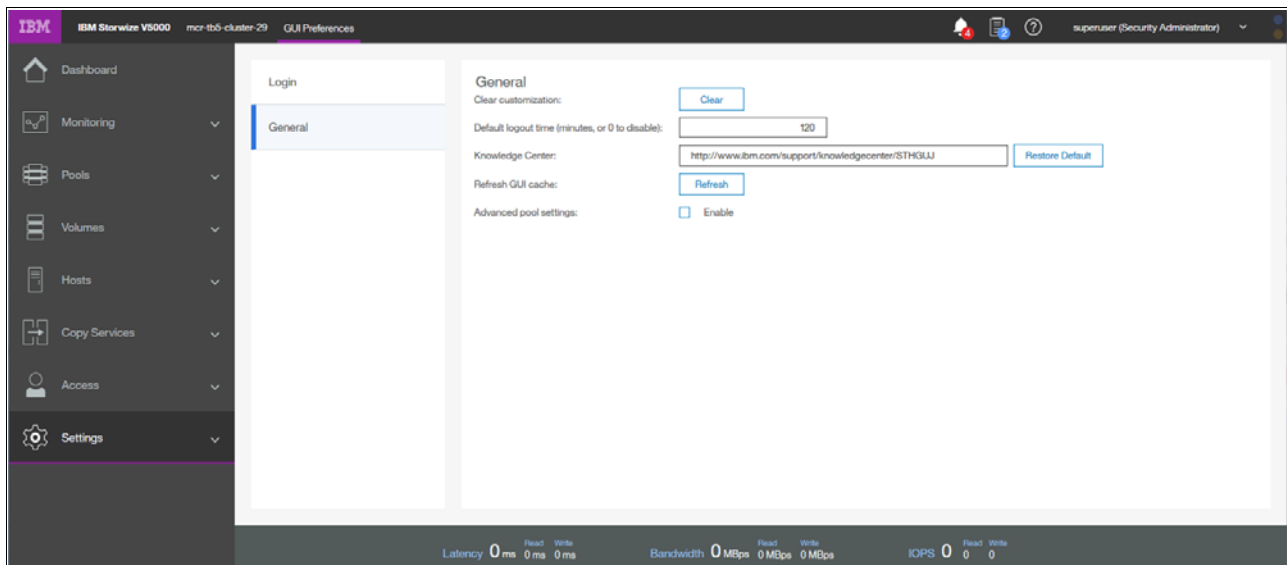


Figure 3-101 General settings

Storage pools

This chapter describes how the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 manage physical storage resources. All storage resources that are under Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 control are managed by using *storage pools*. Storage pools facilitates to dynamically allocate resources, maximize productivity and reduce costs. Internal and external managed disks (MDisks), advanced internal storage, and storage pool management are covered in this chapter. External storage controllers are covered in Chapter 11, “External storage virtualization” on page 607.

Storage pools can be configured through the Initial Setup wizard when the system is first installed, as described in Chapter 2, “Initial configuration” on page 35. They can also be configured after the initial setup through the management GUI, which provides a set of presets to help you configuring different Redundant Array of Independent Disks (RAID) types.

The recommended configuration presets configure all drives into RAID arrays based on drive class and protect them with the correct number of spare drives. Alternatively, you can configure the storage to your own requirements. Selections include the drive class, the number of drives to configure, whether to configure spare drives and optimization for performance or capacity.

Specifically, this chapter provides information about the following topics:

- ▶ 4.1, “Working with internal drives” on page 140
- ▶ 4.2, “Working with storage pools” on page 150
- ▶ 4.3, “Working with managed disks” on page 162
- ▶ 4.4, “Working with external storage controllers” on page 186

4.1 Working with internal drives

This section describes how to configure the internal storage disk drives by using different RAID levels and optimization strategies. For more information about RAID settings see section 4.3.2, “RAID configuration” on page 169.

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems provide an Internal Storage window for managing all internal drives. The Internal Storage window can be accessed by opening the System window, clicking the **Pools** option and then clicking **Internal Storage**, as shown in Figure 4-1.

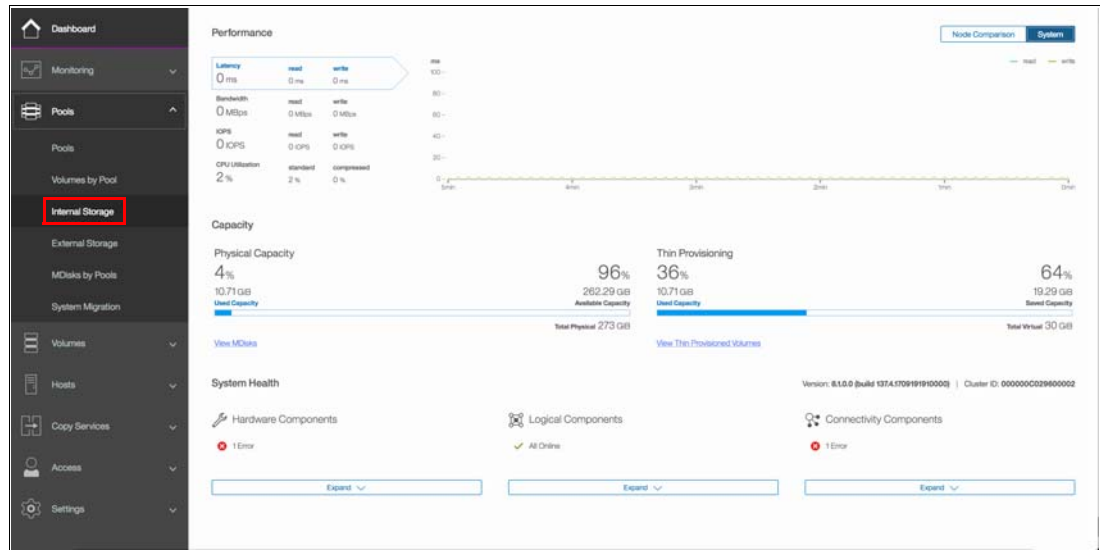


Figure 4-1 Path to Internal Storage window

4.1.1 Internal Storage window

The Internal Storage window (as shown in Figure 4-2 on page 141) provides an overview of the internal drives that are installed in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems. Selecting **All Internal** under the Drive Class Filter shows all the drives that are installed in the managed system, including attached expansion enclosures. Alternatively, you can filter the drives by their type or class. For example, you can choose to show only Enterprise drive class (serial-attached Small Computer System Interface (SCSI) or (SAS)), Nearline SAS, or Flash drives.

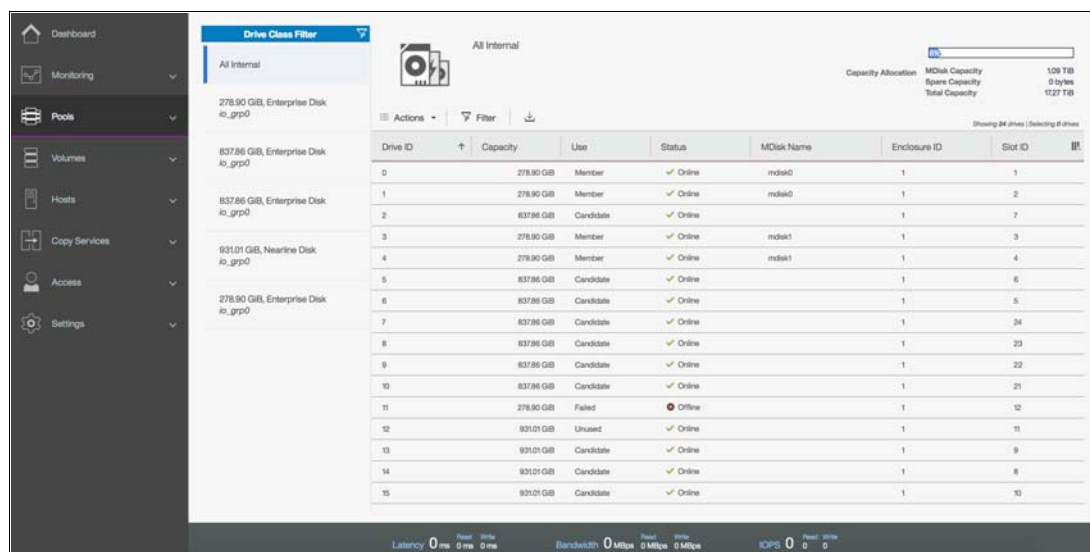


Figure 4-2 Internal Storage window

The right side of the Internal Storage window lists the selected type of internal disk drives. By default, the following information is listed:

- ▶ Logical drive ID
- ▶ Drive capacity
- ▶ Current type of use (unused, candidate, member, spare, or failed)
- ▶ Status (online, offline, and degraded)
- ▶ MDisk name that the drive is a member of
- ▶ Enclosure ID that the drive is installed in
- ▶ Slot ID of the enclosure in which the drive is installed

The default sort order is by enclosure ID. This default can be changed to any other column by left-clicking the column header. To toggle between ascending and descending sort order, left-click the column header again. By hovering over the header names such as Drive ID, it you display a brief description of the items within that column.

Additional columns can be included by right-clicking the gray header bar of the table, which opens the selection panel, as shown in Figure 4-3 on page 142. To restore the default column options, select **Restore Default View**.

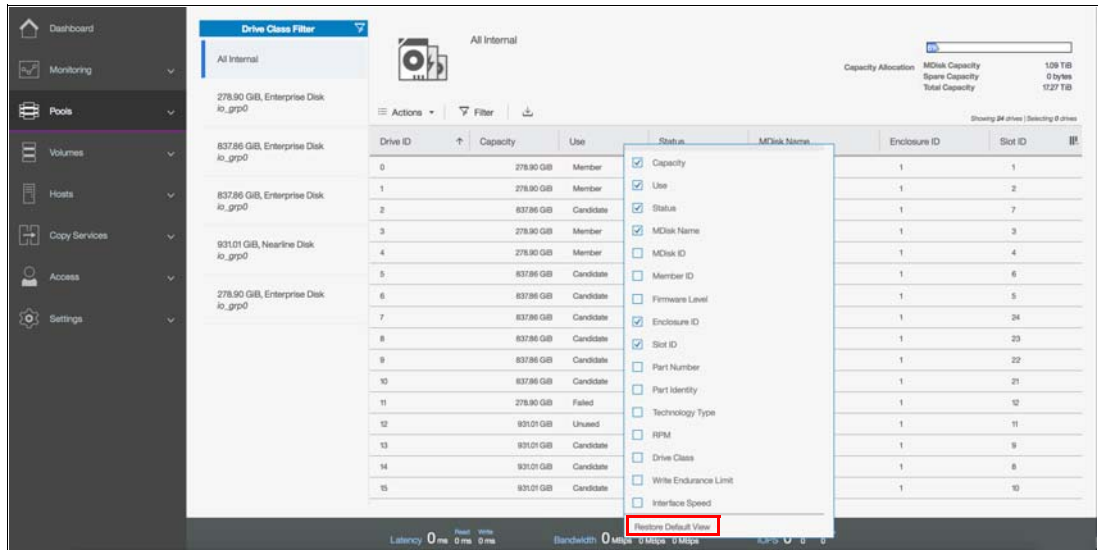


Figure 4-3 Additional column options for Internal Storage window

The overall internal storage capacity allocation indicator is shown in the upper-right corner. The *Total Capacity* shows the overall capacity of the internal storage that is installed in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems. The *MDisk Capacity* shows the internal storage capacity that is assigned to the MDisks. The *Spare Capacity* shows the internal storage capacity that is used for hot spare disks.

The percentage bar that is shown in Figure 4-4 indicates how much capacity is allocated.

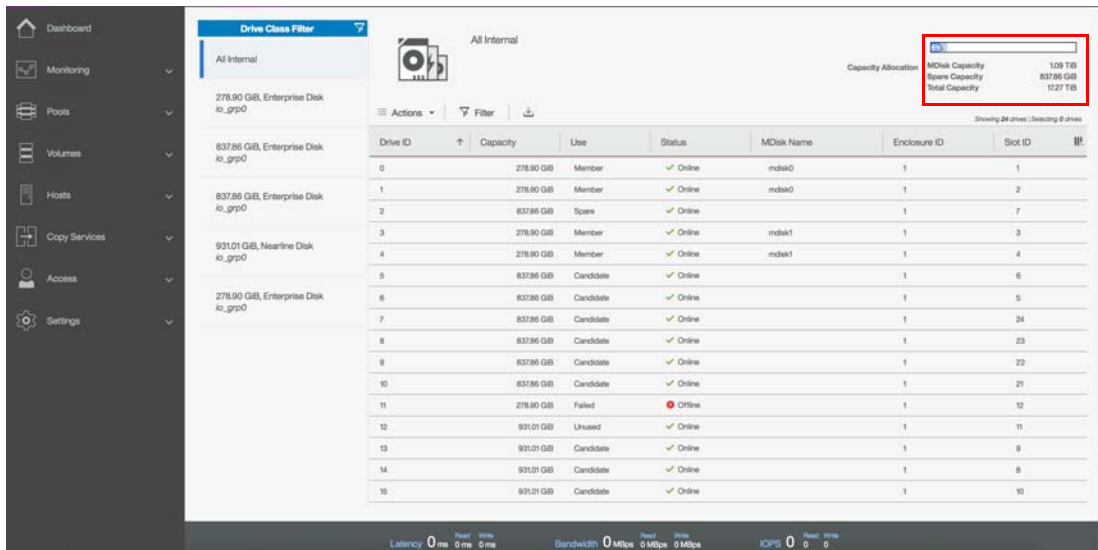


Figure 4-4 Internal Storage allocation indicator

4.1.2 Actions on internal drives

You can perform several actions by right-clicking the internal drives or clicking on the **Actions** drop-down menu, as shown in Figure 4-5 on page 143. If you click Actions without selecting any drive the only option available will be Upgrade All.

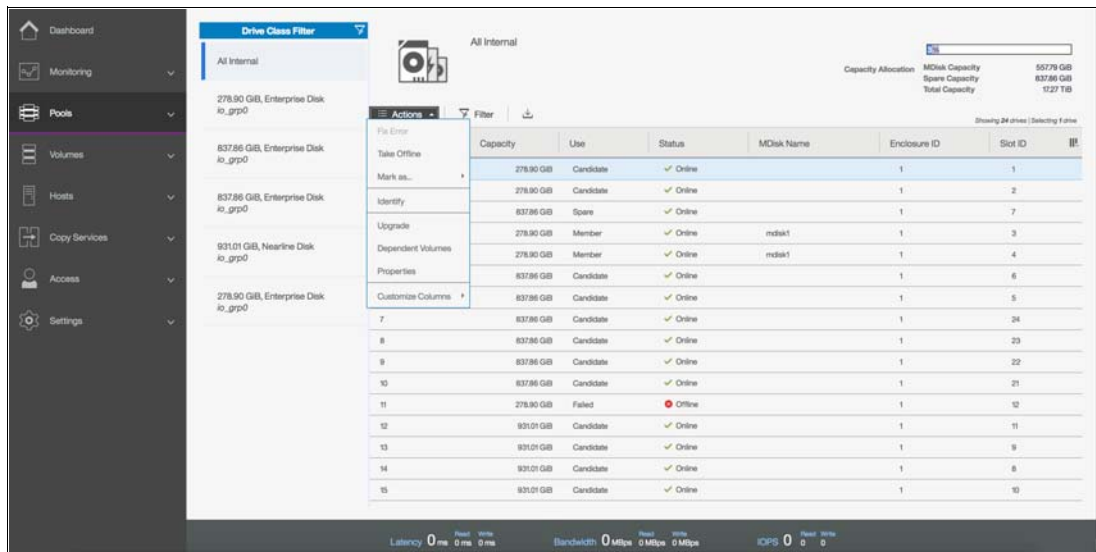


Figure 4-5 Internal drive actions menu

Depending on the status of the selected drive, the following actions are available.

Take Offline

The internal drives can be taken offline if a problem on the drive is identified. A confirmation window opens, as shown in Figure 4-6. The default selection is to only take a drive offline if a spare drive is available, which is strongly recommended and avoids redundancy loss in the *array*. Click **OK** to take the drive offline.

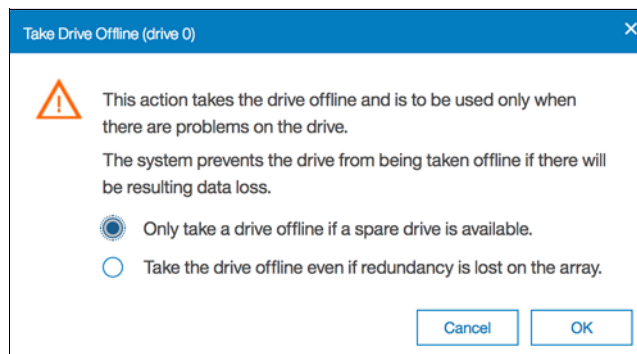


Figure 4-6 Warning before taking offline an internal drive

If the drive fails (as shown in Figure 4-7 on page 144), the MDisk (from which the failed drive is a member) remains online and a hot spare is automatically reassigned.

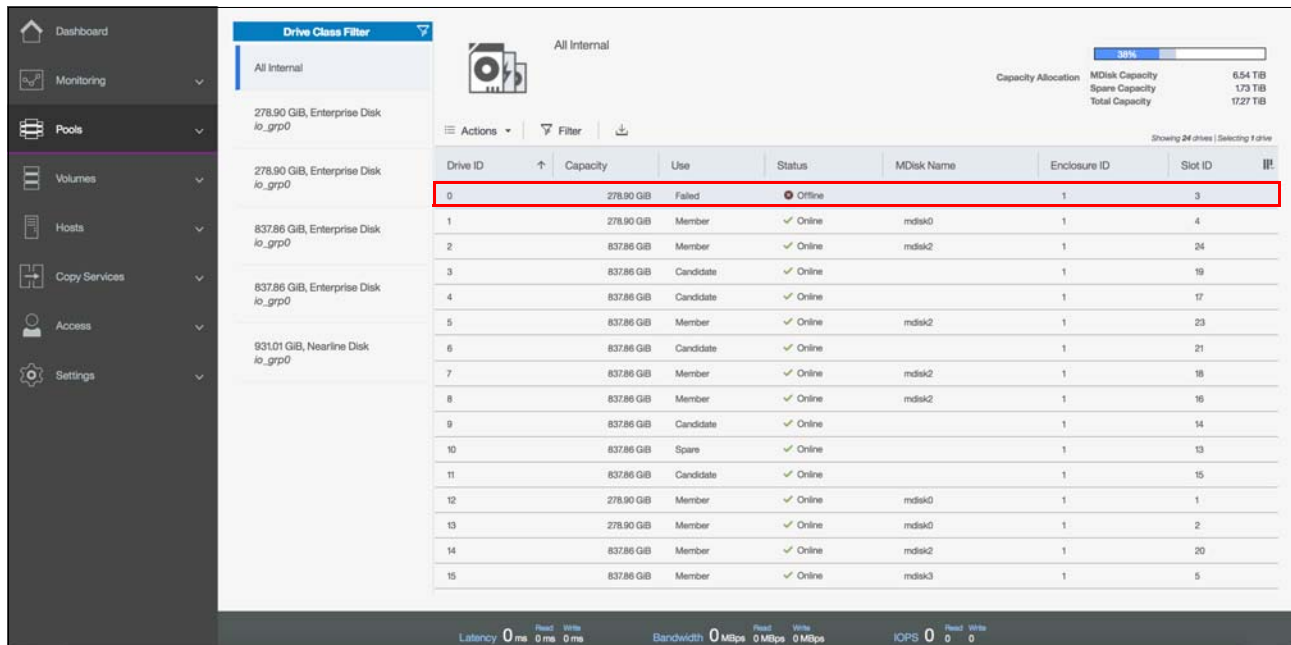


Figure 4-7 Internal drive taken offline

If sufficient spare drives are not available and a drive must be taken offline, the second option for no redundancy (Take the drive offline even if redundancy is lost on the array) must be selected. This option results in a degraded storage pool due to the degraded MDisk, as shown in Figure 4-8.



Figure 4-8 Degraded MDisk

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems prevent the drive from being taken offline if it can result in data loss. A drive cannot be taken offline (as shown in Figure 4-9) if no suitable spare drives are available and based on the RAID level of the MDisk, no sufficient redundancy will be available. Click **Close** to return to the Internal Storage panel.

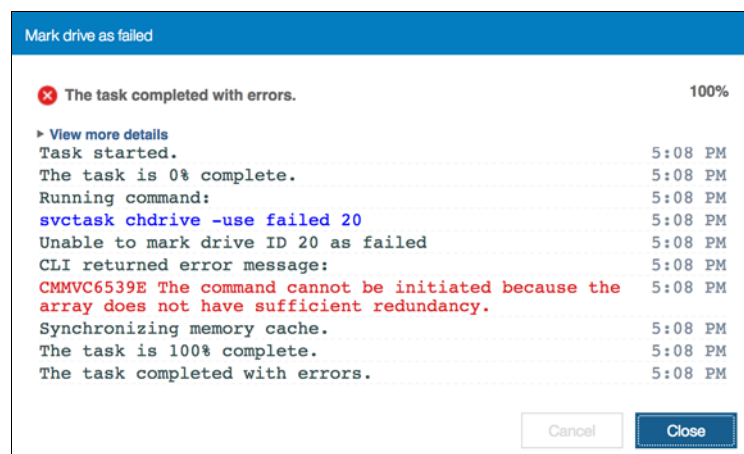


Figure 4-9 Internal drive offline not allowed because of insufficient redundancy

Example 4-1 shows how to use the **chdrive** command-line interface (CLI) command to set the drive to failed.

Example 4-1 The use of the chdrive command to set the drive to failed

```
chdrive -use failed driveID
chdrive -use failed -allowdegraded driveID
```

Mark as

The internal drives in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems can be assigned to the following usage roles by right-clicking the drives and selecting the **Mark as** option, as shown in Figure 4-10:

- ▶ Unused: The drive is not in use, and it cannot be used as a spare.
- ▶ Candidate: The drive is available for use in an array.
- ▶ Spare: The drive can be used as a hot spare, if required.

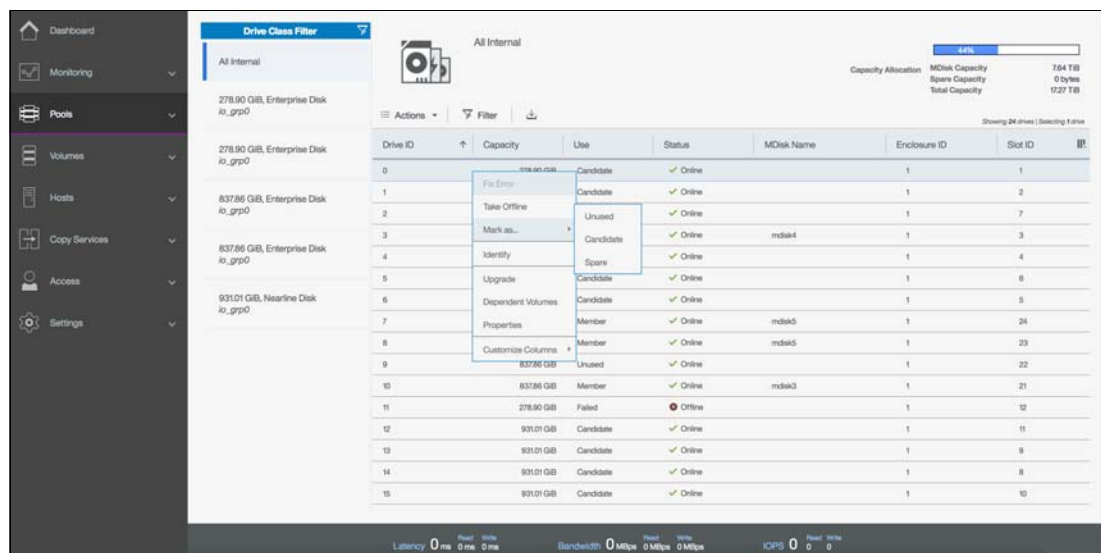


Figure 4-10 Selecting the internal drive “Mark as” action

Defining a new role to a drive depends on the current drive usage role. These dependencies are shown in Figure 4-11.

		To					
		Unused	Candidate	Failed	Member	Spare	
From	Unused	allowed	allowed	no option		not allowed	
	Candidate	allowed	allowed			allowed	
	Failed	allowed	allowed			not allowed	
	Member	No change on member drives					
	Spare	not allowed	allowed	no option		allowed	

Figure 4-11 Internal drive usage role table

Identify

Use the Identify action to turn on the LED light so that you can easily identify a drive that must be replaced or that you want to physically troubleshoot. The panel that is shown in Figure 4-12 appears when the LED is on. Click **Turn LED off** when you are finished to turn the drive LED off and return to the Internal Storage panel.




Figure 4-12 Internal drive identification

Example 4-2 shows how to use the **chenclosureslot** command to turn on and turn off the drive LED.

Example 4-2 The use of the *chenclosureslot* command to turn on and turn off the drive LED

```
chenclosureslot -identify yes/no -slot slot enclosureID
```

Upgrade

From this option, you can easily upgrade the drive firmware. You can use the GUI to upgrade individual drives or upgrade all drives for which updates are available. For more information about upgrading drive firmware, see 12.4.2, “Updating the drive firmware” on page 663 and this web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/lenovo_vseries.html

Dependent Volumes

Clicking **Dependent Volumes** shows the volumes that depend on the selected drive. Volumes depend on a drive only when their underlying MDisk are in a degraded or inaccessible state and when the removal of more hardware causes the volume to go offline. This condition is true for any RAID 0 MDisk because it has no redundancy, or if the associated MDisk is already degraded.

Use the Dependent Volumes option before you perform any drive maintenance to determine which volumes are affected.

Important: A lack of listed dependent volumes does not imply that no volumes are using the drive.

Figure 4-13 shows an example if no dependent volumes are detected for a specific drive. If a dependent volume is identified it will be listed within this panel. When you have volumes listed as dependent you can also check volume saving and throttle by selecting the volume and clicking **Actions**. Click **Close** to return to the Internal Storage panel.

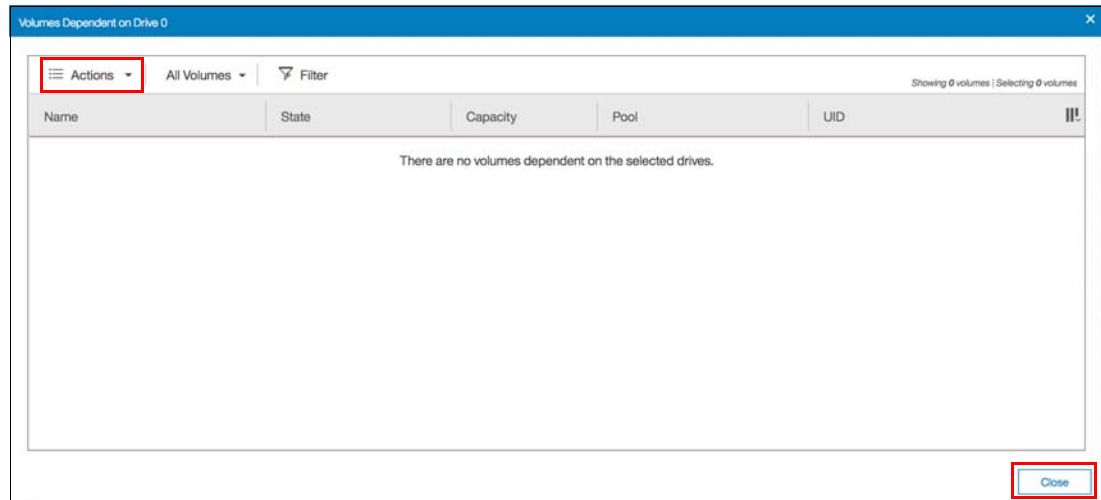


Figure 4-13 Internal drive with no dependent volumes

Example 4-3 shows how to view dependent volumes for a specific drive by using the CLI.

Example 4-3 Command to view dependent virtual disks (VDisks) for a specific drive

```
lsdependentvdisks -drive driveID
```

Properties

Clicking **Properties** in the Actions menu or double-clicking the drive provides the vital product data (VPD) and the configuration information, as shown in Figure 4-14 on page 148. The **Show Details** option in the bottom-left of the Properties panel was selected to show more information.

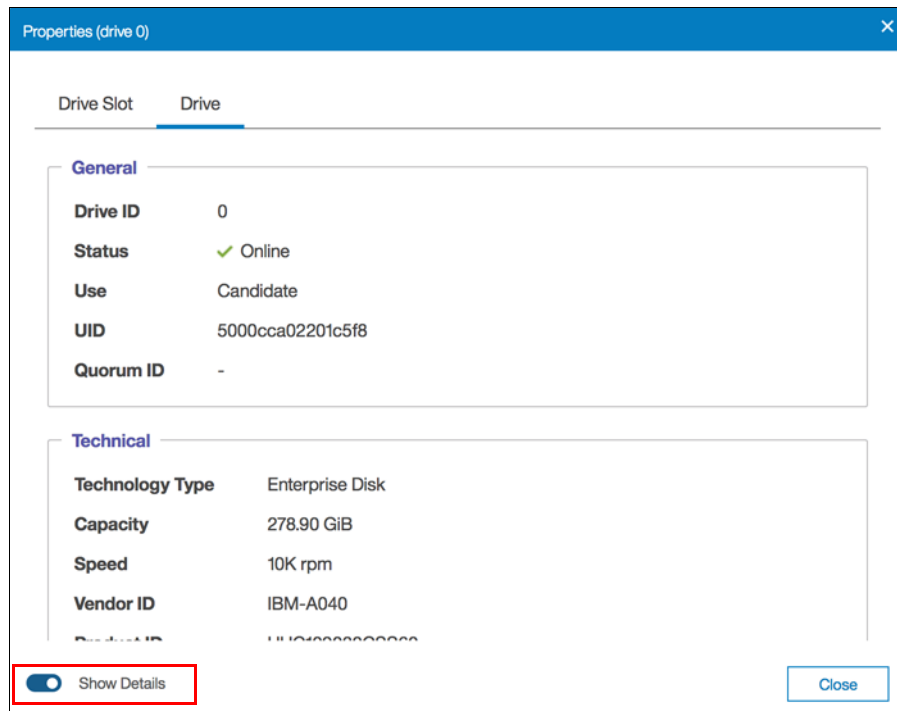


Figure 4-14 Detailed internal drive properties

If the Show Details option is not selected, the technical information section is reduced, as shown in Figure 4-15.

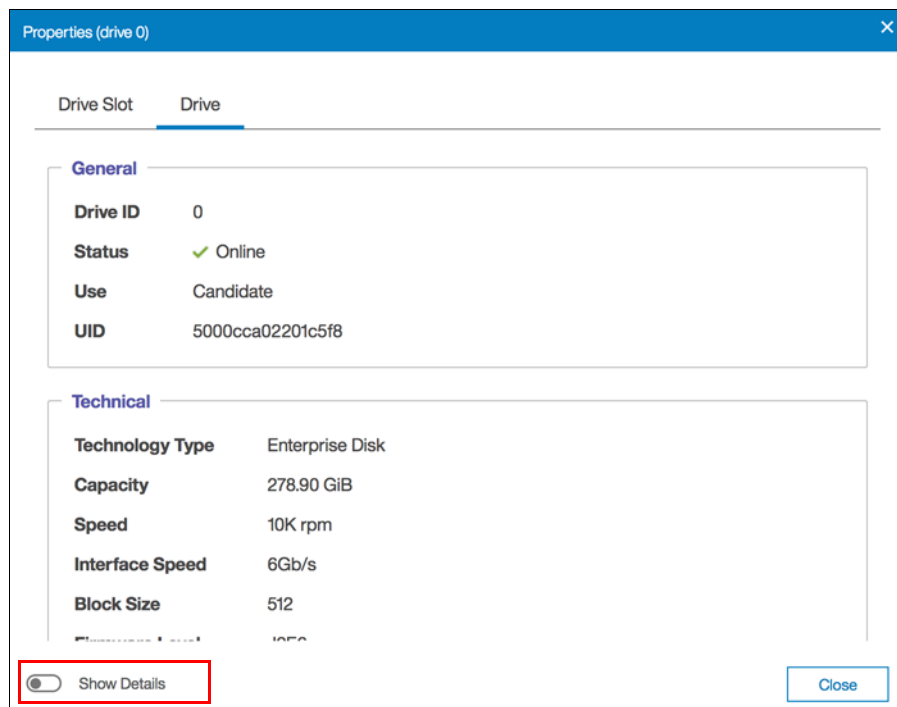


Figure 4-15 Internal drive properties without details

A tab for the Drive Slot is available in the Properties panel (as shown in Figure 4-16 on page 149) to obtain specific information about the slot of the selected drive. The Show Details

option is also applicable to this tab, if you do not select it, the Fault LED information disappears from the panel. Click **Close** to return to the Internal Storage panel.

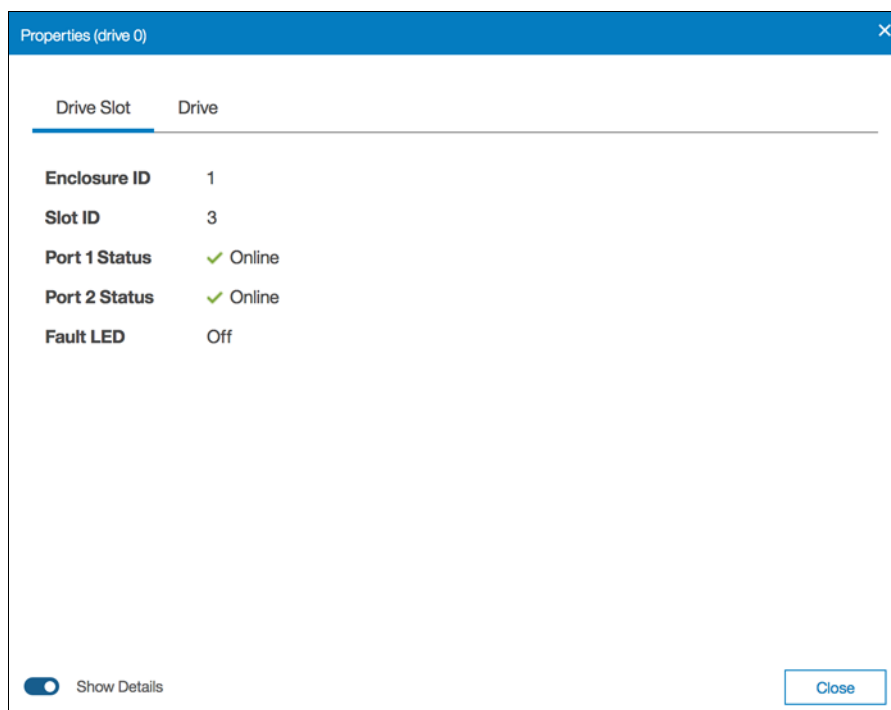


Figure 4-16 Internal drive properties slot

Example 4-4 shows how to use the **lsdrive** command to display the configuration information and drive VPD.

Example 4-4 The use of the lsdrive command to display configuration information and drive VPD

```
IBM_Storwize:ITS0 V5000:superuser>lsdrive 1
id 1
status online
error_sequence_number
use member
UID 5000cca05b1d97b0
tech_type sas_hdd
capacity 278.9GB
block_size 512
vendor_id IBM-E050
product_id HUC156030CSS20
FRU_part_number 01AC594
FRU_identity 11S00D5385YXXX0TGJ8J4P
RPM 15000
firmware_level J2G9
FPGA_level
mdisk_id 0
mdisk_name MDisk_01
member_id 5
enclosure_id 1
slot_id 1
node_id
node_name
```

```

quorum_id
port_1_status online
port_2_status online
interface_speed 12Gb
protection_enabled yes
auto_manage inactive
drive_class_id 145
IBM_Storwize:ITS0 V5000:superuser>

```

Customize Columns

Click **Customize Columns** in the Actions menu to add or remove several columns that are available in the Internal Storage window.

To restore the default column options, select **Restore Default View**, as shown in Figure 4-17.

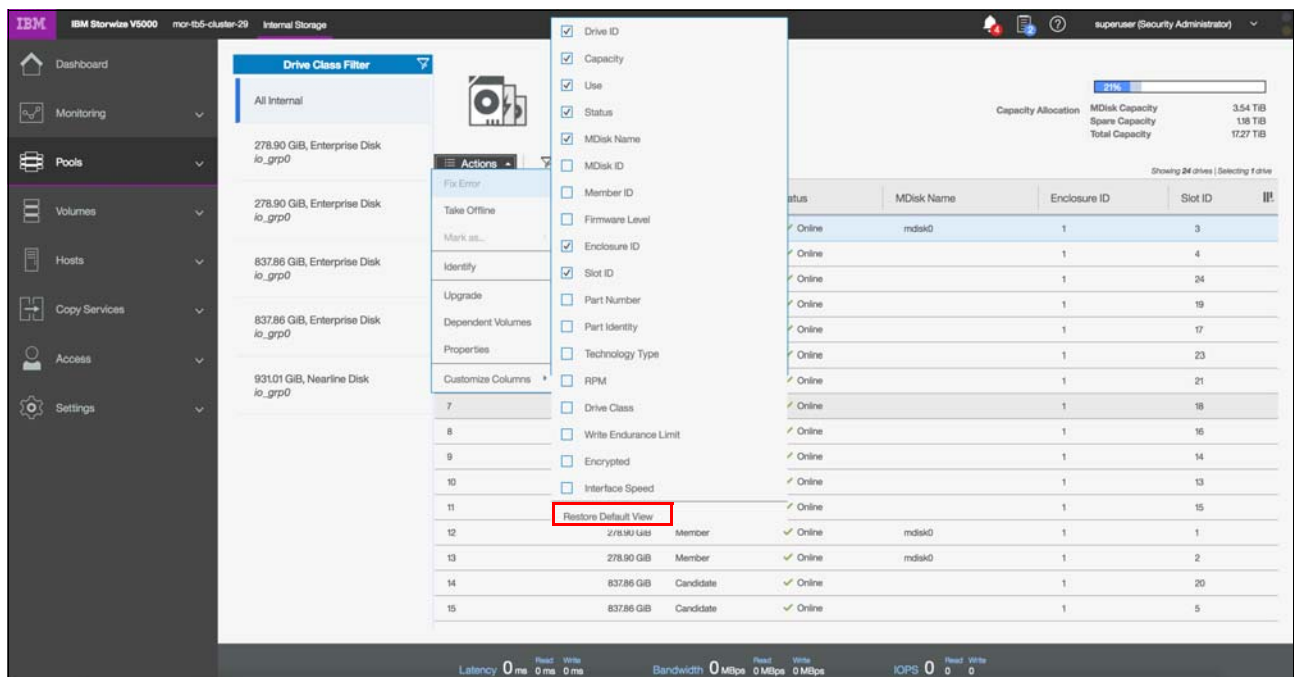


Figure 4-17 Customizing columns on the Internal Storage window

4.2 Working with storage pools

Storage pools (or *pools*) act as containers for MDisks and provision the capacity to volumes. MDisks can be provisioned through internal or external storage. MDisks created from internal storage are created as RAID arrays.

Figure 4-18 on page 151 provides an overview of how storage pools, MDisks, and volumes are related. The numbers in the figure represents the following components:

- ▶ Hosts (1)
- ▶ Volumes (5)
- ▶ Pools (4)
- ▶ External MDisks (0)
- ▶ Arrays (2)

This panel is available by browsing to **Monitoring** → **System** and clicking **Overview** on the upper-right corner of the panel. You can also identify the name of each resource by hovering over the elements on the Overview window.

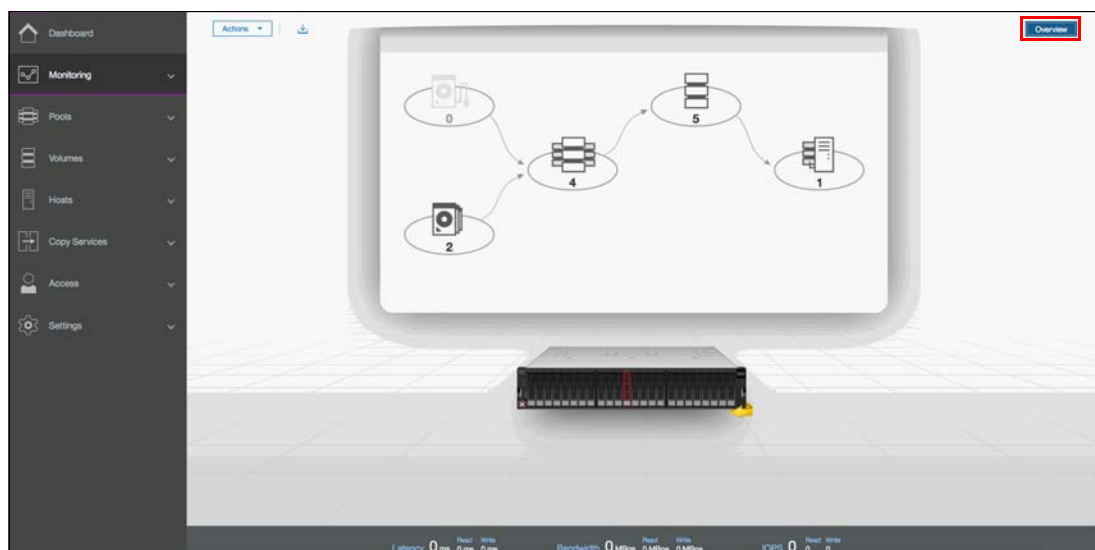


Figure 4-18 System Overview panel

Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 organize storage into pools to ease storage management and make it more efficient. All MDisks in a pool are split into extents of the same size and volumes are created from these available extents. The extent size is a property of the storage pool and when an MDisk is added to a pool the size of the extents that composes it will be based on the attribute of the pool to which the MDisk was added.

Storage pools can be further divided into sub-containers named as *child pools*. Child pools inherit the properties of the parent pool and can also be used to provision volumes.

Storage pools are managed either through the Pools panel or the MDisks by Pool panel. Both panels allow you to execute the same actions; however, actions on child pools can be performed only through the Pools panel. To access the Pools panel browse to **Pools** → **Pools**, as shown in Figure 4-19.




Figure 4-19 Pools panel

The panel lists all storage pools available in the system. If a storage pool has child pools, you can toggle the sign to the left of the storage pool icon to either show or hide the child pools.

4.2.1 Creating storage pools

If you are installing a brand new Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, no pools are created when you first login, so the system automatically suggests a pool creation, which leads directly to the Create Pool panel. You can access the Pools panel in the future through the Pools menu as previously shown in Figure 4-19 on page 151.

To create a new storage pool, you can use one of the following alternatives:

- Navigate to **Pools** → **Pools** and click **Create**, as shown in Figure 4-20.




Figure 4-20 Create button on Pools panel

- Navigate to **Pools** → **MDisks by Pools** and click **Create Pool**, as shown in Figure 4-21.

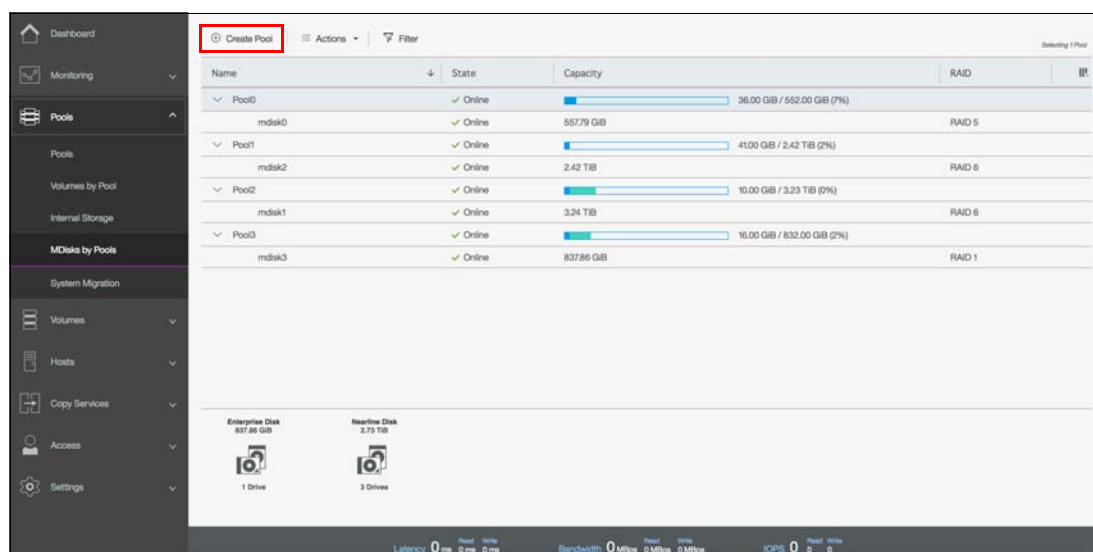


Figure 4-21 Create Pool button on MDisks by Pools panel

Both the alternatives open the dialog box shown in Figure 4-22 on page 153.

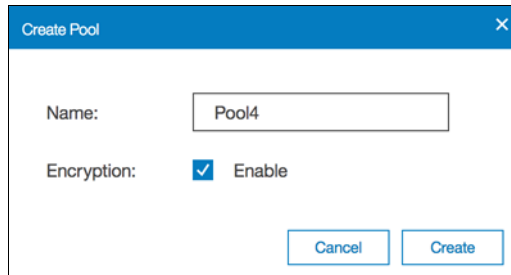


Figure 4-22 Create Pool dialog box

Note: If encryption is enabled, you can additionally select whether the storage pool is encrypted. The encryption setting of a storage pool is selected at creation time and cannot be changed later. By default, if encryption is enabled, encryption is selected.

If advanced pool settings are enabled, you can additionally select an extent size at the time of the pool creation, as shown in Figure 4-23.

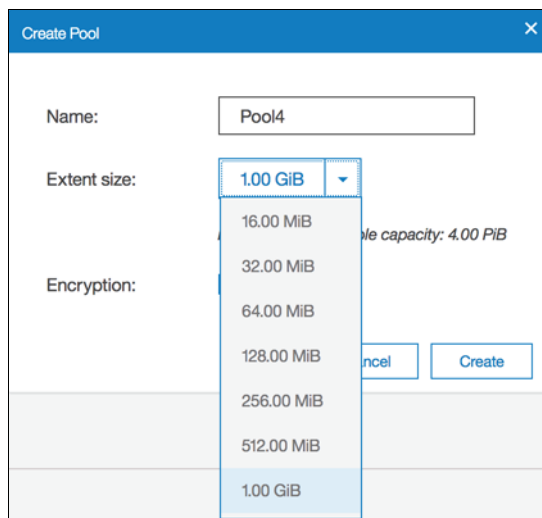


Figure 4-23 Creating pool with advanced pool settings enabled

Note: Every storage pool created through the GUI has a default extent size of 1 GB. The size of the extent is selected at creation time and cannot be changed later. If you want to specify a different extent size at the time of the pool creation, browse to **Settings** → **GUI Preferences** and select **Advanced pool settings**.

In the Create Pool dialog box, enter the pool name and click **Create**. The new pool is created and is included in the list of storage pools with zero bytes, as shown in Figure 4-24 on page 154.

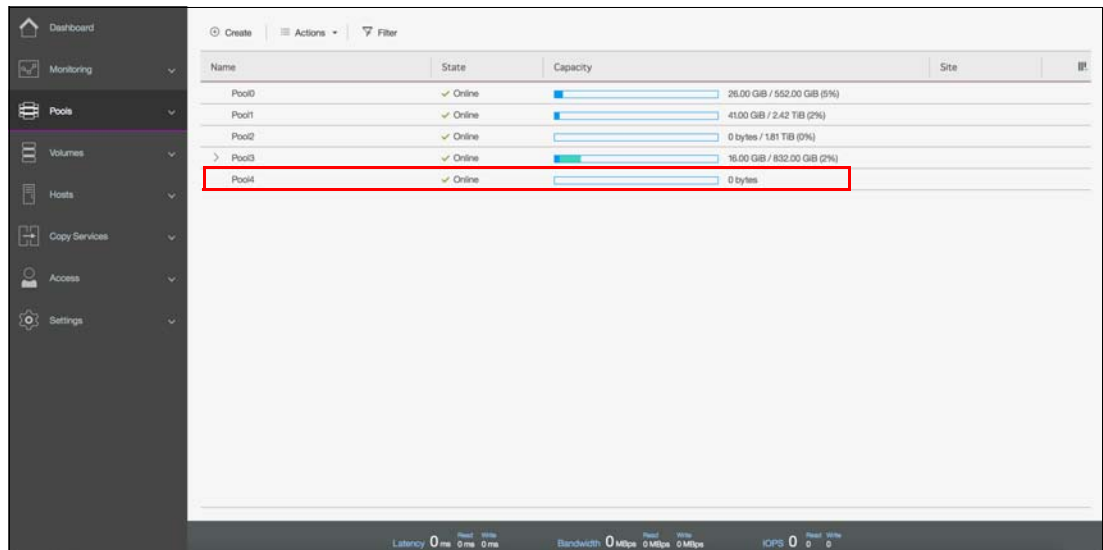


Figure 4-24 New pool with zero bytes included in the list

4.2.2 Actions on storage pools

There are several actions that can be performed on storage pools, which can be accessed through the Pools panel or the Mdisks by Pools panel. To select an action, select the storage pool and click **Actions**. Alternatively, right-click the storage pool.

Figure 4-25 shows the list of available actions for storage pools being accessed through the Pools panel.

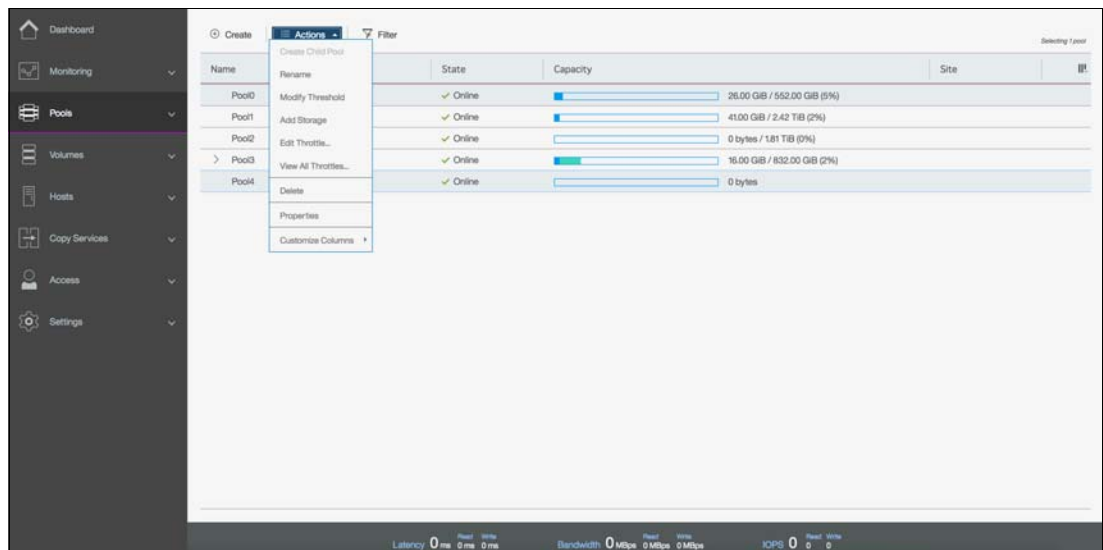


Figure 4-25 Actions list for storage pools

Create child pool

Selecting **Create Child Pool** starts the wizard to create a child storage pool. For information about child storage pools and a detailed description of this wizard, see 4.2.3, “Child storage pools” on page 159

Rename

Selecting **Rename** at anytime allows you to modify the name of a storage pool, as shown in Figure 4-26. Enter the new name and click **Rename**.




Figure 4-26 Renaming pools

Modify threshold

The storage pool threshold refers to the percentage of storage capacity that must be in use for a warning event to be generated. The threshold is especially useful when using thin-provisioned volumes that are configured to expand automatically. The threshold can be modified by selecting **Modify Threshold** and entering the new value, as shown in Figure 4-27. The default threshold is 80%. Warnings can be disabled by setting the threshold to 0%.




Figure 4-27 Modifying pool threshold

Add storage

Selecting **Add Storage** starts the wizard to assign storage to the pool. For a detailed description of this wizard, see 4.3.1, “Assigning managed disks to storage pools” on page 163.

Edit Throttle

You can create, modify, and remove throttles for pools by using the management GUI or the command-line interface. *Throttling* is a mechanism to control the amount of resources that are used when the system is processing I/Os on a specific pool. If a throttle is defined, the system either processes the I/O, or delays the processing of the I/O to free resources for more critical I/O.

There are two parameters that can be defined through the Edit Throttle option:

- ▶ Bandwidth limit defines the maximum amount of bandwidth the pool can process before the system delays I/O processing for this pool.
- ▶ IOPS limit defines the maximum I/O operations per second the pool can process before the system delays I/O processing for this pool

If the pool does not have throttle settings configured, selecting Edit Throttle displays a dialog box with blank fields as shown in Figure 4-28 on page 156. Define the limits and click **Create**.

Figure 4-28 Edit throttle initial configuration

For a pool that already has defined throttle settings, selecting Edit Throttle displays a different dialog box, in which the current bandwidth and IOPS limits will be displayed, as shown in Figure 4-29. You can either change or remove the current bandwidth and IOPS limits by modifying the values and clicking **Save** or clicking **Remove** to disable a limitation.

Figure 4-29 Editing throttles

View all throttles

Selecting **View All Throttles** opens a panel as shown in Figure 4-30 which displays the current throttle information, which includes the limits that were previously applied for bandwidth and IOPS.

Object Name	Status	Throttle Type	Bandwidth Limit	IOPS
ITSO_Pool	Online	Pool	100.00 MBps	100

Figure 4-30 View All Throttles panel

As a default, when the View All Throttles panel is opened through the Pools window, it displays throttle information related to pools, but through the same panel you are allowed to select different objects as shown in Figure 4-31 on page 157. Selecting a different category displays the throttle information for that specific selection.

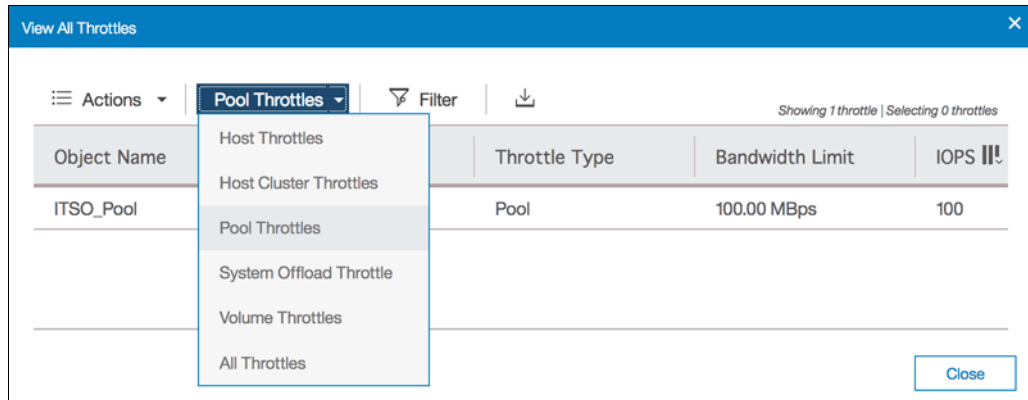


Figure 4-31 Selecting specific throttle information

Delete

Pools can only be deleted through the GUI if no volumes are assigned to the pool. If the pool has any volumes within it the option is not available. Selecting **Delete** immediately deletes the pool without additional confirmation.

Through the CLI, you can delete a pool and all of its contents by using the **-force** parameter. However, all volumes and host mappings are deleted and you cannot recover them.

Important: After you delete the pool through the CLI, all data that is stored in the pool is lost except for the image mode MDisk. The image mode MDisk volume definition is deleted, but the data on the imported MDisk remains untouched.

After deleting a pool, all of the managed or image mode MDisk in the pool return to the unmanaged status.

Properties

Selecting **Properties** displays information about the storage pool as shown in Figure 4-32.

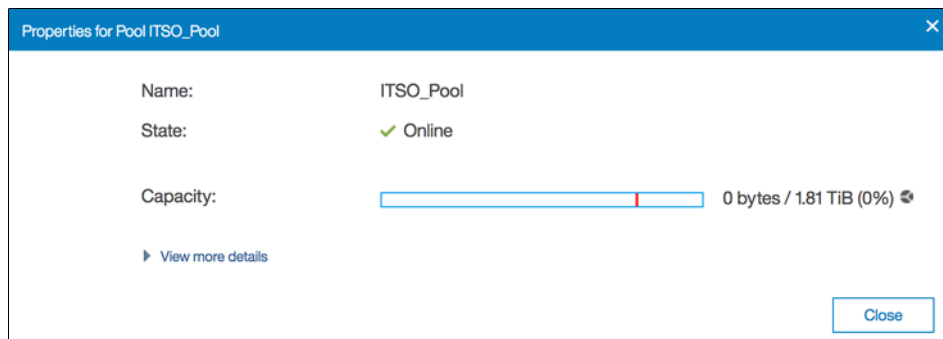


Figure 4-32 Storage pool properties

Additional information is available by clicking **View more details** and by hovering over the elements on the window, as shown in Figure 4-33 on page 158. Click **Close** to return to the Pools panel.

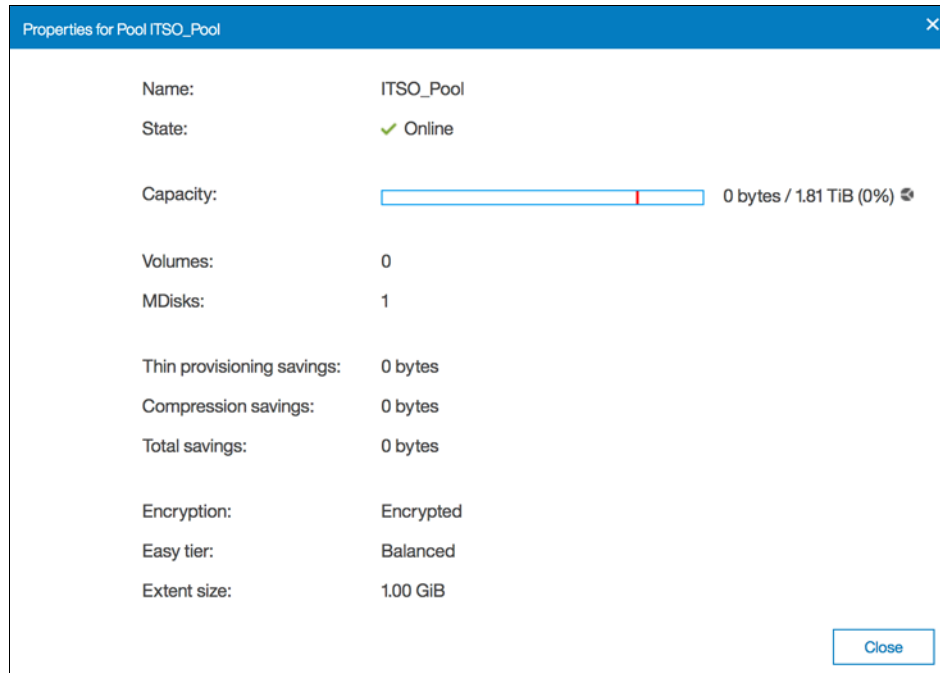


Figure 4-33 Additional details for storage pool properties

Customize columns

Selecting **Customize Columns** in the Actions menu allows you to include additional information fields in the Pools panel as shown in Figure 4-34.

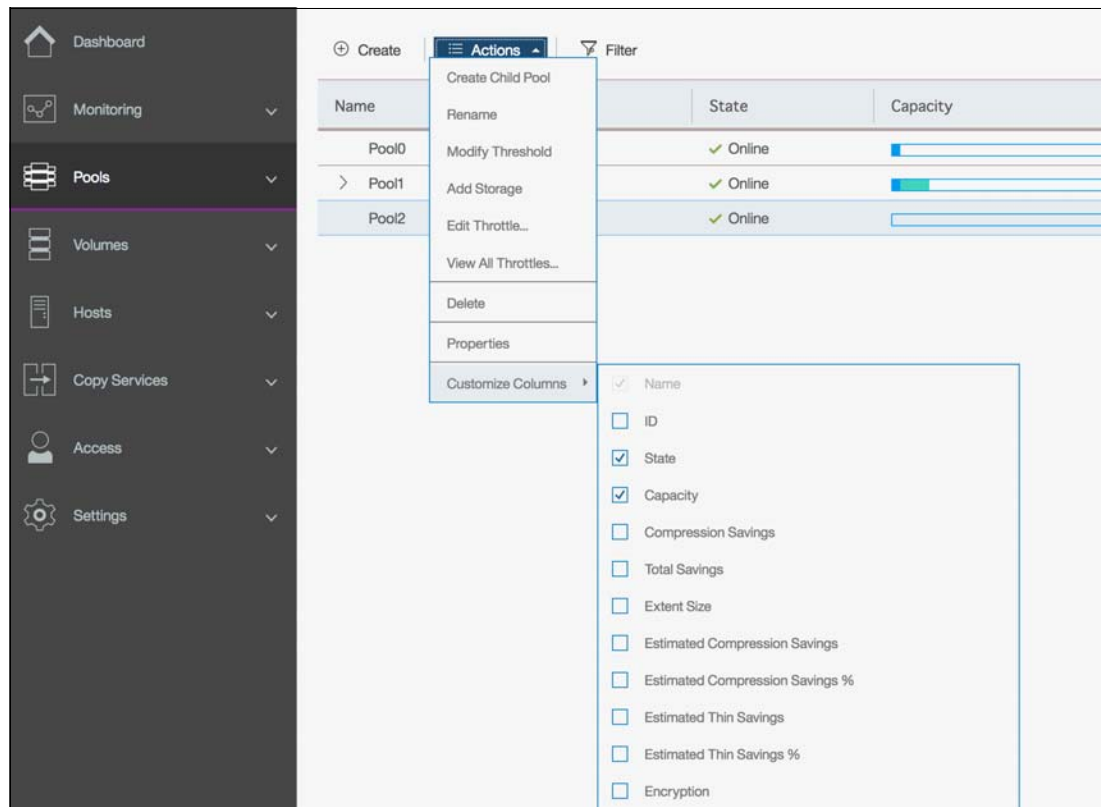


Figure 4-34 Customizing columns in the Pools panel

4.2.3 Child storage pools

A *child storage pool* is a storage pool created within a storage pool. The storage pool in which the child storage pool is created is called *parent storage pool*.

Unlike a parent pool, a child pool does not contain MDisks; its capacity is provided exclusively by the parent pool in the form of extents. The capacity of a child pool is set at creation time, but can be nondisruptively modified later. The capacity must be a multiple of the parent pool extent size and must be smaller than the free capacity of the parent pool.

Child pools are useful when the capacity allocated to a specific set of volumes must be controlled.

Child pools inherit most properties from their parent pools and these cannot be changed. The inherited properties include:

- ▶ Extent size
- ▶ Easy Tier setting
- ▶ Encryption setting, but only if the parent pool is encrypted

Creating a child pool

To create a child pool you can either browse to **Pools** → **Pools** → **Actions** or **Pools** → **MDisks by Pools** → **Actions** and select **Create Child Pool**. Alternatively, you can right-click the parent pool as shown in Figure 4-35.

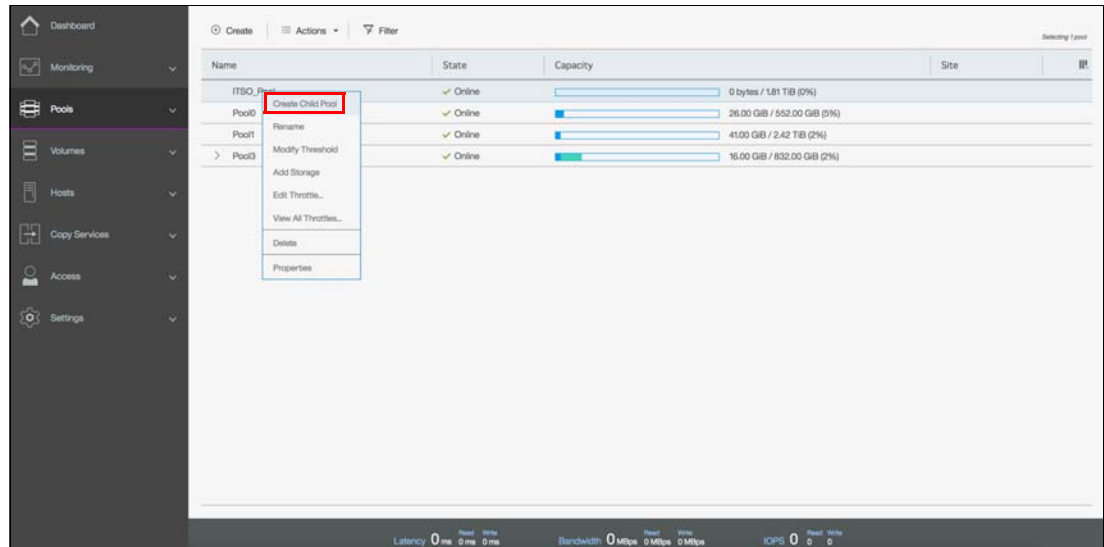
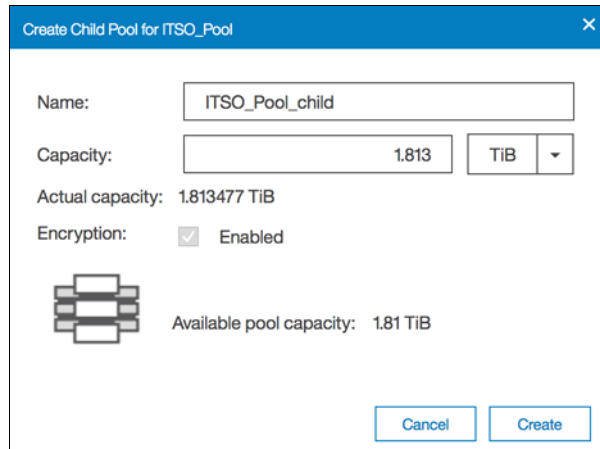


Figure 4-35 Selecting child menu creation

Enter the name and the capacity of the child pool and click **Create**, as shown in Figure 4-36 on page 160.




Create Child Pool for ITSO_Pool

Name:

Capacity: TiB

Actual capacity: 1.813477 TiB

Encryption: ☒ Enabled

 Available pool capacity: 1.81 TiB

Cancel Create

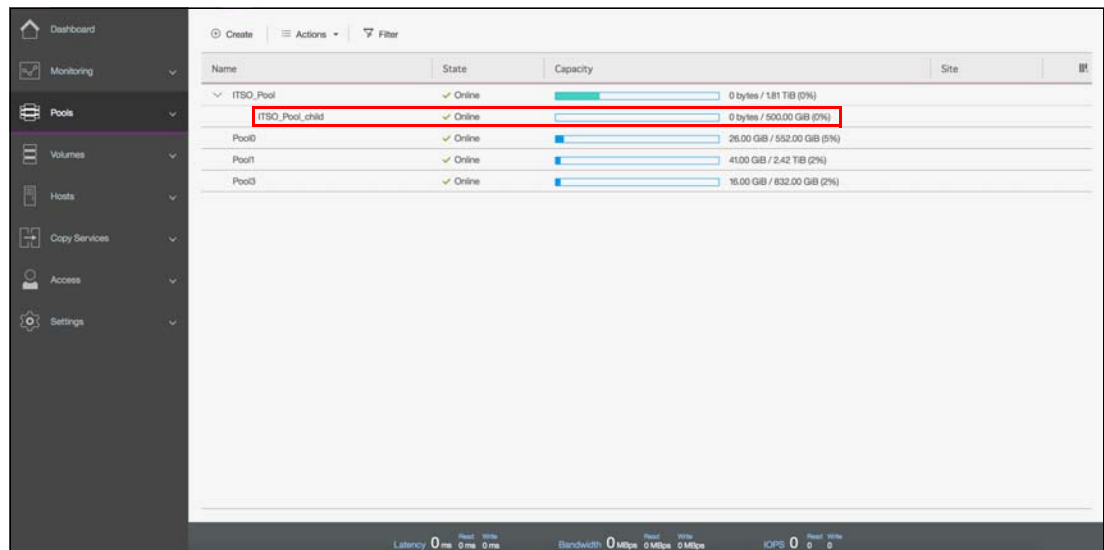
Figure 4-36 Create Child Pool panel

Note: You cannot create an encrypted child pool from an unencrypted parent pool if the parent pool contains any unencrypted array or an MDisk that is not self-encrypting and there are nodes in the system that do not support software encryption (e.g. do not have encryption license enabled).

An encrypted child pool created from an unencrypted parent pool reports as unencrypted if the parent pool contains any unencrypted arrays. Remove these arrays to ensure that the child pool is fully encrypted.

After the child pool is created it is listed in the Pools panel under its parent pool, as shown in Figure 4-37. Toggle the arrow sign in the left of the storage pool name to either show or hide the child pools.

Creating a child pool within a child pool is not possible.



Name	State	Capacity	Site
ITSO_Pool	Online	0 bytes / 1.81 TiB (0%)	
ITSO_Pool_child	Online	0 bytes / 500.00 GiB (0%)	
Pool0	Online	26.00 GiB / 552.00 GiB (5%)	
Pool1	Online	41.00 GiB / 2.42 TiB (2%)	
Pool3	Online	16.00 GiB / 832.00 GiB (2%)	

Latency 0 ms Read 0 ms Write 0 ms Bandwidth 0 MBps Read 0 MBps Write 0 MBps IOPS 0 Read 0 Write 0

Figure 4-37 Child pool list

Actions on child storage pools

All actions supported for parent storage pools are supported for child storage pools, with the exception of **Add Storage**. Child pools additionally support the **Resize** action.

To select an action right-click the child storage pool, as shown in Figure 4-38. Alternatively, select the storage pool and click **Actions**.

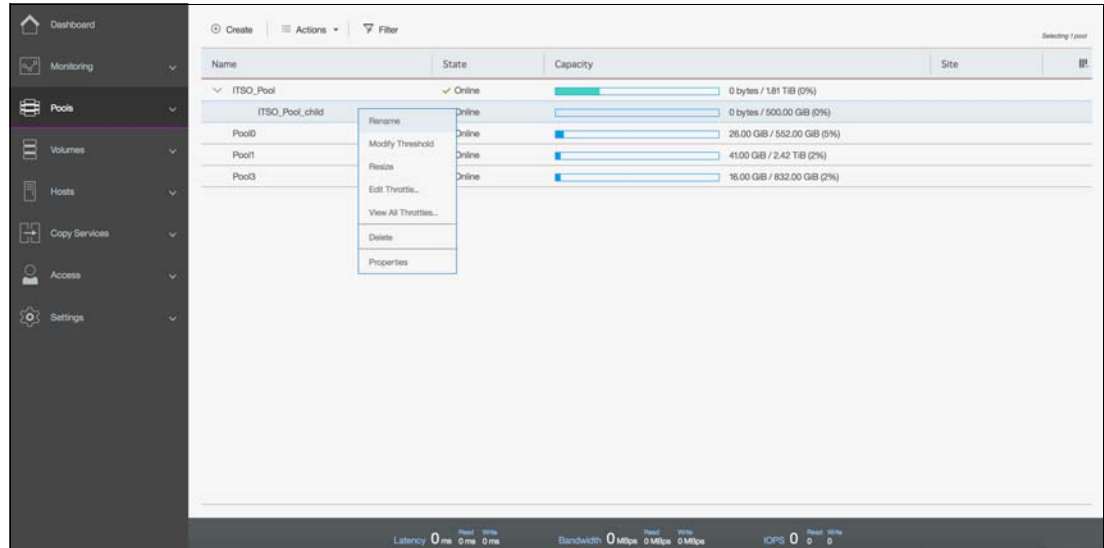


Figure 4-38 Child pools list of actions

Resize

Selecting **Resize** allows you to increase or decrease the capacity of the child storage pool, as shown in Figure 4-39. Enter the new pool capacity and click **Resize**.

Note: You cannot shrink a child pool below its real capacity. Thus, the new size of a child pool needs to be larger than the capacity used by its volumes.

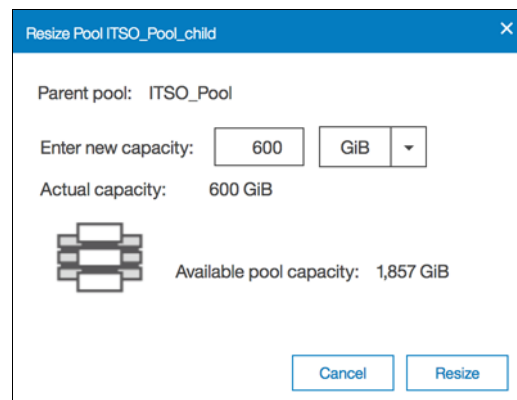


Figure 4-39 Resizing child pools

Delete

Deleting a child pool is a task quite similar to deleting a parent pool. As with a parent pool, the **Delete** action is disabled if the child pool contains volumes. After deleting a child pool the extents that were being occupied return to the parent pool as free capacity.

Note: A volume in a child pool can only be migrated to another child pool within the same parent pool or to its own parent pool. In any other case use volume mirroring instead. During migration from a child pool to its parent pool, or vice versa, there is no real data move. There is only a reassignment of extents between the pools.

4.3 Working with managed disks

A storage pool is created as an empty container, with no storage assigned to it. Storage is then added in the form of MDisks. An MDisk can be either an array from internal storage or an LU from an external storage system. The same storage pool can include both internal and external MDisks.

Arrays are created from internal storage using RAID technology to provide redundancy and increased performance. The system supports two types of RAID: traditional RAID and distributed RAID. Arrays are assigned to storage pools at creation time and cannot be moved between storage pools. It is not possible to have an array that does not belong to any storage pool.

External MDisks can have one of the following modes:

- Unmanaged

External MDisks are discovered by the system as unmanaged MDisks. An unmanaged MDisk is not a member of any storage pool, is not associated with any volumes and has no metadata stored on it. The system does not write to an MDisk that is in unmanaged mode, except when it attempts to change the mode of the MDisk to one of the other modes.

- Managed

When unmanaged MDisks are added to storage pools, they become managed. Managed mode MDisks are always members of a storage pool and provide extents to the storage pool. This mode is the most common and normal mode for an MDisk.

- Image

Image mode provides a direct block-for-block translation from the MDisk to a volume. This mode is provided to satisfy the following major usage scenarios:

- Virtualization of external LUs that contain data not written through the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030
- Exporting MDisks from the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 after volume migrations to image mode MDisks.

MDisks are managed through the MDisks by Pools panel. To access the MDisks by Pools panel browse to **Pools** → **MDisks by Pools**, as shown in Figure 4-40 on page 163.

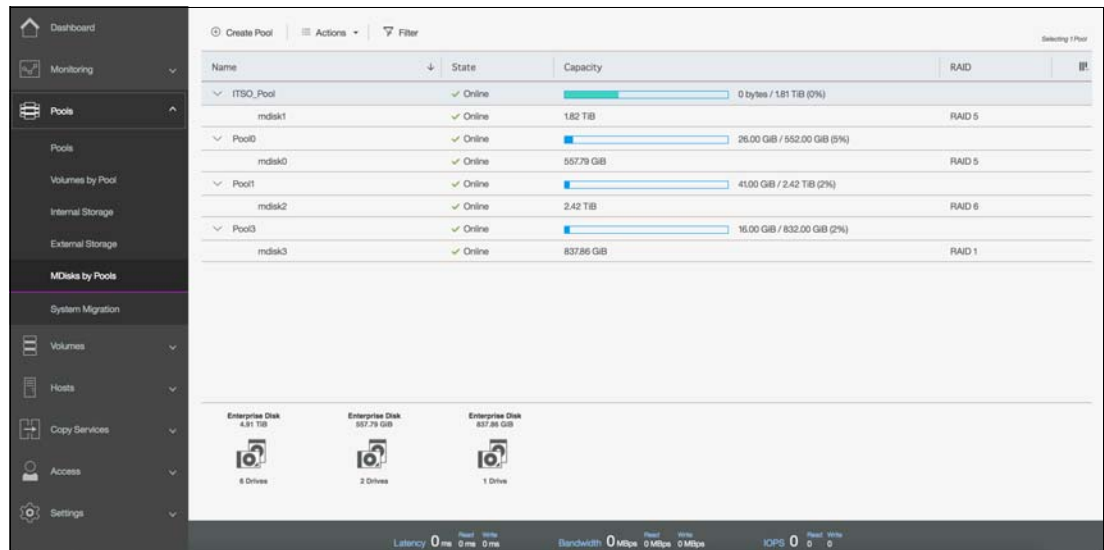


Figure 4-40 MDisk1s by Pools panel

The panel lists all the MDisk1s available in the system under the storage pool to which they belong.

4.3.1 Assigning managed disks to storage pools

MDisk1s can be assigned to a storage pool at any time to increase the number of extents available in the pool. The system automatically balances volume extents between the MDisk1s to provide the best performance to the volumes.

Arrays are created and assigned to a storage pool at the same time.

To assign MDisk1s to a storage pool navigate to **Pools** → **MDisk1s by Pools** and choose one of the following options:

- Option 1: Select **Add Storage** on the right side of the storage pool, as shown in Figure 4-41 on page 164. The Add Storage button is shown only when the pool has no capacity assigned or when the pool capacity usage is over the warning threshold.




Figure 4-41 Add storage: option 1

- Option 2: Right-click the pool and select **Add Storage**, as shown in Figure 4-42. Alternatively, select the a pool and click **Actions**.

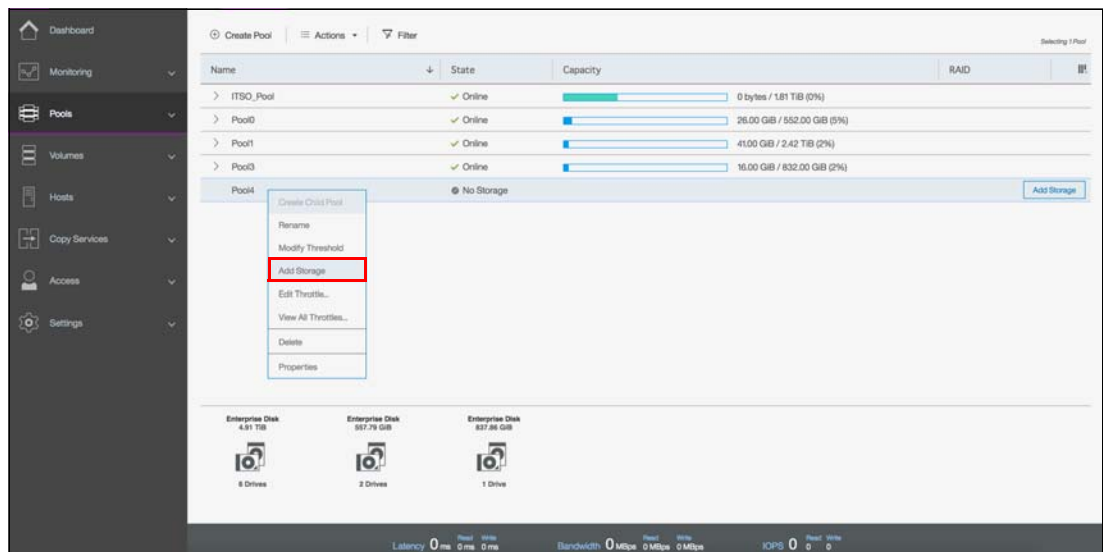


Figure 4-42 Add storage: option 2

- Option 3: Select **Assign** under a specific drive class or external storage controller, as shown in Figure 4-43 on page 165.




Figure 4-43 Add storage: option 3

Both options 1 and 2 start the configuration wizard shown in Figure 4-44.

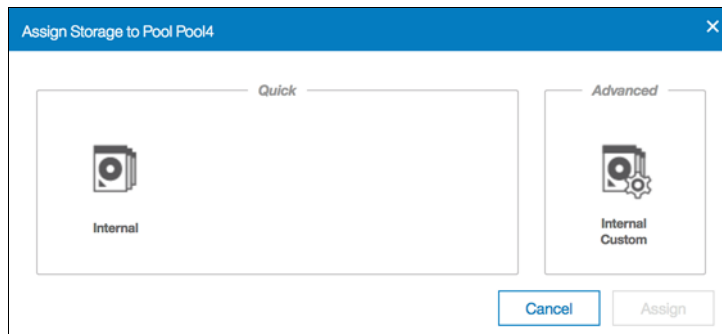


Figure 4-44 Assigning storage to storage pool

Option 3 starts the quick internal wizard for the selected drive class only, as shown in Figure 4-45.

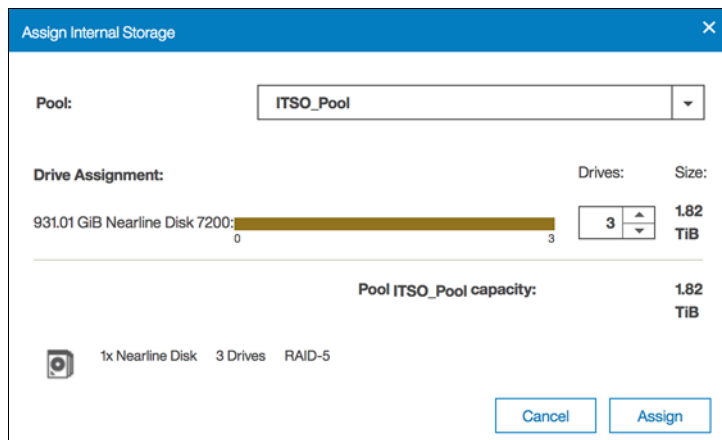


Figure 4-45 Assigning specific storage class

Quick internal configuration

Selecting **Internal** suggests a configuration for internal drives based on RAID configuration presets, considering drive class and number of drives available. It automatically defaults parameters such as stripe width, number of spares (for traditional RAID), number of rebuild areas (for distributed RAID), and number of drives of each class. The number of drives is the only value that can be adjusted.

Figure 4-46 shows an example of a quick configuration.

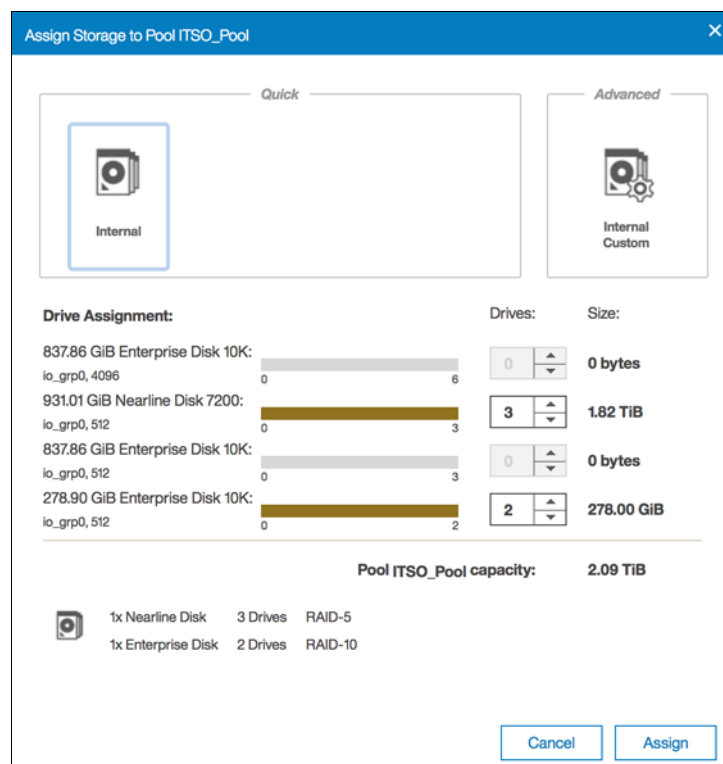


Figure 4-46 Quick configuration wizard

This configuration combines two drive classes, belonging to two different tiers of storage (Nearline and Enterprise). This is the default option and takes advantage of the Easy Tier functionality. However, this can be adjusted by setting the number of drives of different classes to zero as shown in Figure 4-47 on page 167.

Note: If any drive class is not compatible with the drives being assigned that drive class cannot be selected.

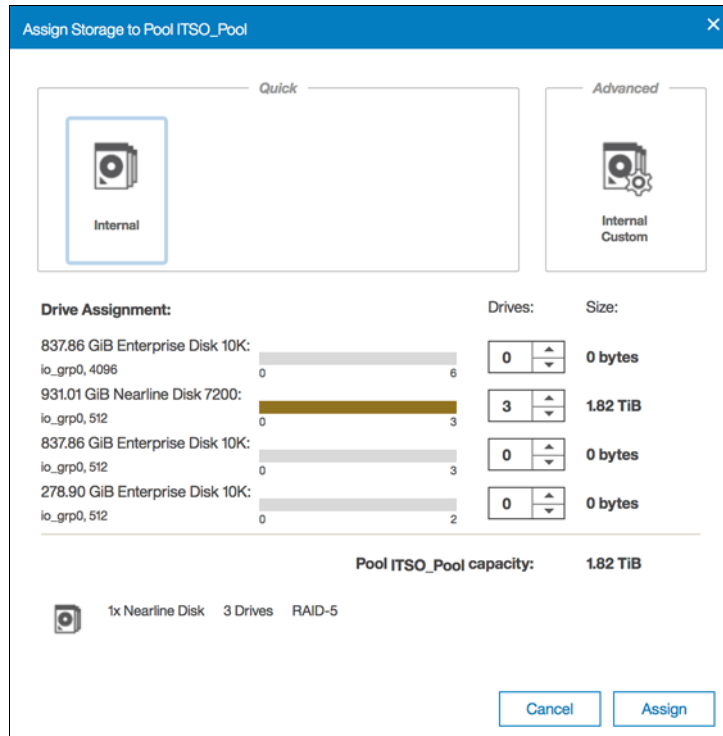


Figure 4-47 Quick configuration wizard with a zeroed storage class

If you are adding storage to a pool with storage already assigned, the existing storage is also taken into consideration, with some properties being inherited from existing arrays for a given drive class. Drive classes incompatible with the classes already in the pool are disabled as well.

When you are satisfied with the presented configuration click **Assign** as shown in Figure 4-48 on page 168. The array MDisks are then created and initializes on the background.

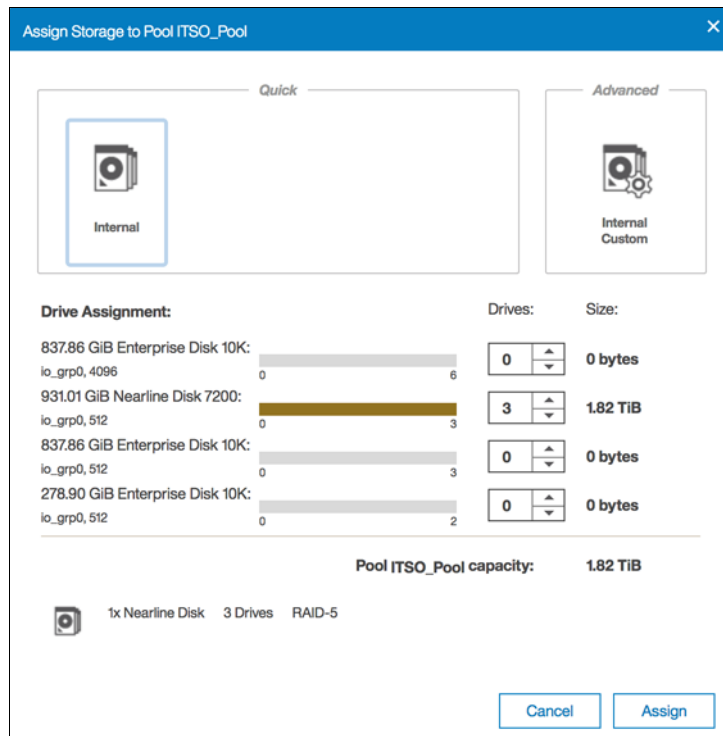


Figure 4-48 Clicking assign on quick configuration wizard

Advanced internal custom configuration

Selecting **Internal Custom** allows the user to customize the configuration for internal drives.

Tip: It is advised to use the advanced configuration only when the quick configuration suggested does not fit your business requirements.

The following values can be customized:

- ▶ RAID level
- ▶ Number of spares
- ▶ Array width
- ▶ Stripe width
- ▶ Number of drives of each class

Figure 4-49 on page 169 shows an example with 6 drives ready to be configured as RAID 6. Click **Summary** to see the list of MDisk arrays to be created. To return to the default settings, select the refresh button next to the pool capacity and to create and assign the arrays, click **Assign**.

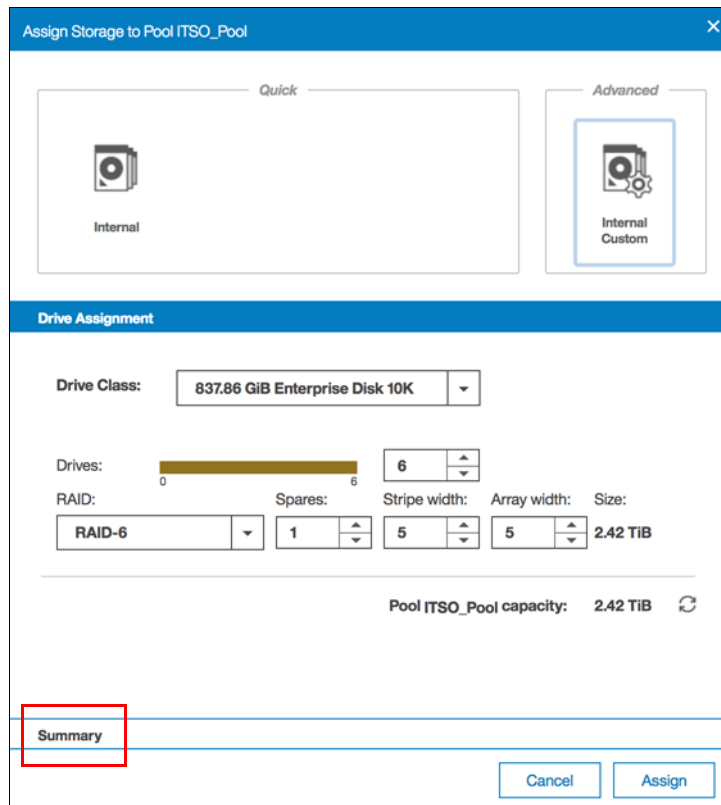


Figure 4-49 Advanced internal custom configuration

4.3.2 RAID configuration

In this topic, we describe the Redundant Array of Independent Disks (RAID) configuration and technology.

Introduction to RAID technology

RAID provides two key design goals:

- ▶ Increased data reliability
- ▶ Increased input/output (I/O) performance

When multiple physical disks are set up to use the RAID technology, they are in a *RAID array*. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 provide multiple, traditional RAID levels:

- ▶ RAID 0
- ▶ RAID 1
- ▶ RAID 5
- ▶ RAID 6
- ▶ RAID 10

RAID technology can provide better performance for data access, high availability for the data, or a combination. RAID levels define a trade-off between high availability, performance, and cost.

The RAID concept must be extended to *disk rebuild time* because of increasing physical disk capacity.

In a disk failure, traditional RAID writes the data to a single spare drive. With increasing capacity, the rebuild time is also increased and the probability of a second failure during the rebuild process becomes more likely, as well. In addition, the spares, when they are not being used, are idle, wasting resources.

Distributed RAID (DRAID) addresses those points and it is available for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 in two types:

- ▶ Distributed RAID 5 (DRAID 5)
- ▶ Distributed RAID 6 (DRAID 6)

Distributed RAID reduces the recovery time and the probability of a second failure during rebuild. Just like traditional RAID, a distributed RAID 5 array can lose one physical drive and survive. If another drive fails in the same array before the bad drive is recovered, the MDisk and the storage pool go offline as they are supposed to. So, distributed RAID does not change the general RAID behavior.

Note: Although Traditional RAID is still supported and is the default choice in the GUI, the recommendation is to use Distributed RAID 6 whenever possible.

4.3.3 Distributed RAID

In distributed RAID, all drives are active, which improves performance. Spare capacity is used instead of the idle spare drives from traditional RAID. As no drives are spare, all drives contribute to performance. The spare capacity is rotated across the disk drives so the write rebuild load is distributed across multiple drives and the bottleneck of one drive is removed.

Figure 4-50 on page 171 shows an example of a distributed RAID with 10 disks. The physical disk drives are divided into multiple packs. The reserved spare capacity (which is marked in yellow) is equivalent to two spare drives, but the capacity distributed across all physical disk drives. The data is distributed across a single row. For simplification, not all packs are shown in Figure 4-50 on page 171.




Figure 4-50 Distributed RAID 6

Figure 4-51 on page 172 shows a single drive failure in the distributed RAID 6 (DRAID 6) environment. Physical disk 3 failed and the RAID 6 algorithm is using the spare capacity for a single spare drive in each pack for rebuild (which is marked in green). All disk drives are involved in the rebuild process, which significantly reduces the rebuild time. For simplification, not all packs are shown in Figure 4-51 on page 172.




Figure 4-51 Single drive failure with DRAID 6

The usage of multiple drives improves the rebuild process, which is up to 10 times faster than traditional RAID. This speed is even more important when you use large drives.

The conversion from traditional RAID to distributed RAID is possible by using volume mirroring or volume migration. Mixing traditional RAID and distributed RAID in the same storage pool is also possible.

Example

The same number of disks can be configured by using traditional or distributed RAID. In our example, we use 6 disk drives and assign those disks as RAID 6 to a single pool.

Figure 4-52 shows the setup for a traditional RAID 6 environment. The pool consists of one MDisk, with 5 disk drives. The spare drive is not listed in this summary.




Figure 4-52 Array configuration for a traditional RAID 6 with 6 disks

Figure 4-53 shows the setup for a distributed RAID 6 environment. The pool consists of a single MDisk with 6 disk drives. The spare drives are included in this summary.




Figure 4-53 Array configuration for a distributed RAID 6 with 6 disks

4.3.4 RAID configuration presets

RAID configuration presets are used to configure internal drives. They are based on the advised values for the RAID level and drive class. Each preset has a specific goal for the number of drives per array and the number of spare drives to maintain redundancy.

For the best performance with solid-state drives (SSDs), arrays with the same number of drives are recommended, which is the same design for traditional RAID arrays.

Table 4-1 describes the presets that are used for Flash drives for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems.

Table 4-1 Flash RAID presets

Preset	Purpose	RAID level	Drives per array goal	Drive count (min - max)	Spare drive goal
Flash RAID 5	Protects against a single drive failure. Data and one stripe of parity are striped across all array members.	5	8	3 - 16	1
Flash Distributed RAID 5	Protects against a single drive failure. Data and one stripe of parity are striped across all array members.	5	8	3 - 16	1
Flash RAID 6	Protects against two drive failures. Data and two stripes of parity are striped across all array members.	6	12	5 - 16	1
Flash Distributed RAID 6	Protects against two drive failures. Data and two stripes of parity are striped across all array members.	6	12	5 - 16	1
Flash RAID 10	Protects against at least one drive failure. All data is mirrored on two array members.	10	8	4 - 16 (even number of drives)	1
Flash RAID 1	Protects against at least one drive failure. All data is mirrored on two array members.	1	2	2	1
Flash RAID 0	Provides no protection against drive failures.	0	8	1 - 8	0
Flash Easy Tier	Mirrors data to protect against drive failure. The mirrored pairs are spread between storage pools to use for the Easy Tier function.	10	2	4 - 16 (even number of drives)	1

Flash RAID instances: In all Flash RAID instances, drives in the array are balanced across enclosure chains, if possible.

Table 4-2 describes the RAID presets that are used for Enterprise SAS and Nearline SAS drives for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems.

Table 4-2 Hard disk drive (HDD) RAID presets

Preset	Purpose	RAID level	Drives per array goal	Drive count (min - max)	Spare goal	Chain balance
Basic RAID 5	Protects against a single drive failure. Data and one stripe of parity are striped across all array members.	5	8	3 - 16	1	All drives in the array are from the same chain wherever possible.
Distributed RAID 5	Protects against a single drive failure. Data and one stripe of parity are striped across all array members.	5	48 - 60	4 - 128	1: 0 - 36 2: 37 - 72 3: 73 - 100 4: 101 - 128	All drives in the array are from the same chain wherever possible.
Basic RAID 6	Protects against two drive failures. Data and two stripes of parity are striped across all array members.	6	12	5 - 16	1	All drives in the array are from the same chain wherever possible.
Distributed RAID 6	Protects against two drive failures. Data and two stripes of parity are striped across all array members.	6	48 - 60	6 - 128	1: 0 - 36 2: 37 - 72 3: 73 - 100 4: 101 - 128	All drives in the array are from the same chain wherever possible.
Basic RAID 10	Protects against at least one drive failure. All data is mirrored on two array members.	10	8	4 - 16 (must be an even number of drives)	1	All drives in the array are from the same chain wherever possible.
Balanced RAID 10	Protects against at least one drive or enclosure failure. All data is mirrored on two array members. The mirrors are balanced across the two enclosure chains.	10	8	4 - 16 (even)	1	Exactly half of the drives are from each chain.
RAID 0	Provides no protection against drive failures.	0	8	1 - 8	0	All drives in the array are from the same chain wherever possible.

4.3.5 Actions on arrays

MDisks created from internal storage are RAID arrays and support specific actions that are not supported on external MDisks. Some actions supported on traditional RAID arrays are not supported on distributed RAID arrays and vice versa.

To choose an action select the array and click **Actions**, as shown in Figure 4-54 on page 175. Alternatively, right-click the array.

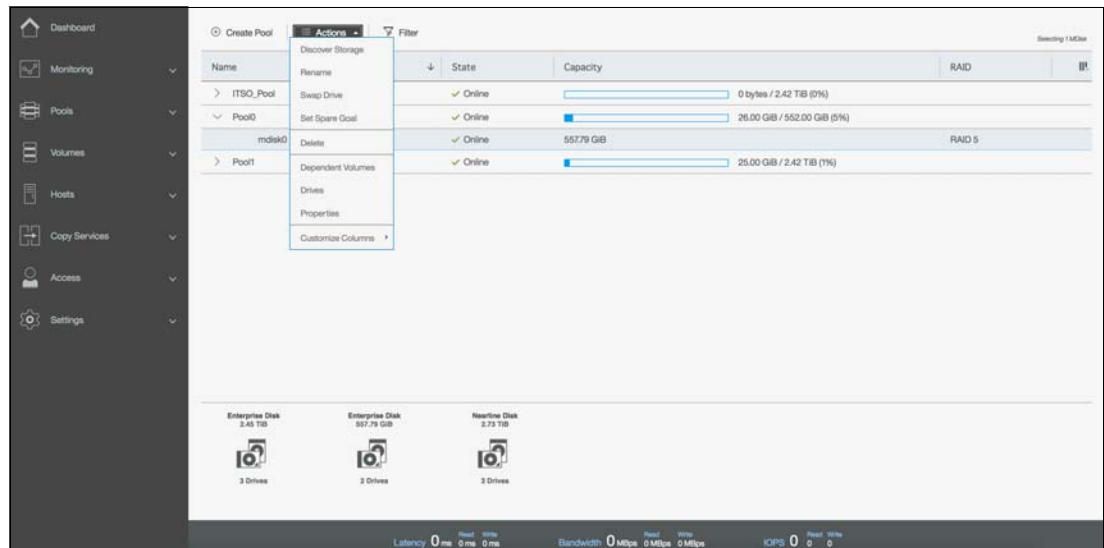


Figure 4-54 Available actions on arrays

Swap drive

Selecting **Swap Drive** allows the user to replace a drive in an array with another drive. The other drive needs to have a use of Candidate or Spare. This action can be used to replace a drive that is expected to fail soon.

Figure 4-55 shows the dialog box that opens. Select the member drive to be replaced and the replacement drive.

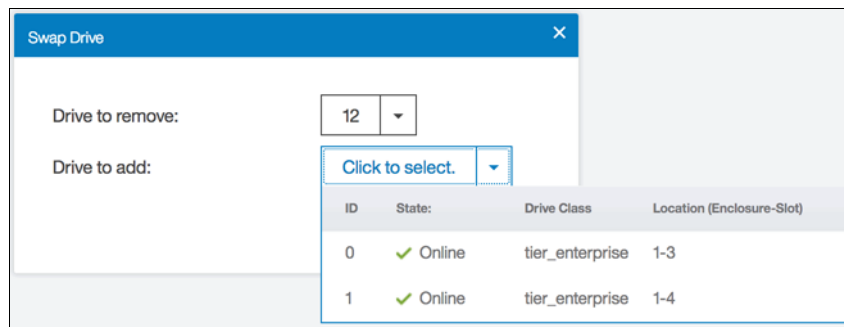


Figure 4-55 Swap Drive panel

After defining the disk to be removed and the disk to be added, click **Swap** as shown in Figure 4-56.

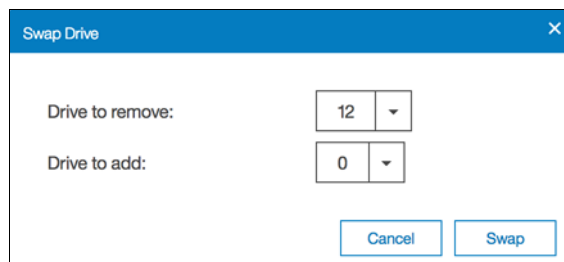


Figure 4-56 Swap button on Swap Drive panel

Set Spare Goal

This action is available only for traditional RAID arrays. Selecting **Set Spare Goal** allows you to set the number of spare drives required to protect the array from drive failures. If the number of spare drives available does not meet the configured goal an error is logged in the event log. This error can be fixed by adding more drives of a compatible drive class as spares.

Figure 4-57 shows the dialog box that opens when you select **Set Spare Goal**. Define the number of required spares and click **Save**.

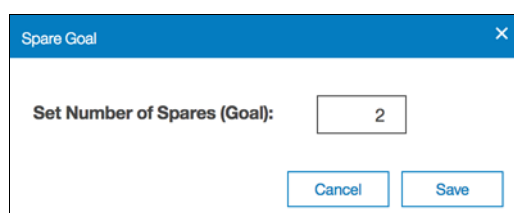
A screenshot of a dialog box titled "Spare Goal" with a close button (X) in the top right corner. The main content area contains the text "Set Number of Spares (Goal):" followed by a text input field containing the number "2". At the bottom right of the dialog are two buttons: "Cancel" and "Save".

Figure 4-57 Spare Goal panel

Set rebuild areas goal

This action is available only for distributed RAID arrays. Selecting **Set Rebuild Areas Goal** allows you to set the number of rebuild areas required to protect the array from drive failures. If the number of rebuild areas available does not meet the configured goal, an error is logged in the event log. This error can be fixed by replacing the failed drives in the array with new drives of a compatible drive class.

Figure 4-58 shows the dialog box that opens when you select **Set Rebuild Areas Goal**. Define the representative number of required spares that will compose the rebuild area and click **Save**.

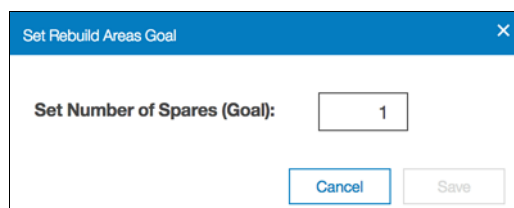
A screenshot of a dialog box titled "Set Rebuild Areas Goal" with a close button (X) in the top right corner. The main content area contains the text "Set Number of Spares (Goal):" followed by a text input field containing the number "1". At the bottom right of the dialog are two buttons: "Cancel" and "Save".

Figure 4-58 Rebuild Areas Goal panel

Delete

Selecting **Delete** removes the array from the storage pool and deletes it.

Remember: An array does not exist outside of a storage pool. Therefore an array cannot be removed from the pool without being deleted.

If there are no volumes using extents from the array the deletion command runs immediately without additional confirmation. If there are volumes using extents from the array, you are prompted to confirm the action as shown in Figure 4-59 on page 177. Click **Yes** to migrate the volumes or **No** to cancel the deletion process.




Figure 4-59 MDisk deletion confirmation panel

Confirming the action starts the migration of the volumes to extents from other MDisks that remain in the pool; after the action completes the array is removed from the storage pool and deleted.

Note: Ensure that you have enough available capacity remaining in the storage pool to allocate the data being migrated from the removed array, otherwise the command fails.

Drives

Selecting **Drives** shows information about the drives that are included in the array, as shown in Figure 4-60.

Member Drives for MDisk mdisk1						
<div> ⋮ Actions Filter Download Showing 6 drives Selecting 0 drives </div>						
Drive ID	Capacity	Use	Status	Enclosure ID	Slot ID	
3	837.86 GiB	Member	✓ Online	1	19	
4	837.86 GiB	Member	✓ Online	1	17	
6	837.86 GiB	Member	✓ Online	1	21	
9	837.86 GiB	Member	✓ Online	1	14	
11	837.86 GiB	Member	✓ Online	1	15	
22	837.86 GiB	Member	✓ Online	1	22	

Figure 4-60 Panel showing the drives that are members of an MDisk

4.3.6 Actions on external MDisks

External MDisks support specific actions that are not supported on arrays. Some actions are supported only on unmanaged external MDisks and some are supported only on managed external MDisks.

To choose an action right-click the external MDisk, as shown in Figure 4-61 on page 178. Alternatively, select the external MDisk and click **Actions**.

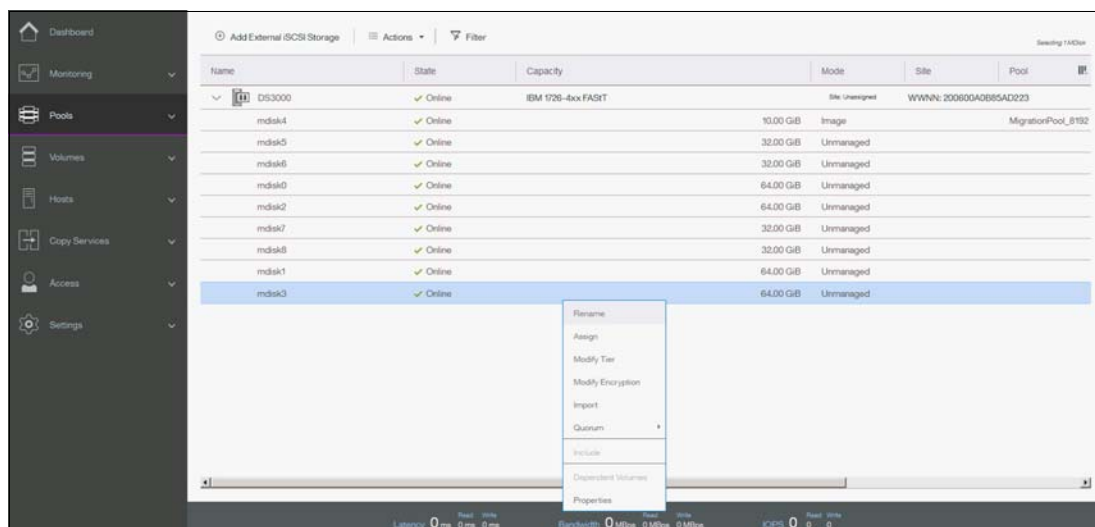


Figure 4-61 Available actions for external MDisks

Assign

This action is available only for unmanaged MDisks. Selecting **Assign** opens the dialog box shown in Figure 4-62. This action acts only on the selected MDisk or MDisks.

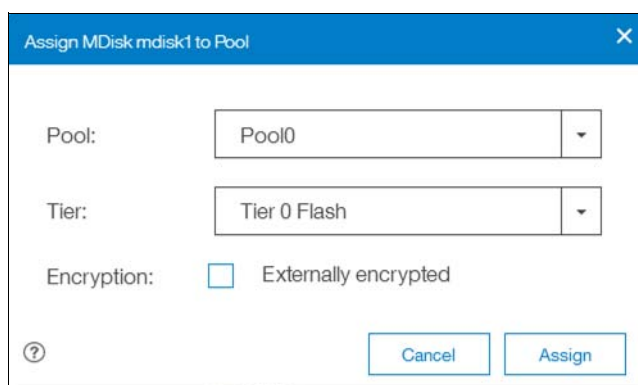


Figure 4-62 Assigning external MDisks to a pool

Important: If you need to preserve existing data on an unmanaged MDisk do *not* assign it to a storage pool because this action *deletes the data* on the MDisk. Use **Import** instead.

Modify tier

Selecting **Modify Tier** allows the user to modify the tier to which the external MDisk is assigned, as shown in Figure 4-63 on page 179. This setting is adjustable because the system cannot detect the tiers associated with external storage automatically. **Enterprise Disk** (Tier 2) is the option selected by default.

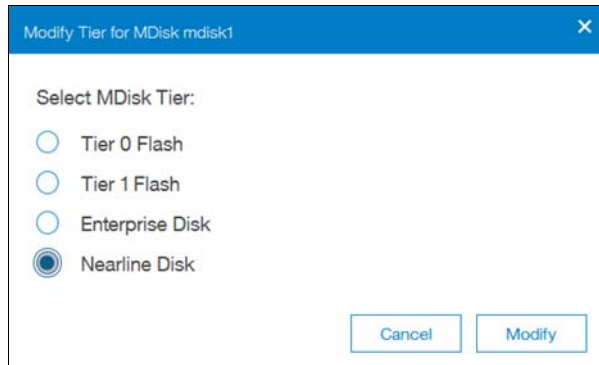


Figure 4-63 Modifying external MDisk tier

This option is available only when encryption is enabled. Selecting **Modify Encryption** allows the user to modify the encryption setting for the MDisk, as shown in Figure 4-64.

For example, if the external MDisk is already encrypted by the external storage system, change the encryption state of the MDisk to **Externally encrypted**. This stops the system from encrypting the MDisk again if the MDisk is part of an encrypted storage pool.



Figure 4-64 Modifying external MDisk encryption

Import

This action is available only for unmanaged MDisks. Importing an unmanaged MDisk allows the user to preserve the data on the MDisk, either by migrating the data to a new volume or by keeping the data on the external system.

Selecting **Import** allows you to choose one of the following migration methods:

- **Import to temporary pool as image-mode volume** does not migrate data from the source MDisk. It creates an *image-mode volume* that has a direct block-for-block translation of the MDisk. The existing data is preserved on the external storage system, but it is also accessible from the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems.

If this method is selected the image-mode volume is created in a temporary migration pool and presented through the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. Choose the extent size of the temporary pool and click **Import**, as shown in Figure 4-65 on page 180.

Import MDisk mdisk1 As Volume

Volume name:

Import method: ☒ Import to temporary pool as image-mode volume
☐ Migrate to an existing pool

Extent Size:

☐ Copy Services on the external storage system are used with this volume

Figure 4-65 Importing an external MDisk as an image-mode volume

The MDisk is imported and listed as an image-mode MDisk in the temporary migration pool, as shown in Figure 4-66.

Name	State	Capacity	Mode	Site	Pool
DS3000	Online	IBM 128-bx FASST	Site Unmanaged	WWNN: 200820A0B85AD223	
mdisk2	Online	64.00 GiB	Unmanaged		
mdisk4	Online	10.00 GiB	Image		MigrationPool_8192
mdisk5	Online	32.00 GiB	Unmanaged		
mdisk6	Online	32.00 GiB	Unmanaged		
mdisk3	Online	64.00 GiB	Unmanaged		
mdisk5	Online	32.00 GiB	Unmanaged		
mdisk6	Online	64.00 GiB	Unmanaged		
mdisk1	Online	64.00 GiB	Image		MigrationPool_302
mdisk7	Online	32.00 GiB	Unmanaged		

Figure 4-66 Image-mode MDisk

A corresponding image-mode volume is now available in the same migration pool, as shown in Figure 4-67 on page 181.

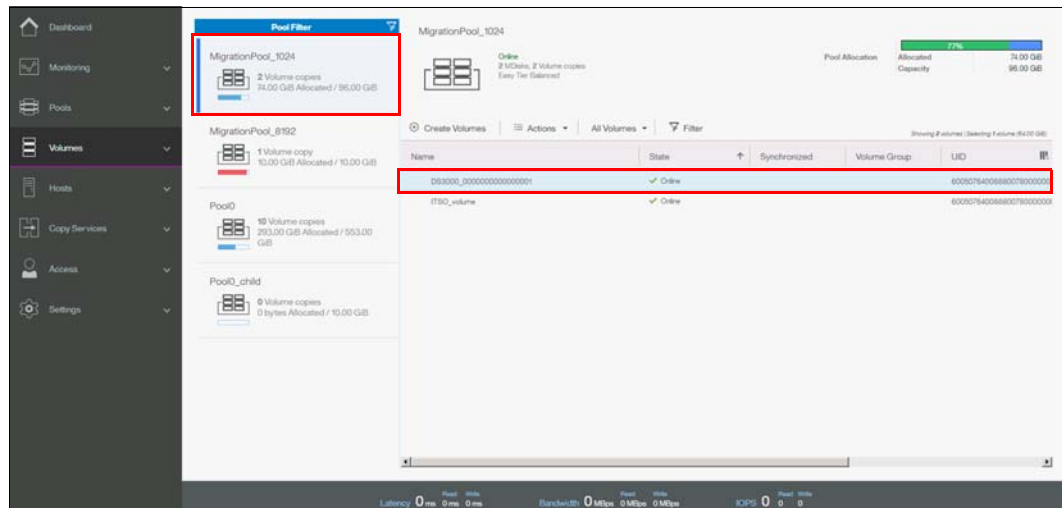


Figure 4-67 Corresponding image-mode volume

The image-mode volume can then be mapped to the original host mode. The data is still physically present on the physical disk of the original external storage controller system and no automatic migration process is currently running. If needed, the image-mode volume can be migrated manually to another storage pool using volume migration or volume mirroring later.

- **Migrate to an existing pool** starts by creating an image-mode volume as the first method. However, it then migrates the data from the image-mode volume onto another volume in the selected storage pool. After the migration process completes the image-mode volume and temporary migration pool are deleted.

If this method is selected, choose the storage pool to hold the new volume and click **Import**, as shown in Figure 4-68. Only pools with sufficient free extent capacity are listed.

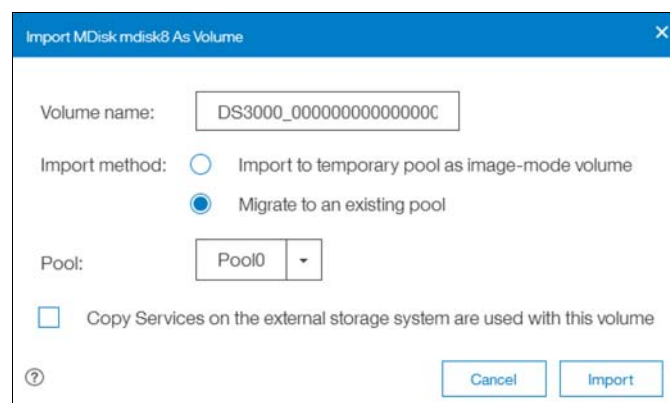


Figure 4-68 Importing an external MDisk to an existing pool

The data migration begins automatically after the MDisk is imported successfully as an image-mode volume. You can check the migration progress by navigating to **Pools** → **System Migration**, as shown in Figure 4-69 on page 182.


[illegible]

Figure 4-70 Striped volume after migration

At this point all data has been migrated from the source MDisk and the MDisk is no longer in image mode, as shown in Figure 4-71 on page 183. The MDisk can be removed from the temporary pool and used as a regular MDisk to host volumes.

Name	State	Capacity	Mode	Site	Pool
DISK000	✓ Online	IBM 526-4xx FAST	Site Unmanaged	WWN: 200600A0B85AD223	
mdisk2	✓ Online	64.00 GB	Unmanaged		
mdisk5	✓ Online	32.00 GB	Unmanaged		
mdisk3	✓ Online	64.00 GB	Unmanaged		
mdisk0	✓ Online	64.00 GB	Unmanaged		
mdisk7	✓ Online	32.00 GB	Unmanaged		
mdisk1	✓ Online	64.00 GB	Unmanaged		
mdisk6	✓ Online	32.00 GB	Managed		MigrationPool_1024
mdisk4	✓ Online	10.00 GB	Image		MigrationPool_8192

Figure 4-71

Alternatively, import and migration of external MDisks to another pool can be done by selecting **Pools** → **System Migration**. Migration and the system migration wizard are described in more detail in Chapter 7, “Storage migration” on page 323.

Include

The system can exclude an MDisk with multiple I/O failures or persistent connection errors from its storage pool to ensure these errors do not interfere with data access. If an MDisk has been automatically excluded, run the fix procedures to resolve any connection and I/O failure errors. Drives used by the excluded MDisk with multiple errors might require replacing or reseating.

After the problems have been fixed, select **Include** to add the excluded MDisk back into the storage pool.

Remove

In some cases you may want to remove external MDisks from storage pools to reorganize your storage allocation. Selecting **Remove** removes the MDisk from the storage pool. After the MDisk is removed it goes back to unmanaged. If there are no volumes in the storage pool to which this MDisk is allocated the command runs immediately without additional confirmation. If there are volumes in the pool, you are prompted to confirm the action, as shown in Figure 4-72. Click **Yes** to migrate the volumes or **No** to cancel the deletion process.




Figure 4-72 Removing an external MDisk

Confirming the action starts the migration of the volumes to extents from other MDisks that remain in the pool; when the action completes the MDisk is removed from the storage pool and returns to unmanaged.

Note: Ensure that you have enough available capacity remaining in the storage pool to allocate the data being migrated from the removed MDisk or else the command fails.

Important: The MDisk being removed must remain accessible to the system while all data is copied to other MDisk in the same storage pool. If the MDisk is unmapped before the migration finishes all volumes in the storage pool go offline and remain in this state until the removed MDisk is connected again.

4.3.7 More actions on MDisk

There are a few additional actions supported both on arrays and external MDisk.

Discover storage

The Discover storage option in the upper left of the MDisk by Pools window is useful if external storage controllers are in your environment. (For more information, see Chapter 11, “External storage virtualization” on page 607). The Discover storage action starts a rescan of the Fibre Channel network. It discovers any new MDisk that were mapped to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems and rebalances MDisk access across the available controller device ports.

This action also detects any loss of controller port availability and updates the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 configuration to reflect any changes.

When external storage controllers are added to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 environment, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 automatically discover the controllers, and the logical unit numbers (LUNs) that are presented by those controllers are listed as unmanaged MDisk.

However, if you attached new storage and the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 did not detect it, you might need to use the Discover storage option before the system detects the new LUNs. If the configuration of the external controllers is modified afterward, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 might be unaware of these configuration changes. Use Detect MDisk to rescan the Fibre Channel network and update the list of unmanaged MDisk.

Figure 4-73 on page 185 shows the Discover storage option.

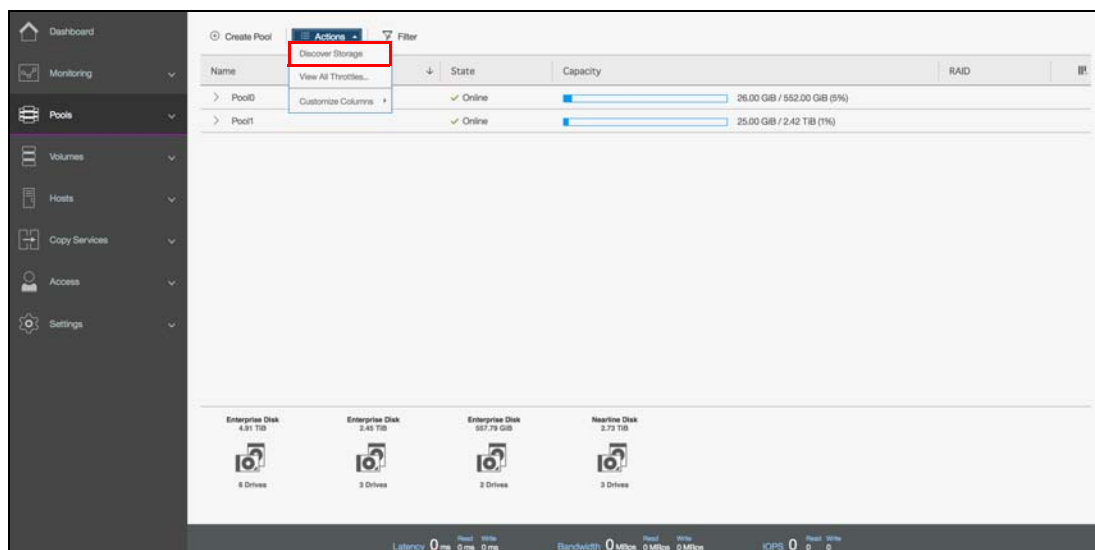


Figure 4-73 Discover storage action

Note: The Discover storage action is asynchronous. Although the task appears to be finished, it might still be running in the background.

Rename

MDisks can be renamed by selecting the MDisk and clicking **Rename** from the Actions menu. Enter the new name of your MDisk (as shown in Figure 4-74) and click **Rename**.

The screenshot shows a dialog box titled 'Rename MDisk mdisk2'. It contains a text input field with the text 'ITSO_MDisk'. Below the input field are two buttons: 'Cancel' and 'Rename'.

Figure 4-74 Rename MDisk

Show Dependent Volumes

Figure 4-75 on page 186 shows the volumes that depend on an MDisk. The volumes can be displayed by selecting the MDisk and clicking **Show Dependent Volumes** from the Actions menu. The volumes are listed with general information.

Name	State	Capacity	Pool	UID
ITSO_BASIC_VOL	✓ Online	1.00 GiB	Pool1	6005076300A600000800000000000002
ITSO_COMPR_VOL	✓ Online	1.00 GiB	Pool1	6005076300A600000800000000000005
ITSO_MIG_BY_VOLCOPY_V...	✓ Online	3.00 GiB	Pool1	6005076300A600000800000000000016
ITSO_TP_VOL	✓ Online	1.00 GiB	Pool1	6005076300A600000800000000000004

Figure 4-75 Show MDisk dependent volumes

Properties

The Properties action for an MDisk shows the information that you need to identify it. In the MDisks by Pools window, select the MDisk and click **Properties** from the Actions menu. Alternatively, right-click the MDisk and select **Properties**. For additional information related to the selected MDisk, click **View more details** as shown in Figure 4-76.

Name: mdisk2
 State: ✓ Online
 ID: 2
 Capacity: 2.42 TiB
 Pool: Pool1
[View more details](#)

Figure 4-76 MDisk properties

4.4 Working with external storage controllers

After the internal storage configuration is complete, you can find the MDisks that were created by using the internal drives in the MDisks by Pools window. When you use this window, you can manage all MDisks that are made up of internal and external storage.

Logical unit numbers (LUNs) that are presented by external storage systems to Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are discovered as unmanaged MDisks. Initially, the MDisk is not a member of any storage pools, which means that it is not used by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems.

To learn more about external storage, see Chapter 11, “External storage virtualization” on page 607.

Host configuration

This chapter describes how to use the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 graphical user interface (GUI) to create hosts and how to prepare a host to access the volumes that are created. (Volume creation is described in Chapter 6, “Volume configuration” on page 269.)

Specifically, this chapter provides information about the following topics:

- ▶ 5.1, “Host attachment overview” on page 190
- ▶ 5.2, “Planning for direct-attached hosts” on page 191
- ▶ 5.3, “Preparing the host operating system” on page 191
- ▶ 5.4, “N-Port Virtualization ID (NPIV) Support” on page 218
- ▶ 5.5, “Creating hosts by using the GUI” on page 226
- ▶ 5.6, “Host Clusters” on page 245
- ▶ 5.7, “Proactive Host Failover” on page 267

5.1 Host attachment overview

A host system is an open-systems computer that is connected to the switch through a Fibre Channel (FC), direct-attached, serial-attached SCSI (SAS) connection or an Internet Small Computer System Interface (iSCSI).

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support the following host attachment protocols:

- ▶ 16 Gb FC **or** 10 Gb iSCSI/FC over Ethernet (FCoE) as an optional host interface
- ▶ 12 Gb SAS (standard host interface)
- ▶ 1 Gb or 10 Gb iSCSI (standard host interface, depending on the model)

This chapter describes the following topics:

- ▶ Prepare the host operating system:
 - Microsoft Windows:
 - FC
 - iSCSI
 - SAS
 - VMware:
 - FC
 - iSCSI
 - SAS
- ▶ Create hosts by using the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI:
 - Creating FC hosts
 - Creating iSCSI hosts
 - Creating SAS hosts

In this chapter, we assume that your hosts are connected to your FC, SAS, or Internet Protocol (IP) network and you completed the steps that are described in Chapter 2, “Initial configuration” on page 35. Follow basic zoning recommendations to ensure that each host has at least two network adapters, that each adapter is on a separate network (or at minimum in a separate zone), and that each adapter is connected to both canisters. This setup ensures four paths for failover and failback.

For SAS attachment, ensure that each host has at least one SAS host bus adapter (HBA) connection to each Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 canister. Further information about configuring SAS attached hosts is provided in 2.4, “SAS direct-attach planning” on page 45.

Before you map the newly created volumes on the host of your choice, preparation goes a long way toward ease of use and reliability. Several steps are required on a host in preparation for mapping new Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 volumes to the host. Use the Lenovo Storage interoperability matrix to check the code levels that are supported to attach your host to your storage. The Lenovo interoperability matrix is a web tool that checks the interoperation of host, storage, switches, and multipathing drivers. The interoperability matrix can be obtained at the following addresses:

<https://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v3700v2/6535/documentation>

<https://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v5030/6536/documentation>

This chapter focuses on Windows and VMware. If you want to attach any other hosts, for example, IBM AIX, Linux, or an Apple system, you can find the required information at the following address:

http://systemx.lenovofiles.com/mobile/help/topic/com.lenovo.storage.v3700.doc/lenovo_vseries.html

5.2 Planning for direct-attached hosts

Starting with V7.5, we supported direct-attached Fibre Channel hosts. A direct-attached configuration dedicates the entire port bandwidth for use for a specific connection. Planning must account for the volume of expected traffic before you decide how many ports are used in a direct-attached configuration. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems offer multiple options for you to decide how to create a direct-attached configuration.

Because of the bandwidth requirements when you use a direct-attached configuration, it is important to determine the volume of expected traffic when you decide the number of required ports. For example, if a Lenovo Storage V5030 to Lenovo Storage V5030 direct-attached link is configured between nodes, the link might carry intra-node traffic, such as FlashCopy data. Therefore, enough Lenovo Storage V5030 to Lenovo Storage V5030 bandwidth needs to be available so that it can carry all possible intra-node traffic without any bandwidth bottleneck.

The following guidelines are provided for direct-attached configurations.

5.2.1 Fibre Channel direct attachment to host systems

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support direct attachment connectivity between its FC ports and host ports. Host systems can connect to 16 Gb FC ports on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. No special configuration is required for host systems that use this configuration.

5.2.2 FC direct attachment between nodes

Direct connection of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 FC ports without using an FC switch is supported. Such connections between the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 nodes might be useful in small configurations where no FC switch exists. It can also be used to connect nodes in the same input/output (I/O) group to provide a dedicated connection for mirroring the fast write cache data. Ensure that sufficient bandwidth is provisioned between nodes to accommodate all of the intra-node traffic.

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support 16 Gb Fibre Channel ports and supports Fibre Channel direct attachment on all 16 Gb ports.

Note: Be careful about the maximum length of the Fibre Channel links in this configuration.

5.3 Preparing the host operating system

The following steps are required to prepare a host to connect to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030:

1. Ensure that the latest supported system updates are applied to your host operating system.
2. Ensure that the HBAs are physically installed in the host.
3. Ensure that the latest supported HBA firmware and driver levels are installed in your host.
4. Configure the HBA parameters. Although settings are provided for each host OS in the following sections, review the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Information Center to obtain the latest supported settings.
5. Configure the host I/O parameters, such as the disk I/O time-out value.
6. Install and configure multipath software.
7. Determine the host worldwide port names (WWPNs).
8. Connect the HBA ports to switches by using the correct cables, or directly attach to the ports on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.
9. Configure the switches, if applicable.
10. Optional: Configure SAN Boot.

5.3.1 Windows 2008 R2 and 2012 R2: Preparing for Fibre Channel attachment

Complete the following steps to prepare a Windows 2008 R2 or Windows 2012 R2 host to connect to an Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 by using Fibre Channel:

1. Ensure that the current O/S service pack and test fixes are applied to your server.
2. Use the current firmware and driver levels on your host system.
3. Install a host bus adapter (HBA) or HBAs on the Windows server by using the current basic input/output system (BIOS) and drivers.
4. Connect the FC host adapter ports to the switches, or use direct connections.
5. Configure the switches (zoning).
6. Configure the HBA for hosts that run Windows.
7. Set the Windows time-out value.
8. Install the multipath module.

Downloading and installing the supported drivers and firmware

Install a supported HBA driver for your configuration. Use the Windows Device Manager or vendor tools, such as QLogic Converged Console (QCC) or HBAnyware (Emulex), to install the driver. Brocade adapter software is now maintained by QLogic, so check the QLogic web pages to obtain the correct support for your Brocade adapters. Also, check and update the BIOS (firmware) level of the HBA by using the manufacturer's provided tools. Check the readme file to see whether any Windows registry parameters must be set for the HBA driver.

You can obtain the current supported levels by navigating from the following address:

<https://ibm.biz/BdHKW8>

Configuring Brocade HBAs for Windows

This section applies to Windows hosts with installed Brocade HBAs. After you install the device driver and firmware, you must configure the HBAs. To perform this task, either use the Brocade Host Connectivity Manager (HCM) software or restart into the HBA BIOS, load the adapter defaults, and set the following values:

- ▶ Host Adapter BIOS: Disabled (unless the host is configured for storage area network (SAN) Boot)
- ▶ Queue depth: 4

Configuring QLogic HBAs for Windows

This section applies to Windows hosts with installed QLogic HBAs.

After you install the device driver and firmware, you must configure the HBAs. To perform this task, either use the QLogic QConverge Console (QCC) command-line interface (CLI) software or restart into the HBA BIOS, load the adapter defaults, and set the following values:

- ▶ Host Adapter BIOS: Disabled (unless the host is configured for SAN Boot)
- ▶ Adapter Hard Loop ID: Disabled
- ▶ Connection Options: 1 (only point to point)
- ▶ Logical unit numbers (LUNs) Per Target: 0
- ▶ Port Down Retry Count: 15

The QCC Control Center Software is at this web page:

<https://ibm.biz/BdHjUg>

Configuring Emulex HBAs for Windows

This section applies to Windows hosts with installed Emulex HBAs.

After you install the device driver and firmware, you must configure the HBAs. To perform this task, either use the Emulex HBAnyware software or restart into the HBA BIOS, load the defaults, and set topology to 1 (10F_Port Fabric).

Setting the Windows time-out value

For Windows hosts, the disk I/O time-out value needs to be set to 60 seconds. To verify this setting, complete the following steps:

1. Click **Start** → **Run**. Alternatively, open a Power Shell window.
2. In the dialog box or Power Shell window, type `regedit` and press Enter.
3. In the registry editor, locate the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\disk\TimeOutValue` key.
4. Confirm that the value for the key is 60 (decimal value), and, if not, change the value to 60 (Figure 5-1).




Figure 5-1 Windows time-out value

Installing the multipathing software

Microsoft Multipath Input/Output (MPIO) solutions work with device-specific modules (DSMs) that are written by vendors, but the MPIO driver package does not, by itself, form a complete solution. This joint solution enables the storage vendors to design device-specific solutions that are tightly integrated with the Windows operating system. MPIO is not included with the Windows operating system. Storage vendors must pack the MPIO drivers with their own DSM.

Lenovo Subsystem Device Driver DSM (SDDDSM) is the Lenovo multipath I/O solution that is based on Microsoft MPIO technology. It is a device-specific module that supports Lenovo storage devices on Windows hosts. The intent of MPIO is to provide better integration of a multipath storage solution with the operating system, and it supports the use of multipath in the SAN infrastructure during the startup process for SAN Boot hosts.

To ensure correct multipathing with the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, SDDDSM must be installed on Windows hosts. To install SDDDSM, complete the following steps:

1. Go to the following SDDDSM download matrix to determine the correct level of SDDDSM to install for Windows 2008 R2 or Windows 2012 R2, and then download the package:
http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/svc_w2kmultipathovr_21osvf.html
2. Extract the package to your hard disk drive, and run **setup.exe** to install SDDDSM. A command prompt window opens (Figure 5-2). Confirm the installation by entering Yes.


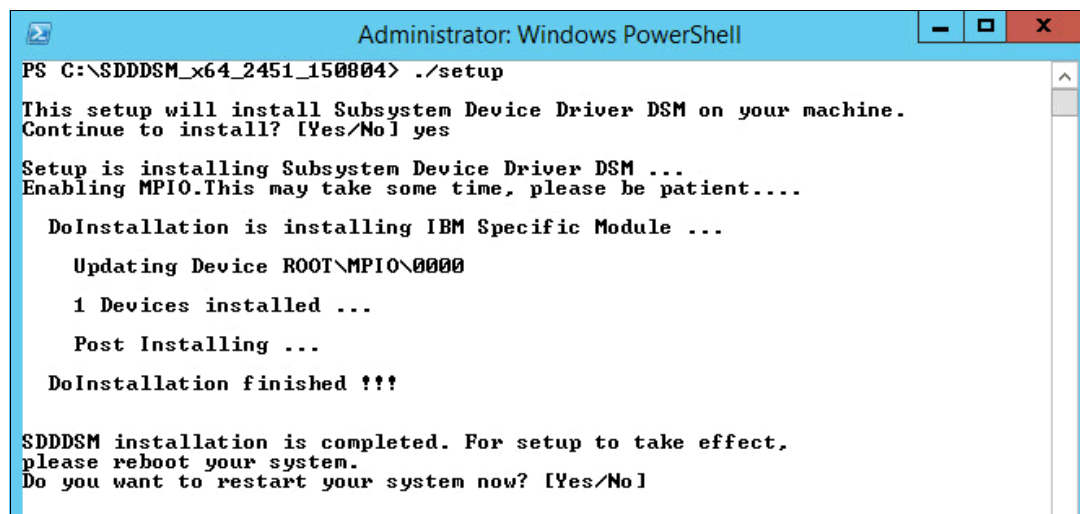


Figure 5-2 SDDDSM setup

3. During the setup, SDDDSM also determines whether an older version is installed and prompts you to upgrade to the current version.

4. After the setup completes, you are prompted to restart the system. Confirm this restart by typing Yes and pressing Enter (Figure 5-3).



```
Administrator: Windows PowerShell
PS C:\SDDDSM_x64_2451_150804> ./setup
This setup will install Subsystem Device Driver DSM on your machine.
Continue to install? [Yes/No] yes
Setup is installing Subsystem Device Driver DSM ...
Enabling MPIO.This may take some time, please be patient....
DoInstallation is installing IBM Specific Module ...
Updating Device ROOT\MPIO\0000
1 Devices installed ...
Post Installing ...
DoInstallation finished !!!
SDDDSM installation is completed. For setup to take effect,
please reboot your system.
Do you want to restart your system now? [Yes/No]
```

Figure 5-3 Answer Yes to restart the host

You successfully installed SDDDSM. To check whether SDDDSM is installed correctly, see the following sections about Windows 2008 R2 and Windows 2012 R2.

Windows 2008 R2

To check the installed driver version, complete the following steps:

1. Select **Start** → **All Programs** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**.
2. A command prompt opens. Run **datapath query version** to determine the version that is installed (Example 5-1) for this Windows 2008 R2 host.

Example 5-1 The datapath query version command

```
C:\Program Files\IBM\SDDDSM>datapath query version
IBM SDDDSM Version 2.4.5.1
Microsoft MPIO Version 6.1.7601.17514
```

3. The worldwide port names (WWPNs) of the FC HBA are required to correctly zone switches and configure host attachment on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. You can obtain the WWPNs by using vendor tools, the HBA BIOS, the native Windows command line, or SDDDSM. This command can be used to determine the worldwide port names (WWPNs) of the host. Run **datapath query wwpn** (Example 5-2) and note the WWPNs of your host because you need them later.

Example 5-2 The datapath query wwpn command

```
C:\Program Files\IBM\SDDDSM>datapath query wwpn
Adapter Name      PortWWN
Scsi Port3:      21000024FF35B960
Scsi Port4:      21000024FF25B961
```

If you need more detailed information about SDDDSM, see *Multipath Subsystem Device Driver User's Guide*, GC52-1309. The guide is available at this web page:

Windows 2012 R2

To check the installed driver version, complete the following steps:

1. First, select the **Windows Start** icon in the lower-left corner. See Figure 5-4.




Figure 5-4 Windows 2012 R2 start

2. Expand the view by clicking the down arrow that is shown in Figure 5-5.




Figure 5-5 Expand view to see all programs that are installed

3. Search for the section *Subsystem Device Driver DSM*. See Figure 5-6.




Figure 5-6 Subsystem Device Driver DSM in the all programs menu

4. Click **Subsystem Device Driver DSM**. See Figure 5-7.




Figure 5-7 Subsystem Device Driver DSM

5. A command prompt opens. Run **datapath query version** to determine the version that is installed. See Figure 5-8.

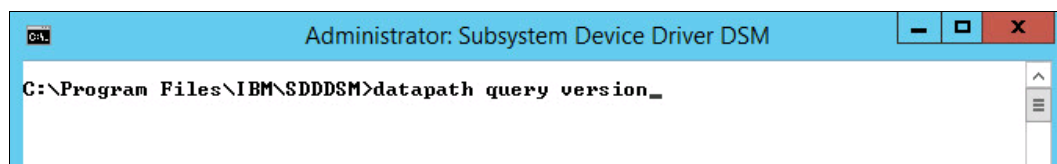


Figure 5-8 Datapath query version

6. See Example 5-3 for the Windows 2012 R2 host.

Example 5-3 The datapath query version command

```
C:\Program Files\IBM\SDDDSM>datapath query version
IBM SDDDSM Version 2.4.5.1
Microsoft MPIO Version 6.3.9600.16384
C:\ProgramFiles\IBM\SDDDSM>
```

7. The WWPNNs of the FC HBA are required to correctly zone switches and configure host attachment on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. You can obtain the WWPNNs by using vendor tools, the HBA BIOS, the native Windows command line, or SDDDSM. This command can be used to determine the WWPNNs of the host. Run **datapath query wwpn** (Example 5-4) and document the WWPNNs of your host because you need them later.

Example 5-4 The datapath query wwpn command

```
C:\Program Files\IBM\SDDDSM>datapath query wwpn
Adapter Name      PortWWN
Scsi Port 7       100000051EC76B89
Scsi Port 7       100000051EC76B8A
```

If you need more detailed information about SDDDSM, see *Multipath Subsystem Device Driver User's Guide*, GC52-1309. A PDF version of the *Multipath Subsystem Device Driver User's Guide*, GC52-1309, is at the following web page:

<https://ibm.biz/BdEqeZ>

The Windows host was prepared to connect to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, and you know the WWPNs of the host. The next step is to configure a host object for the WWPNs by using the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI. This task is explained in 5.5.1, “Creating Fibre Channel hosts” on page 228.

SAN Boot hosts are beyond the intended scope of this book. For more information about that topic, search for SAN Boot on the following web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/lenovo_vseries.html

Windows 2003: The examples focus on Windows 2008 R2 and Windows 2012, but the procedure for Windows 2003 is similar. If you use Windows 2003, do not forget to install Microsoft hotfix 908980. If you do not install it before you perform this procedure, preferred pathing is not available. You can download this hotfix from the following address:

<http://support.microsoft.com/kb/908980>

5.3.2 Windows 2008 R2 and Windows 2012 R2: Preparing for iSCSI attachment

This section details iSCSI attachment.

Installing and updating supported HBAs

Install a supported HBA model with the latest supported firmware and drivers for your configuration. The latest supported HBAs and levels for Windows 2008 R2 and 2012 R2 are available at interoperability matrix in the following web pages:

<https://datacentersupport.lenovo.com/us/en/products/storage/lenovo-storage/v3700v2/6535/documentation>

<https://datacentersupport.lenovo.com/us/en/products/storage/lenovo-storage/v5030/6536/documentation>

Install the driver by using Windows Device Manager or vendor tools. Also, check and update the firmware level of the HBA by using the manufacturer's provided tools. Always check the readme file to see whether any Windows registry parameters must be set for the HBA driver.

Important: For converged network adapters (CNAs), which can support both FC and iSCSI, it is important to ensure that the Ethernet networking driver is installed in addition to the FCoE driver. You are required to install the Ethernet networking driver and the FCoE driver before you configure iSCSI.

If you use a hardware iSCSI HBA, refer to the manufacturer's documentation and the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Lenovo Information Center for further details and the latest information about the hardware and host OS configuration. The following section describes how to configure iSCSI by using the software initiator.

In Windows 2008 R2 and 2012, the Microsoft iSCSI software initiator is preinstalled.

Complete the following steps:

1. Enter `iscsi` in the search field of the Windows 2008 R2 Start window (Figure 5-9) and click **iSCSI Initiator**.

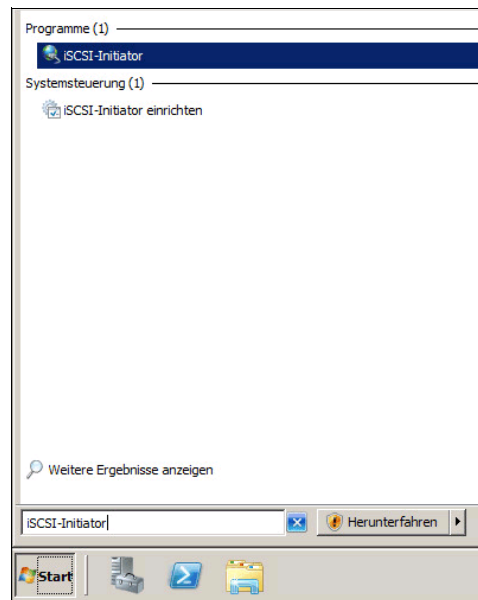


Figure 5-9 Windows 2008 R2 iSCSI Initiator

2. For Windows 2012 R2, go to the all programs menu and enter `iSCSI` in the search field at the top of the window. See Figure 5-10.



Figure 5-10 iSCSI Initiator Windows 2012 R2

3. Confirm the automatic start of the iSCSI service (Figure 5-11).

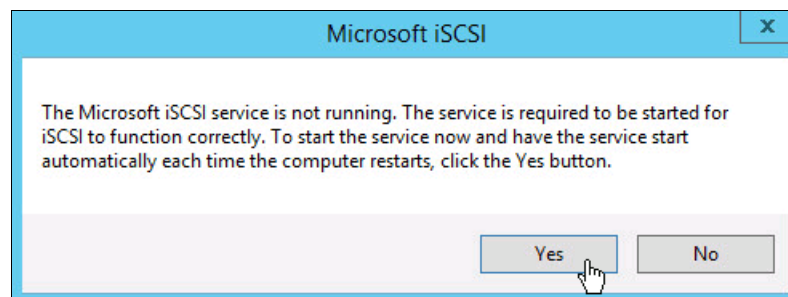


Figure 5-11 Automatic start of the iSCSI service

4. The iSCSI Initiator Properties window opens. Select the **Configuration** tab (Figure 5-12 on page 200). Write down the initiator name of your Windows host because you use it later.

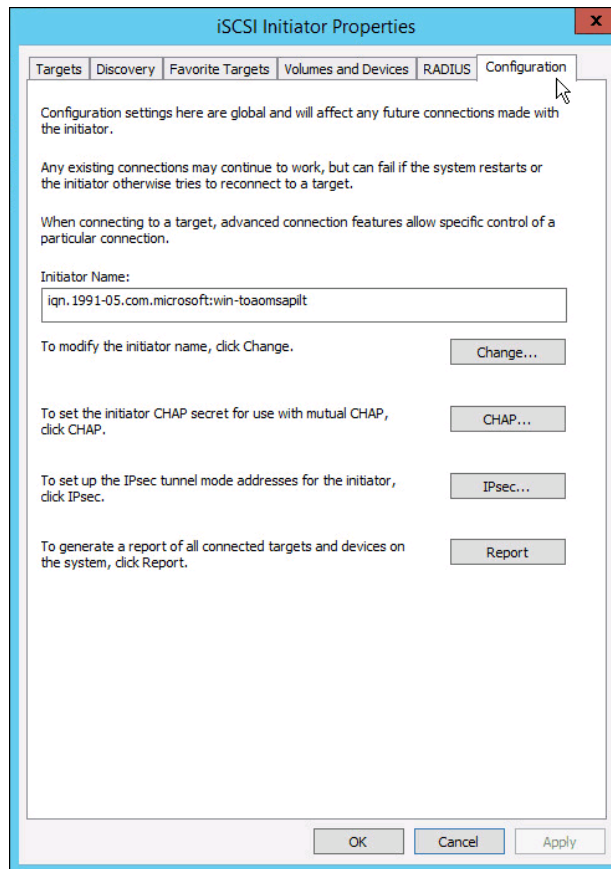


Figure 5-12 iSCSI Initiator Properties window

5. You can change the initiator name, or enable advanced authentication, but these actions are out of the scope of our basic setup. By default, iSCSI authentication is not enabled. More details are available at the Lenovo Information Center for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 at the following address:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/lenovo_v_series.html

6. From the Configuration tab, you can change the initiator name, enable CHAP authentication, and more. However, these tasks are beyond the scope of our basic setup. CHAP authentication is disabled, by default. For more information, see the *Microsoft iSCSI Initiator Step-by-Step Guide*:

<http://technet.microsoft.com/en-us/library/ee338476%28v=ws.10%29.aspx>

Configuring Ethernet ports

We suggest that you use separate dedicated ports for host management and iSCSI. In this case, we need to configure IPs on the iSCSI ports in the same subnet and virtual LAN (VLAN) as the external storage that we want to attach to.

To configure Ethernet port IPs on Windows 2008 R2 and 2012 R2, complete the following steps:

1. Go to **Control Panel** → **Network and Internet** → **Network and Sharing Center**. The window that is shown in Figure 5-13 opens.

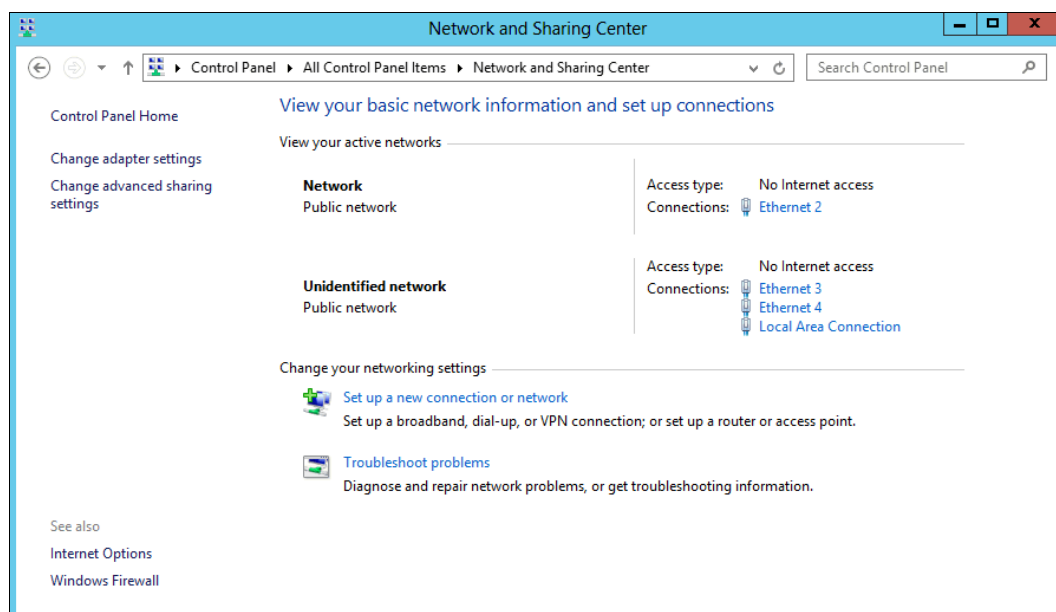


Figure 5-13 Network and Sharing Center in Windows 2012 R2

In this case, two networks are visible to the system. We use the first network to connect to the server. It consists of a single dedicated Ethernet port for management. The second network is our iSCSI network. It consists of two dedicated Ethernet ports for iSCSI. We suggest that you use at least two dedicated ports for failover purposes.

2. To configure an IP address, click one of the iSCSI Ethernet connections (in this case, Ethernet 3 or Ethernet 4). Figure 5-14 shows the window that displays the Ethernet status.




Figure 5-14 Ethernet status

3. To configure the IP address, click **Properties** (Figure 5-15).

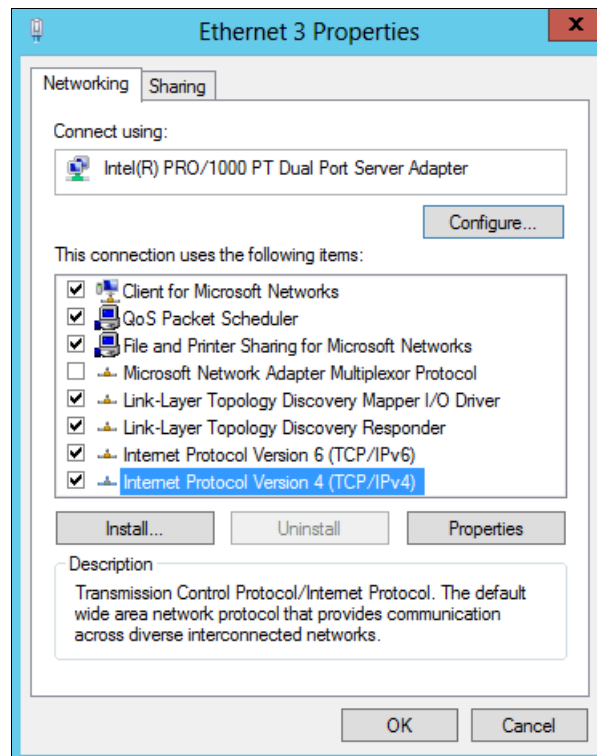


Figure 5-15 Ethernet properties

4. If you use IPv6, select **Internet Protocol Version 6 (TCP/IPv6)** and click **Properties**. Otherwise, select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties** to configure an IPv4 address.
5. For IPv4, the window that is shown in Figure 5-16 opens. To manually set the IP, select “Use the following address” and enter an IP address, subnet mask, and gateway. Set the DNS server address, if required. Click **OK** to confirm.

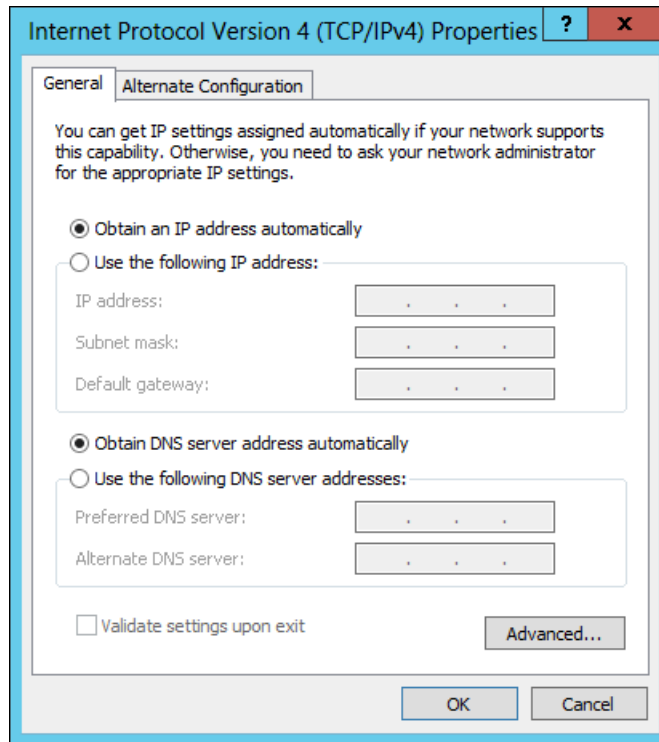


Figure 5-16 Configuring an IPv4 address in Windows 2012 R2

6. Repeat the previous steps for each port that you want to configure for iSCSI attachment.

The Ethernet ports are now prepared for iSCSI attachment.

Setting the Windows registry keys

Modify the system registry so that your iSCSI operations are more reliable:

1. In the search field of the Windows Start window, type **regedit** and click **regedit.exe**.
2. In the registry editor, locate the following key:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\LinkDownTime`
 Confirm that the value for the **LinkDownTime** key is 120 (decimal value) and, if not, change the value to 120.
3. In the registry editor, locate the following key:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\MaxRequestHoldTime`
 Confirm that the value for the **MaxRequestHoldTime** key is 120 (decimal value) and, if not, change the value to 120.

4. In the registry editor, locate the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\MaxPendingRequests

Confirm that the value for the MaxPendingRequests key is 2048 (decimal value) and, if not, change the value to 2048.

5. In the registry editor, locate the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk\TimeOutValue

Confirm that the value for the TimeOutValue key is 60 (decimal value) and, if not, change the value to 60.

6. Restart your host for these changes to take effect.

Multipath support for iSCSI on Windows

For multipathing with iSCSI, we need to enable Microsoft Multipath Input/Output (MPIO). See “Installing the multipathing software” on page 194 for instructions to enable MPIO.

Important: Subsystem Device Driver DSM (SDDDSM) is **not** supported for iSCSI attachment. Do not follow the steps to install SDDDSM that you follow to install FC or SAS.

These basic steps are to prepare a Windows 2008 R2 or Windows 2012 R2 host for iSCSI attachment. To configure Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 for iSCSI connections, see 5.5.3, “Creating iSCSI hosts” on page 236.

5.3.3 Windows 2012 R2: Preparing for SAS attachment

This procedure is described in the following sections.

Installing and updating supported HBAs

Install a supported SAS HBA with the latest supported firmware and drivers for your configuration. A list of the latest supported HBAs and levels for Windows 2008 R2, Windows 2012 R2, and other operating systems is available at the interoperability matrix of the following link:

<https://datacentersupport.lenovo.com/us/en/products/storage/lenovo-storage/v3700v2/6535/documentation>

<https://datacentersupport.lenovo.com/us/en/products/storage/lenovo-storage/v5030/6536/documentation>

Install the driver by using Windows Device Manager or vendor tools. Also, check and update the firmware level of the HBA by using the manufacturer’s provided tools. Always check the readme file to see whether any Windows registry parameters must be set for the HBA driver.

Determining host WWPNS

The worldwide port names (WWPNs) of the SAS HBA are required to configure host attachment on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

You can obtain the host WWPNs through vendor tools or the HBA BIOS. However, the easiest way is to connect the SAS cables to the ports on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, log on to the Storwize CLI through Secure Shell (SSH), and run **svcinfo lssasportcandidate**, as shown in Example 5-5.

Example 5-5 Finding host WWPNs

```
IBM_Storwize:ITS0_V5000:superuser>svcinfo lssasportcandidate
sas_wwpn
500062B200556140
500062B200556141
```

Configuring SAS HBAs on Windows

We suggest these settings:

- ▶ I/O Timeout for Block Devices: 10
- ▶ I/O Timeout for Sequential Devices: 10
- ▶ I/O Timeout for Other Devices: 10
- ▶ LUNs to Scan for Block Devices: All
- ▶ LUNs to Scan for Sequential Devices: All
- ▶ LUNs to Scan for Other Devices: All

Multipath support for SAS on Windows

For multipathing with SAS, we need to enable Microsoft Multipath Input/Output (MPIO) and install Lenovo Subsystem Device Driver DSM (SDDDSM). For instructions, see “Installing the multipathing software” on page 194.

We described the basic steps to prepare a Windows 2008 R2 and 2012 R2 host for SAS attachment. For information about configuring SAS attachment on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 side, see 5.5.5, “Creating SAS hosts” on page 243.

5.3.4 VMware ESXi: Preparing for Fibre Channel attachment

Complete the following steps to prepare a VMware ESXi host to connect to an Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 by using Fibre Channel:

1. Install the HBA or HBAs on the ESXi server.
2. Ensure that the current firmware levels are applied on your host system.
3. Update and configure the HBA for hosts that are running ESXi.
4. Connect the FC host adapter ports to the switches.
5. Configure the switches (zoning).
6. Install the VMware ESXi Hypervisor and load additional drivers, if required.

Downloading and installing the supported firmware

Install the current firmware levels to your host server. For the HBAs, check the interoperability matrix at the following address:

<https://datacentersupport.lenovo.com/us/en/products/storage/lenovo-storage/v3700v2/6535/documentation>

<https://datacentersupport.lenovo.com/us/en/products/storage/lenovo-storage/v5030/6536/documentation>

Download the current supported HBA firmware for your configuration and apply it to your system. Certain HBAs and especially the new converged network adapters (CNAs) require an additional driver to be loaded into ESXi. Check the VMware Compatibility Guide to see whether any requirements exist for your configuration by going to the following address:

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

Configuring QLogic HBAs for VMware ESXi

This section applies to ESXi hosts with installed QLogic HBAs. After you install the firmware, you must configure the HBAs. To perform this task, either use the QCC software, or use the HBA BIOS, load the adapter defaults, and set the following values:

- ▶ Host adapter settings:
 - Host Adapter BIOS: Disabled (unless the host is configured for SAN Boot)
 - Frame size: 2048
 - Loop Reset Delay: 5 (minimum)
 - Adapter Hard Loop ID: Disabled
 - Hard Loop ID: 0
 - Spinup Delay: Disabled
 - Connection Options 1: Point to point only
 - Fibre Channel Tape Support: Disabled
 - Data Rate: 2
- ▶ Advanced adapter settings:
 - Execution throttle: 100
 - LUNs per Target: 0
 - Enable LIP Reset: No
 - Enable LIP Full Login: Yes
 - Enable Target Reset: Yes
 - Login Retry Count: 8
 - Link Down Timeout: 10
 - Command Timeout: 20
 - Extended event logging: Disabled (Enable it for debugging only.)
 - RIO Operation Mode: 0
 - Interrupt Delay Timer: 0

The QCC management software delivers a unified web-based single-pane-of-glass management console across the QLogic family of storage and networking adapters. A graphical user interface (GUI) or command-line interface (CLI) is available. A VMware vCenter plug-in is also available. You can obtain the QCC for Windows at the following web page:

<https://ibm.biz/BdEfsj>

Configuring Emulex HBAs for VMware ESXi

This section applies to ESXi hosts with installed Emulex HBAs. After you install the firmware, load the default settings of all of your adapters that are installed on the host system, and ensure that the Adapter BIOS is disabled, unless you use SAN Boot.

VMware ESXi installation

To install VMware ESXi, complete the following steps:

1. Install your VMware ESXi server and load any additional drivers and patches, if required. If you are not familiar with the procedure, see the installation guide at the following address:

<https://www.vmware.com/support/pubs/>

2. After you complete your ESXi installation, connect to your ESXi server by using the vSphere web client and go to the **Configuration** tab.
3. Click **Storage Adapters**, and scroll down to your FC HBAs (Figure 5-17). Document the WWPNs of the installed adapters for later use.

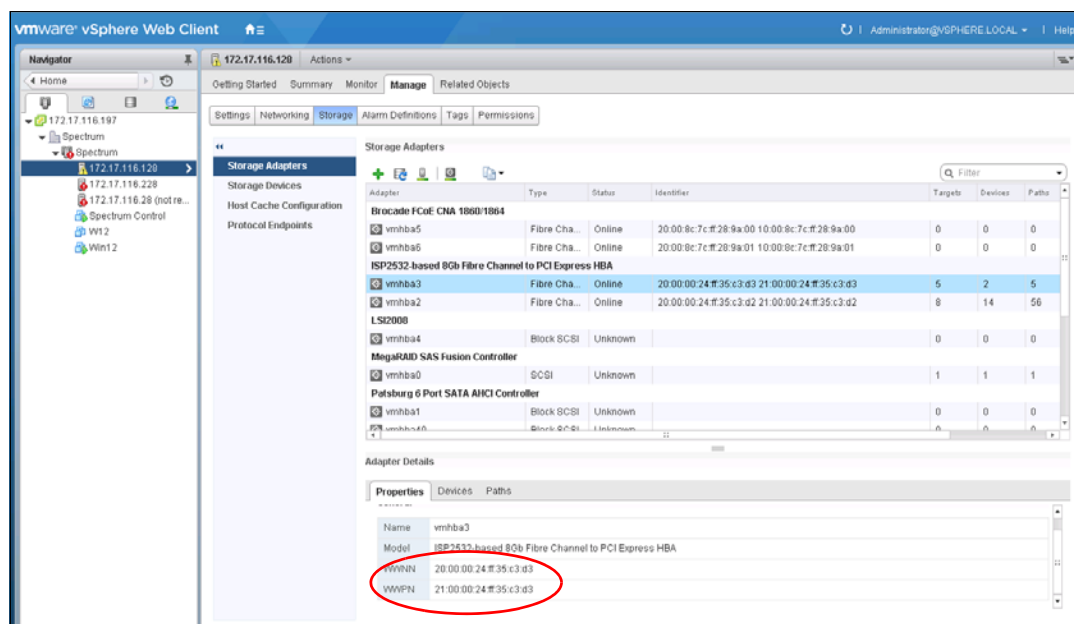


Figure 5-17 Show WWPNs in VMware ESXi

VMware ESXi multipathing

The ESXi server has its own multipathing software. You do not need to install a multipathing driver, either on the ESXi server or on the guest operating systems. The ESXi multipathing policy supports several operating modes:

- ▶ Round Robin
- ▶ Fixed
- ▶ Most Recently Used (MRU)

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are an active/active storage device. The suggested multipathing policy is *Round Robin*. Round Robin performs static load balancing for I/O. If you do not want the I/O balanced over all available paths, the *Fixed* policy is supported also. This policy setting can be selected for every volume.

Set this policy after you attach the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 LUNs to the ESXi host. For more information, see Chapter 6, "Volume configuration" on page 269. If you use an older version of VMware ESX (up to version 3.5), *Fixed* is the suggested policy setting.

MRU selects the first working path, which is discovered at system start time. If this path becomes unavailable, the ESXi/ESX host switches to an alternative path and continues to use the new path while it is available. This policy is the default policy for LUNs that are presented from an Active/Passive array. ESXi/ESX does not return to the previous path if, or when, it returns. It remains on the working path until it, for any reason, fails.

Determining host WWPNs

The worldwide port names (WWPNs) of the FC HBA are required to correctly zone switches and configure host attachment on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. On VMware ESXi, you can obtain these WWPNs through the VMware vSphere Client.

Note: Beginning with VMware ESXi version 5.5, certain new features can be accessed through the vSphere Web Client only. However, we do not demonstrate any of these features. All of the following examples continue to focus on the use of the desktop client.

Connect to the ESXi server (or VMware vCenter) by using the VMware vSphere Client and browse to the **Configuration** tab. Click **Storage Adapters** to see the HBA WWPNs, as shown in Figure 5-18.

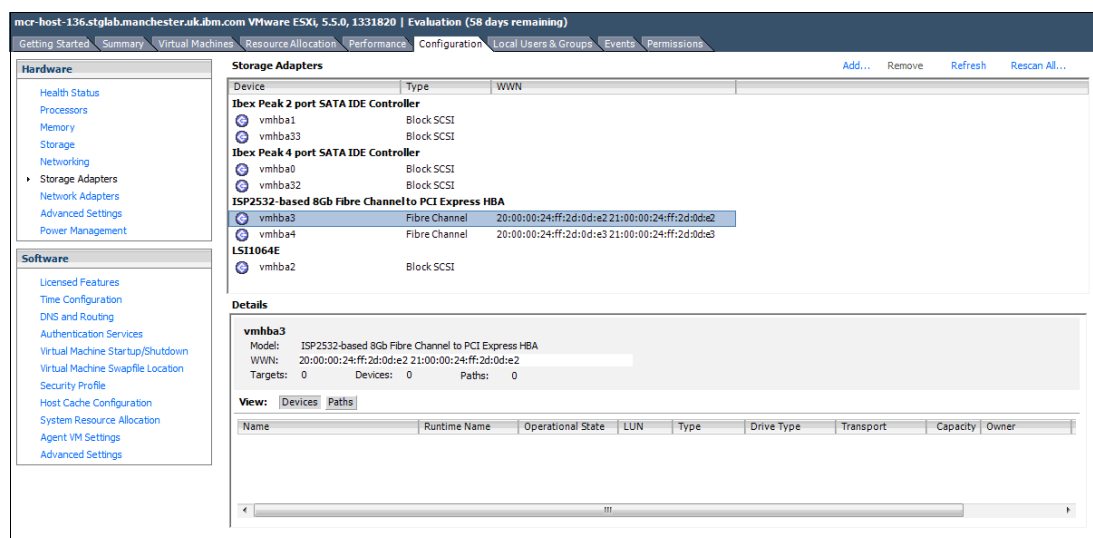


Figure 5-18 FC WWPNs in VMware vSphere Client

After all of these steps are completed, the ESXi host is prepared to connect to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. Go to 5.5.1, “Creating Fibre Channel hosts” on page 228 to create the ESXi FC host in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI.

5.3.5 VMware ESXi: Preparing for iSCSI attachment

This section describes how to enable iSCSI on VMware ESXi hosts. We focus on vSphere because the complete iSCSI stack was rewritten in this level. This level offers improved performance and supports useful features, such as jumbo frames and Transmission Control Protocol (TCP) Segmentation Offload. We focus on the basic ESXi iSCSI setup. More detailed information is provided in the VMware *vSphere Documentation Center*, which is available on the following web pages, depending on which version you use:

<https://docs.vmware.com/en/VMware-vSphere/index.html>

For more information, see the VMware Compatibility Guide:

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

Important: For converged network adapters (CNAs) that support both FC and iSCSI, it is important to ensure that the Ethernet networking driver is installed in addition to the FCoE driver. The Ethernet networking driver and the FCoE driver are required for the configuration of iSCSI.

For more information about the hardware and host OS configuration, if you use a hardware iSCSI HBA, see the manufacturer's documentation and the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Lenovo Information Center. The following section describes how to configure iSCSI by using the software initiator.

Complete the following steps to prepare a VMware ESXi host to connect to a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 by using iSCSI:

1. Ensure that the current firmware levels are applied on your host system.
2. Install VMware ESXi and load additional drivers if required.
3. Connect the ESXi server to your network. You need to use separate network interfaces for iSCSI traffic.
4. Configure your network to fulfill your security and performance requirements.

The iSCSI initiator is installed by default on your ESXi server, but you must enable it. To enable it, complete the following steps:

1. Connect to your ESXi server by using the vSphere Client. Go to **Manage**, and select **Networking** (Figure 5-19).




Figure 5-19 Select VMware networking

2. Click **Add Networking** to start the Add Networking wizard (Figure 5-20). Select **VMkernel Network Adapter** and click **Next**.

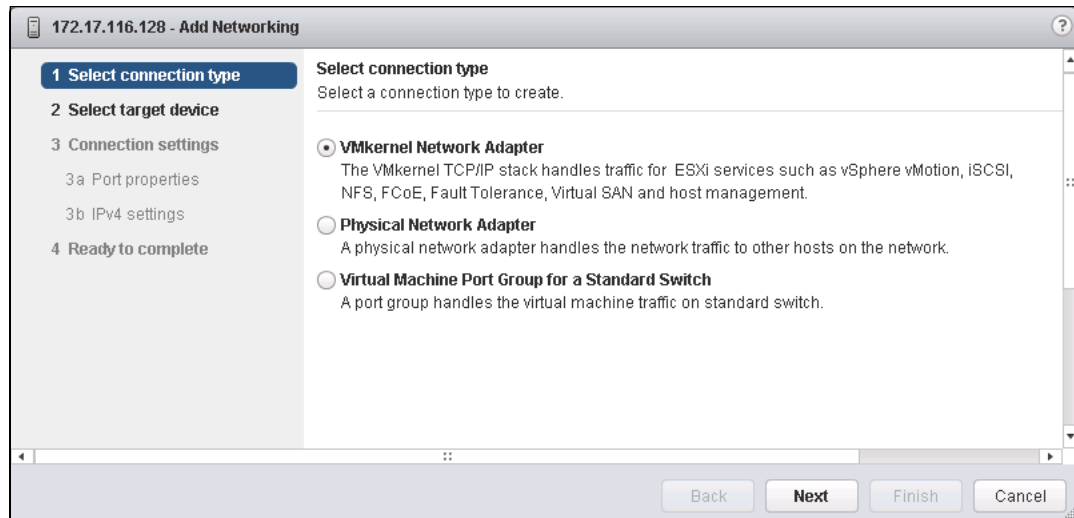


Figure 5-20 VMware: Add Networking wizard

3. Click **Select target device**, as shown in Figure 5-21.

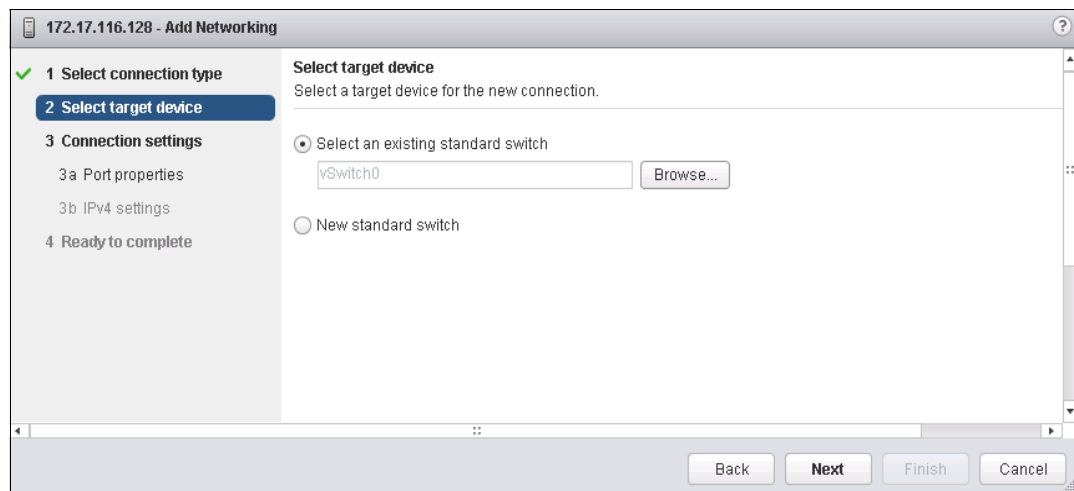


Figure 5-21 Select target device

4. Select one or more network interfaces that you want to use for iSCSI traffic and click **Next** (Figure 5-22).

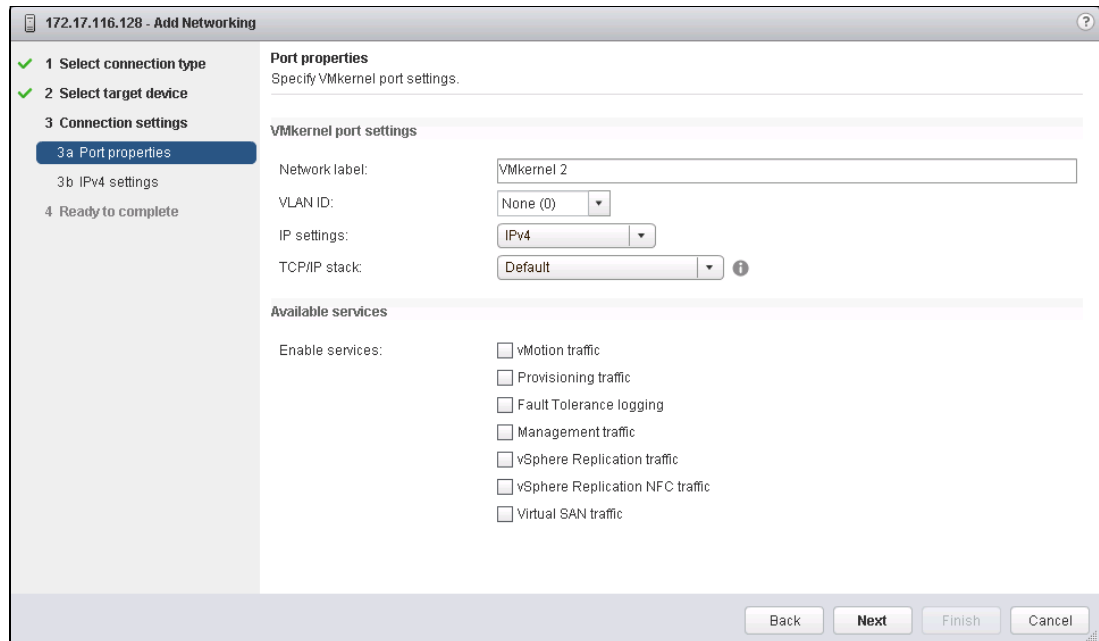


Figure 5-22 VMware: Select an iSCSI interface

5. Enter a meaningful network label and click **Next** (Figure 5-23).

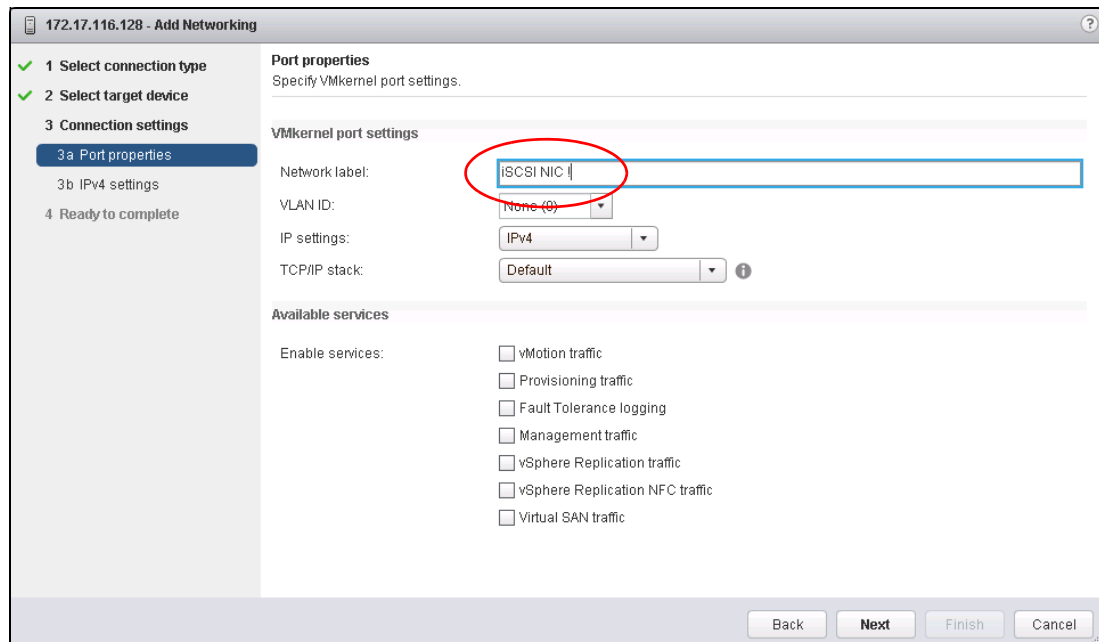


Figure 5-23 VMware: Enter a network label

6. Enter an IP address for your iSCSI network. Use a dedicated network for iSCSI traffic (Figure 5-24).




Figure 5-24 VMware: Enter an iSCSI network IP

7. Click **Next**, as shown in Figure 5-25.

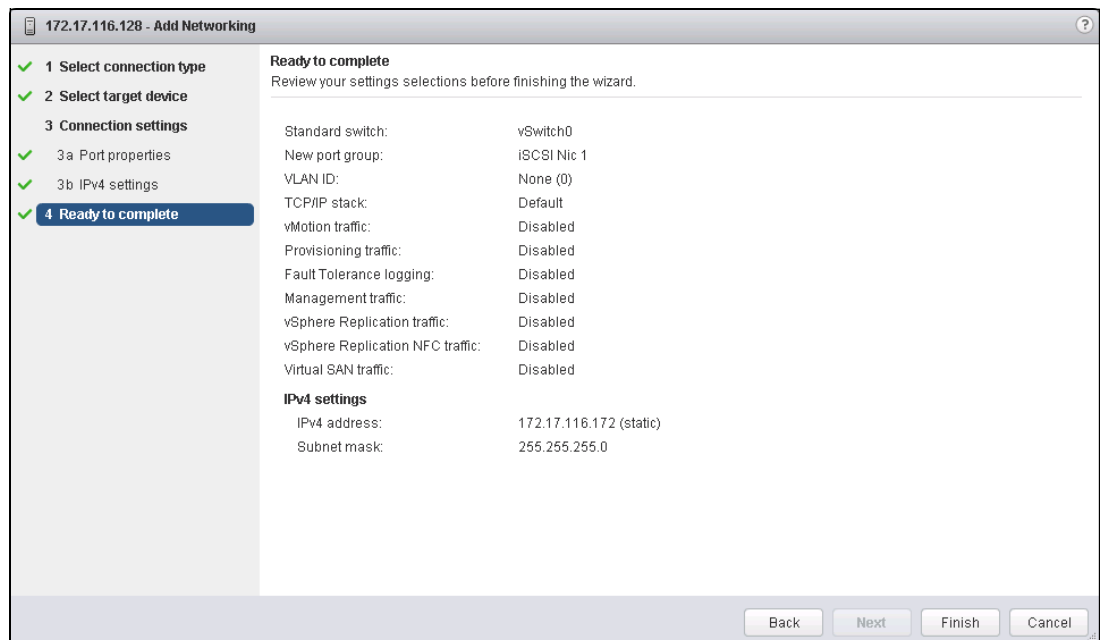


Figure 5-25 VMware: Ready to complete

8. Click **Finish** to complete the setup.

- Check whether an iSCSI software adapter is available. Select **Storage Adapters** on the Manage tab. See Figure 5-26.

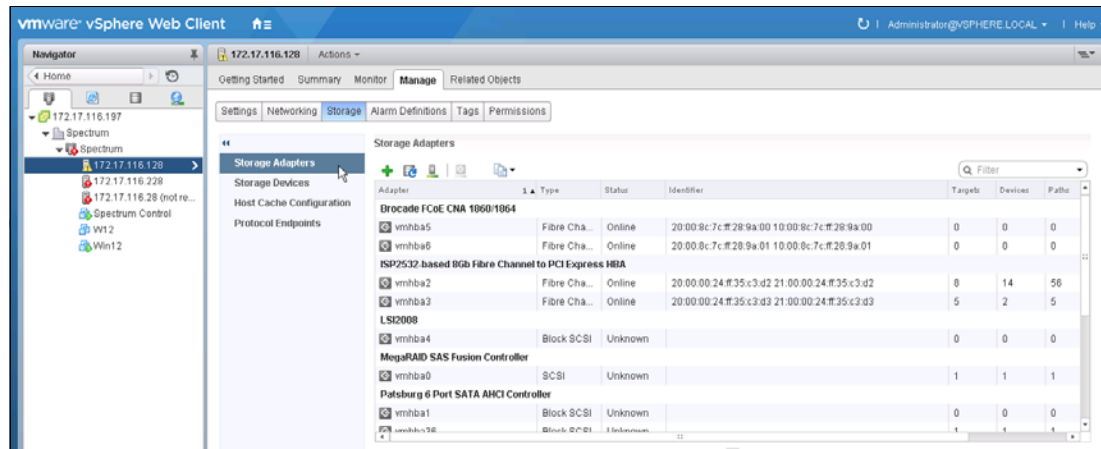


Figure 5-26 VMware: Select new iSCSI software adapter

- Click the plus sign (+) to add an iSCSI software adapter. See Figure 5-27.

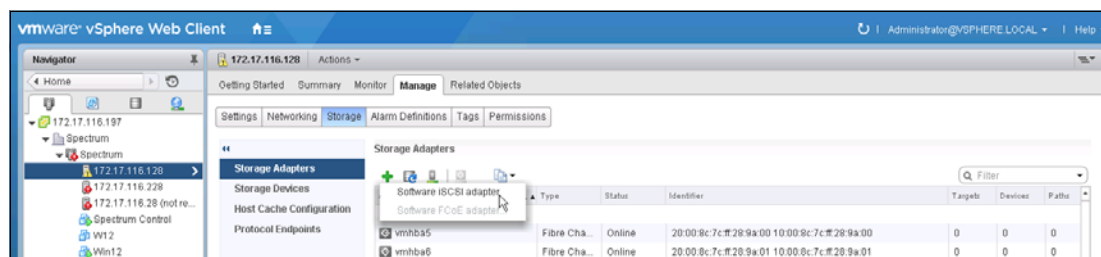


Figure 5-27 VMware: Add an iSCSI software adapter

- The Add Software iSCSI Adapter window opens. See Figure 5-28.




Figure 5-28 VMware: Add Software iSCSI Adapter window

- Click **OK**. A message displays that prompts you to configure the adapter after it is added.

13. A new iSCSI adapter is added to the Storage Adapters window. See Figure 5-29.

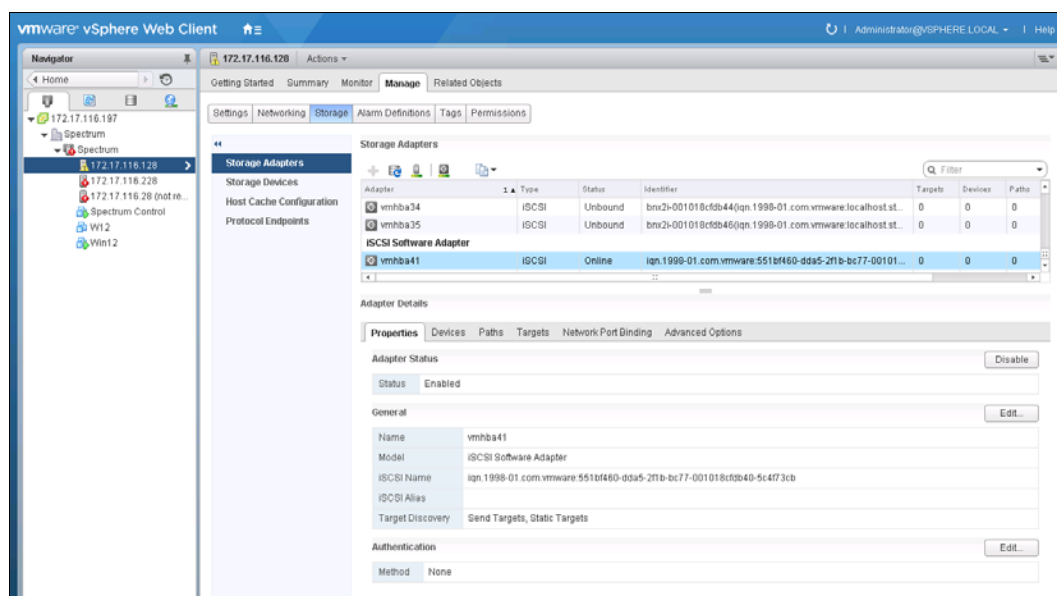


Figure 5-29 VMware: New iSCSI Software Adapter

14. Select **Storage Adapters** and scroll to the iSCSI Software Adapter (Figure 5-30). Highlight it and you see the Adapter Details in the lower part of the window.



Figure 5-30 VMware: iSCSI Software Adapter

15. The iSCSI Software Adapter Properties window opens. Figure 5-31 shows that the initiator is enabled by default. To change this setting, click **Disable**.




Figure 5-31 VMware: iSCSI Software Adapter properties

16.The VMware ESX iSCSI initiator is successfully enabled (Figure 5-32). Document the initiator name for later use.




Figure 5-32 Enabled VMware iSCSI initiator

Multipath support for iSCSI on VMware ESXi

As explained in 5.3.4, “VMware ESXi: Preparing for Fibre Channel attachment” on page 205, the ESXi server uses its own multipathing software.

For iSCSI, extra configuration is required in the VMkernel port properties to enable path failover. Each VMkernel port must map to one physical adapter port, which is not the default setting. Complete the following steps:

- 1. Browse to the **Configuration** tab and select **Networking**. Click **Properties** next to the vSwitch that you configured for iSCSI to open the window that is shown in Figure 5-33.




Figure 5-33 View the vSwitch properties with listed VMkernel ports

2. Select one of the VMkernel ports and click **Edit**. The window that is shown in Figure 5-34 opens.

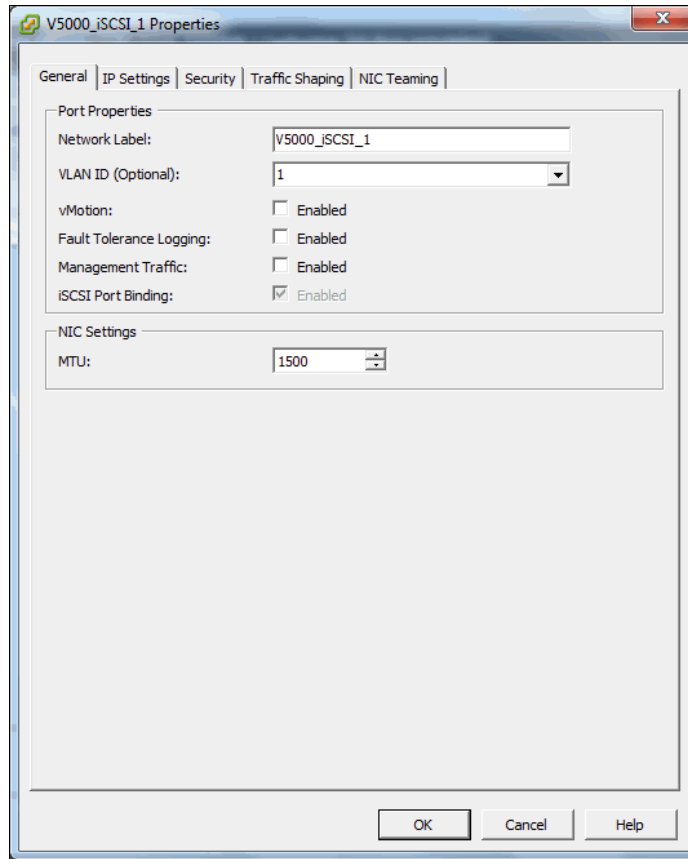


Figure 5-34 Editing a VMkernel port

3. Click the **NIC Teaming** tab. Select **Override switch failover order** and ensure that each port is tied to one physical adapter port, as shown in Figure 5-35.

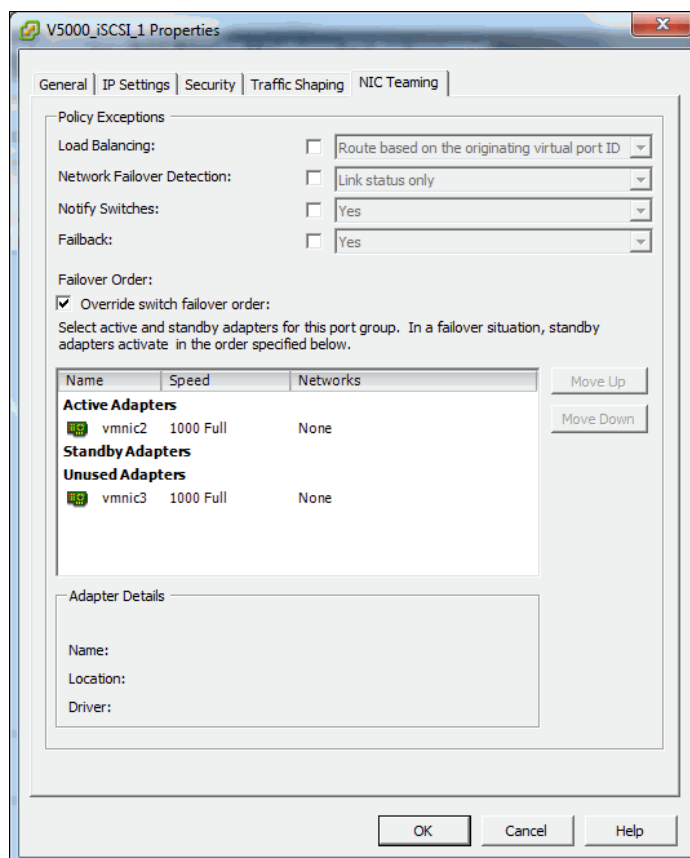


Figure 5-35 Configuring a VMkernel port to bind to a single physical adapter port

These basic steps are required to prepare a VMware ESXi host for iSCSI attachment. For information about configuring iSCSI attachment on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, see 5.5.3, “Creating iSCSI hosts” on page 236.

For more information about configuring iSCSI attachment on the VMware ESXi side, the following white paper, which was published by VMware, is a useful resource:

<http://ibm.biz/Bd4ND6>

5.3.6 VMware ESXi: Preparing for SAS attachment

This procedure is described in the following sections.

Installing and updating supported HBAs

Install a supported HBA with the latest supported firmware and drivers for your configuration. A list of the latest supported HBAs and levels for VMware ESXi is available at the interoperability matrix:

<https://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v3700v2/6535/documentation>

<https://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v5030/6536/documentation>

Install the driver by using VMware vSphere Client, the ESXi CLI, or vendor tools. Also, check and update the firmware level of the HBA by using the manufacturer's provided tools. Always check the readme file to see whether more configuration is required for the HBA driver.

For more information, see the VMware Compatibility Guide:

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

Configuring SAS HBAs on VMware ESXi

In this example, we used an LSI 9207-8e card and did not need to configure HBA parameters beyond the default settings. We advise that you check the parameters through the HBA BIOS or vendor tools to confirm that they are suitable for your requirements.

Multipath support for SAS on VMware ESXi

As with FC, we can use native ESXi multipathing for SAS attachment on VMware ESXi 5.5. For more information, see 5.3.4, "VMware ESXi: Preparing for Fibre Channel attachment" on page 205.

Determining host WWPNS

The worldwide port names (WWPNs) of the SAS HBA are required to configure host attachment on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030..

The host WWPNS are not directly available through VMware vSphere Client. However, you can obtain them by using vendor tools or the HBA BIOS. The method that is described in 5.3.3, "Windows 2012 R2: Preparing for SAS attachment" on page 204 also works.

These basic steps are required to prepare a VMware ESXi host for SAS attachment. For information about configuring SAS attachment on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. side, see 5.5.5, "Creating SAS hosts" on page 243.

For more information and guidance about attaching storage with VMware ESXi, the following document, which was published by VMware, is a useful resource:

<http://ibm.biz/Bd4ND5>

5.4 N-Port Virtualization ID (NPIV) Support

The usage model for all Lenovo storage V-series products is based around two-way active/active node models. That is a pair of distinct control modules that share active/active access for a given volume. These nodes each have their own Fibre Channel WWNN, and thus all ports that are presented from each node have a set of WWPNS that are presented to the fabric.

Traditionally, should one node fail or be removed for some reason, the paths that are presented for volumes from that node would go offline, and it is up to the native OS multipathing software to failover from using both sets of WWPN to just those that remain online. While this is exactly what multipathing software is designed to do, occasionally it can be problematic, particularly if paths are not seen as coming back online for some reason.

Starting with V7.7.0, Lenovo storage V-series system can be enabled into NPIV mode. When NPIV mode is enabled on the Lenovo storage V-series system, ports do not come up until they are ready to service I/O, which improves host behavior around node unpends. In addition, path failures due to an offline node are masked from host multipathing.

When NPIV is enabled on Lenovo storage V-series system nodes, each physical WWPN reports up to three virtual WWPNs, as shown in Table 5-1.

Table 5-1 Spectrum Virtualize NPIV Ports's

NPIV port	Port description
Primary NPIV Port	This is the WWPN that communicates with backend storage, and might be used for node to node traffic. (Local or remote.)
Primary Host Attach Port	This is the WWPN that communicates with hosts. It is a target port only, and this is the primary port, so it represents this local node's WWNN.
Failover Host Attach Port	This is a standby WWPN that communicates with hosts and is only brought online on this node if the partner node in this I/O Group goes offline. This is the same as the Primary Host Attach WWPN on the partner node.

Figure 5-36 depicts the three WWPNs associated with an SVC port when NPIV is enabled.




Figure 5-36 Allocation of NPIV virtual WWPN ports per physical port

The failover host attach port (in pink) is not active at this time. Figure 5-37 on page 220 shows what happens when the second node fails. Subsequent to the node failure, the failover host attach ports on the remaining node are active and have taken on the WWPN of the failed node's primary host attach port.

Note: Figure 5-37 on page 220 shows only two ports per node in detail, but the same applies for all physical ports.

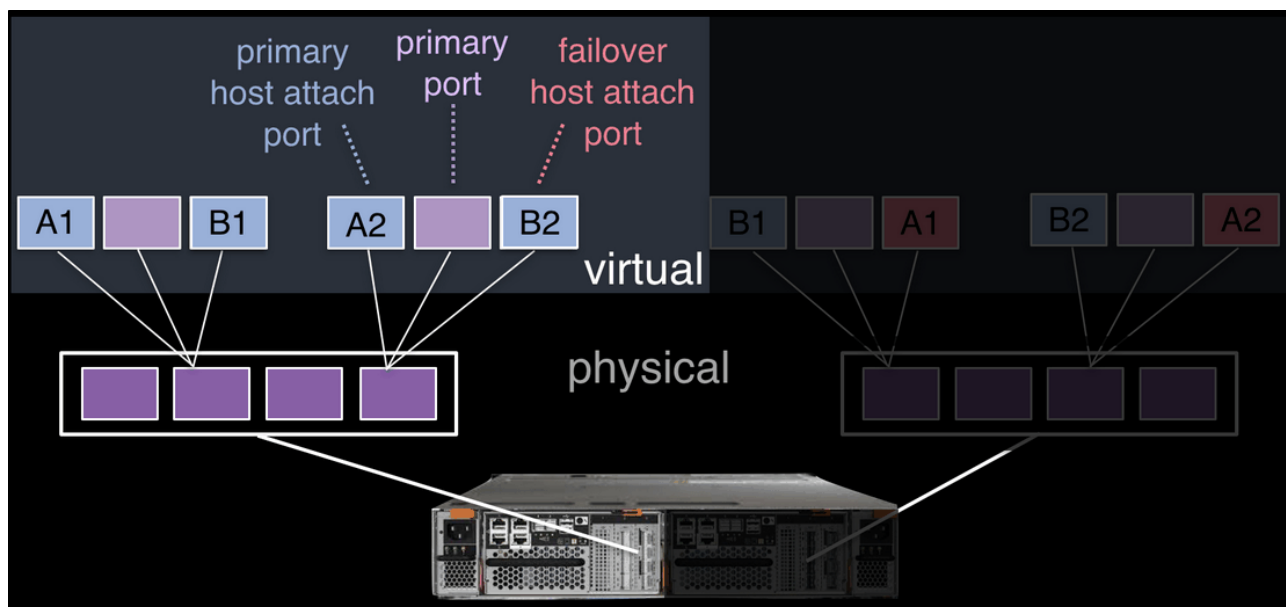


Figure 5-37 Allocation of NPIV virtual WWPN ports per physical port after a node failure

With V7.7.0 onwards, this all happens automatically when NPIV is enabled at a system level . At this time, the failover only happens automatically between the two nodes in an I/O Group.

There is a transitional mode for compatibility with an earlier version during the transition period as well.

The processes for enabling NPIV on a new Lenovo storage V-series system is slightly different than on an existing system. For more information, see Lenovo Information Center: http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_icconfiguringnpiv.html

Note: NPIV is only supported for Fibre Channel protocol. It is not supported for FCoE protocol or iSCSI.

5.4.1 NPIV Prerequisites

The following key points should be considered for NPIV enablement:

- ▶ For NPIV enablement, the system must be at version 7.7.0 or later.
- ▶ A version 7.7.0 or later system with NPIV enabled as backend storage for a system that is earlier than version 7.7.0 is not supported.
- ▶ Both nodes in an IO group should have identical hardware to allow failover to work as expected.
- ▶ Fibre Channel switches must permit each physically connected system port the ability to create two additional NPIV ports.
- ▶ Check the configuration and restriction webpage at <https://ibm.biz/BdJSJx> to ensure your operating environment is supported for NPIV.

5.4.2 Enabling NPIV on a new system

For v7.7.0 and later, a new system should have NPIV enabled by default. For any case where it is not enabled by default and NPIV is wanted, then NPIV can be enabled on a new system by running the following steps:

1. Run the **lsiogrp** command to list the I/O groups present in a system, as shown in Example 5-6.

Example 5-6 Listing the I/O groups in a system

```
IBM_2145:ITS0:superuser>lsiogrp
id name          node_count vdisk_count host_count site_id site_name
0  io_grp0        2          2          2          0
1  io_grp1        0          0          2          0
2  io_grp2        0          0          2          0
3  io_grp3        0          0          2          0
4  recovery_io_grp 0          0          0          0
```

2. Run the **lsiogrp** command to view the status of N_Port ID Virtualization (NPIV), as shown in Example 5-7.

Example 5-7 Checking NPIV mode with the fctargetportmode field

```
IBM_2145:ITS0:superuser>lsiogrp 0
id 0
name io_grp0
node_count 2
vdisk_count 2
host_count 2
flash_copy_total_memory 20.0MB
flash_copy_free_memory 20.0MB
remote_copy_total_memory 20.0MB
remote_copy_free_memory 20.0MB
mirroring_total_memory 20.0MB
mirroring_free_memory 20.0MB
raid_total_memory 40.0MB
raid_free_memory 38.8MB
maintenance no
compression_active yes
accessible_vdisk_count 2
compression_supported yes
max_enclosures 10
encryption_supported no
flash_copy_maximum_memory 552.0MB
site_id
site_name
fctargetportmode enabled
compression_total_memory 2047.9MB
```

3. If the resulting output is `fctargetportmode enabled`, as shown in Example 5-7, then NPIV is enabled.
4. The virtual WWPNs can be listed using the **lstargetportfc** command, as shown in Example 5-8 on page 222.

Example 5-8 Listing the virtual WWPNs

```
IBM_2145:ITS0:superuser>lstargetportfc
```

id	WWPN	WWNN	port_id	owning_node_id	current_node_id	nportid	host_io_permitted
virtualized							
1	50050768021000EF	50050768020000EF	1	1	1	010200	no
2	50050768029900EF	50050768020000EF	1	1	1	010201	yes
3	50050768022000EF	50050768020000EF	2	1	1	020200	no
4	5005076802A900EF	50050768020000EF	2	1	1	020201	yes
5	50050768023000EF	50050768020000EF	3	1	1	0A83C0	no
6	5005076802B900EF	50050768020000EF	3	1	1	0A83C1	yes
7	50050768024000EF	50050768020000EF	4	1	1	011400	no
8	5005076802C900EF	50050768020000EF	4	1	1	011401	yes
33	50050768021000F0	50050768020000F0	1	2	2	010300	no
34	50050768029900F0	50050768020000F0	1	2	2	010301	yes
35	50050768022000F0	50050768020000F0	2	2	2	020300	no
36	5005076802A900F0	50050768020000F0	2	2	2	020301	yes
37	50050768023000F0	50050768020000F0	3	2	2	011500	no
38	5005076802B900F0	50050768020000F0	3	2	2	011501	yes
39	50050768024000F0	50050768020000F0	4	2	2	0A82C0	no
40	5005076802C900F0	50050768020000F0	4	2	2	0A82C1	yes

5. At this point you can configure zones for hosts using the primary host attach ports (virtual WWPNs) of the Lenovo storage V-series ports, as shown in **bold** in the output of Example 5-8.
6. If the status of `fctargetportmode` is disabled, run the **chiogrp** command to get into transitional mode for NPIV, as shown in Example 5-9.

Example 5-9 Change the NPIV mode to transitional

```
IBM_2145:ITS0:superuser>chiogrp -fctargetportmode transitional 0
```

7. The transitional mode can be verified using the **lsiogrp** command, as shown in Example 5-10.

Example 5-10 NPIV transitional mode

```
IBM_2145:ITS0:superuser>lsiogrp 0
id 0
name io_grp0
node_count 2
vdisk_count 2
host_count 2
flash_copy_total_memory 20.0MB
flash_copy_free_memory 20.0MB
remote_copy_total_memory 20.0MB
remote_copy_free_memory 20.0MB
mirroring_total_memory 20.0MB
mirroring_free_memory 20.0MB
raid_total_memory 40.0MB
raid_free_memory 38.8MB
maintenance no
compression_active yes
accessible_vdisk_count 2
compression_supported yes
max_enclosures 10
encryption_supported no
flash_copy_maximum_memory 552.0MB
```

```

site_id
site_name
fctargetportmode transitional
compression_total_memory 2047.9MB

```

8. In transitional mode, host I/O is permitted on primary ports and primary host attach ports (virtual WWPN), as shown in Example 5-11 under the host_io_permitted column.

Example 5-11 WWPNs in transitional mode

```
IBM_2145:ITS0:superuser>ls targetportfc
```

id	WWPN	WNNN	port_id	owning_node_id	current_node_id	nportid	host_io_permitted
----	------	------	---------	----------------	-----------------	---------	-------------------

```
virtualized
```

1	50050768021000EF	50050768020000EF	1	1	1	010200	yes no
2	50050768029900EF	50050768020000EF	1	1	1	010201	yes yes
3	50050768022000EF	50050768020000EF	2	1	1	020200	yes no
4	5005076802A900EF	50050768020000EF	2	1	1	020201	yes yes
5	50050768023000EF	50050768020000EF	3	1	1	0A83C0	yes no
6	5005076802B900EF	50050768020000EF	3	1	1	0A83C1	yes yes
7	50050768024000EF	50050768020000EF	4	1	1	011400	yes no
8	5005076802C900EF	50050768020000EF	4	1	1	011401	yes yes
33	50050768021000F0	50050768020000F0	1	2	2	010300	yes no
34	50050768029900F0	50050768020000F0	1	2	2	010301	yes yes
35	50050768022000F0	50050768020000F0	2	2	2	020300	yes no
36	5005076802A900F0	50050768020000F0	2	2	2	020301	yes yes
37	50050768023000F0	50050768020000F0	3	2	2	011500	yes no
38	5005076802B900F0	50050768020000F0	3	2	2	011501	yes yes
39	50050768024000F0	50050768020000F0	4	2	2	0A82C0	yes no
40	5005076802C900F0	50050768020000F0	4	2	2	0A82C1	yes yes

9. Enable NPIV by changing the mode from transitional to enabled, as shown in Example 5-12.

Example 5-12 Enabling NPIV

```
IBM_2145:ITS0:superuser>chiogrp -fctargetportmode enabled 0
```

10. NPIV enablement can be verified by checking the fctargetportmode field, as shown in Example 5-13.

Example 5-13 NPIV enablement verification

```
IBM_2145:ITS0:superuser>lsiogrp 0
id 0
name io_grp0
node_count 2
vdisk_count 2
host_count 2
flash_copy_total_memory 20.0MB
flash_copy_free_memory 20.0MB
remote_copy_total_memory 20.0MB
remote_copy_free_memory 20.0MB
mirroring_total_memory 20.0MB
mirroring_free_memory 20.0MB
raid_total_memory 40.0MB
raid_free_memory 38.8MB
maintenance no
compression_active yes
accessible_vdisk_count 2
compression_supported yes
max_enclosures 10
```

```

encryption_supported no
flash_copy_maximum_memory 552.0MB
site_id
site_name
fctargetportmode enabled
compression_total_memory 2047.9MB

```

At this point, you can configure zones for hosts by using the primary host attach ports (virtual WWPNs) of the Lenovo Storage V-series ports, as shown in **bold** in the output of Example 5-8 on page 222.

5.4.3 Enabling NPIV on an existing system

When systems that are running code before 7.7.1 are upgraded to version 7.7.1 or higher, the NPIV feature is not turned on by default, because it might require changes to host side zoning.

Enabling N_Port ID Virtualization (NPIV) on an existing system requires that you complete the following steps after meeting the prerequisites:

1. Audit your SAN fabric layout and zoning rules, because NPIV has stricter requirements. Ensure that equivalent ports are on the same fabric and in the same zone. For more information, see the topic about zoning considerations for N_PortID Virtualization in Lenovo Information Center:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_icconfiguringnpiv.html

2. Check the path count between your hosts and the Lenovo storage V series system to ensure that the number of paths is half of the usual supported maximum. For more information, see the topic about zoning considerations for N_Port ID Virtualization in Lenovo Information Center:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_icconfiguringnpiv.html

3. Run the **lstorageportfc** command to note down the primary host attach WWPNs (virtual WWPNs), as shown in **bold** in Example 5-14.

Example 5-14 Using the lstorageportfc command to get primary host WWPNs (virtual WWPNs)

```

IBM_2145:ITS0:superuser>lstorageportfc
id WWPN                WWNN                port_id owning_node_id current_node_id nportid host_io_permitted
virtualized
1  50050768021000EF  50050768020000EF 1      1                1                010200 yes          no
2  50050768029900EF  50050768020000EF 1      1                1                000000 no         yes
3  50050768022000EF  50050768020000EF 2      1                1                020200 yes          no
4  5005076802A900EF  50050768020000EF 2      1                1                000000 no         yes
5  50050768023000EF  50050768020000EF 3      1                1                0A83C0 yes          no
6  5005076802B900EF  50050768020000EF 3      1                1                000000 no         yes
7  50050768024000EF  50050768020000EF 4      1                1                011400 yes          no
8  5005076802C900EF  50050768020000EF 4      1                1                000000 no         yes
33 50050768021000F0  50050768020000F0 1      2                2                010300 yes          no
34 50050768029900F0  50050768020000F0 1      2                2                000000 no         yes
35 50050768022000F0  50050768020000F0 2      2                2                020300 yes          no
36 5005076802A900F0  50050768020000F0 2      2                2                000000 no         yes
37 50050768023000F0  50050768020000F0 3      2                2                011500 yes          no
38 5005076802B900F0  50050768020000F0 3      2                2                000000 no         yes
39 50050768024000F0  50050768020000F0 4      2                2                0A82C0 yes          no

```

40 5005076802C900F0 50050768020000F0 4 2 000000 no yes

4. Include the primary host attach ports (virtual WWPNs) to your host zones.
5. Enable transitional mode for NPIV on Lenovo storage V-series system (Example 5-15).

Example 5-15 NPIV in transitional mode

```
IBM_2145:ITS0:superuser>chiogrp -fctargetportmode transitional 0
```

6. Ensure that the primary host attach WWPNs (virtual WWPNs) now allows host traffic, as shown in **bold** in Example 5-16.

Example 5-16 Host attach WWPNs (virtual WWPNs) permitting host traffic

```
IBM_2145:ITS0:superuser>lsstargetportfc
```

id	WWPN	WWNN	port_id	owning_node_id	current_node_id	nportid	host_io_permitted
virtualized							
1	50050768021000EF	50050768020000EF	1	1	1	010200	yes no
2	50050768029900EF	50050768020000EF	1	1	1	010201	yes yes
3	50050768022000EF	50050768020000EF	2	1	1	020200	yes no
4	5005076802A900EF	50050768020000EF	2	1	1	020201	yes yes
5	50050768023000EF	50050768020000EF	3	1	1	0A83C0	yes no
6	5005076802B900EF	50050768020000EF	3	1	1	0A83C1	yes yes
7	50050768024000EF	50050768020000EF	4	1	1	011400	yes no
8	5005076802C900EF	50050768020000EF	4	1	1	011401	yes yes
33	50050768021000F0	50050768020000F0	1	2	2	010300	yes no
34	50050768029900F0	50050768020000F0	1	2	2	010301	yes yes
35	50050768022000F0	50050768020000F0	2	2	2	020300	yes no
36	5005076802A900F0	50050768020000F0	2	2	2	020301	yes yes
37	50050768023000F0	50050768020000F0	3	2	2	011500	yes no
38	5005076802B900F0	50050768020000F0	3	2	2	011501	yes yes
39	50050768024000F0	50050768020000F0	4	2	2	0A82C0	yes no
40	5005076802C900F0	50050768020000F0	4	2	2	0A82C1	yes yes

7. Ensure that the hosts are using the NPIV ports for host I/O.

Remember:

The following information can be useful:

- ▶ You can verify that you are logged in to them by entering the **lsfabric -host host_id_or_name** command. If I/O activity is occurring, each host has at least one line in the command output that corresponds to a host port and shows active in the activity field:
 - Hosts where no I/O was issued in the past 5 minutes do not show active for any login.
 - Hosts that do not adhere to preferred paths might still be processing I/O to primary ports.
- ▶ Depending on the host operating system, rescanning of the SAN might be required on some hosts to recognize additional paths now provided via primary host attach ports (virtual WWPNs).

8. After a minimum of 15 minutes has passed since entering transitional mode, change the system to enabled mode by entering the command, as shown in Example 5-17 on page 226.

Example 5-17 Enabling the NPIV

```
IBM_2145:ITS0:superuser>chlogrp -fctargetportmode enabled 0
```

Now NPIV has been enabled on the Lenovo storage V-series systems, and hosts should also be using the virtualized WWPNs for I/O. At this point, the host zones can be amended appropriately to use primary host attach port WWPNs (virtual WWPNs) only.

5.5 Creating hosts by using the GUI

This section describes how to create Fibre Channel, iSCSI, and SAS hosts by using the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI. We assume that the hosts are prepared for attachment, as described in 5.3, “Preparing the host operating system” on page 191, and that you know the host WWPNs and their iSCSI initiator names.

Considerations when you configure hosts in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030

When you create a host object in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, it is important to verify the configuration limits and restrictions, which are published at the following web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/lenovo_vseries.html

1. Open the Hosts configuration window by clicking **Hosts** (Figure 5-38).

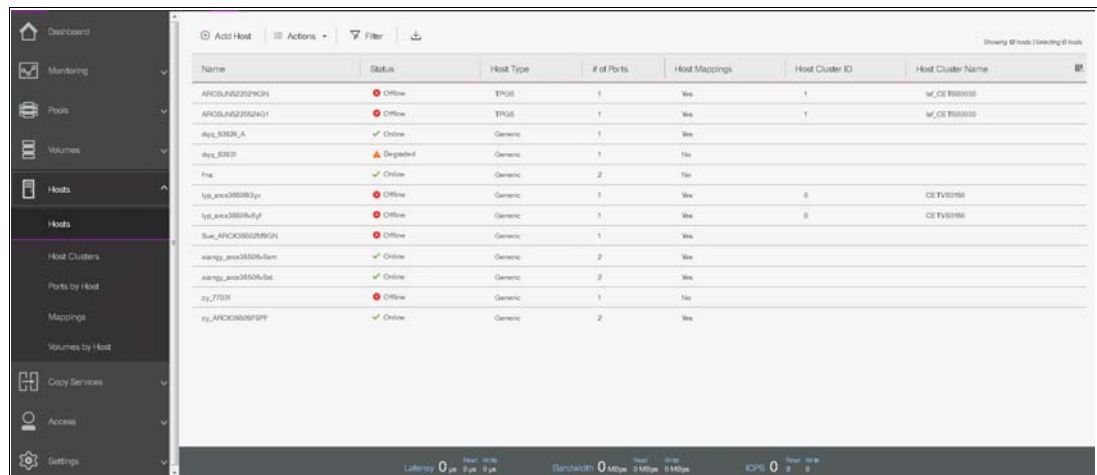


Figure 5-38 Open the Hosts window

2. To create a host, click **Add Host** to start the wizard (Figure 5-39).

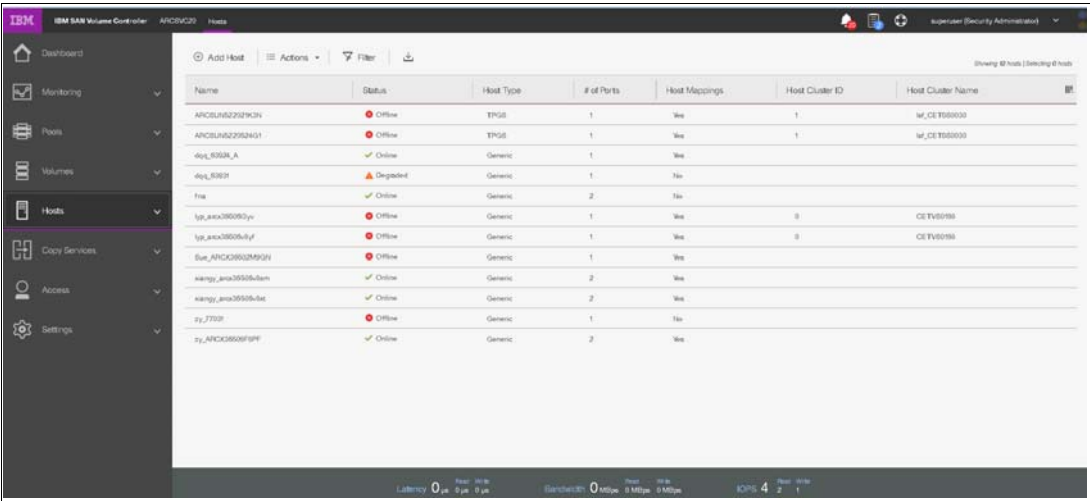


Figure 5-39 Add Host

3. If you want to create a Fibre Channel host, continue with 5.5.1, “Creating Fibre Channel hosts” on page 228. To create iSCSI hosts, go to 5.5.3, “Creating iSCSI hosts” on page 236. To create SAS hosts, go to 5.5.5, “Creating SAS hosts” on page 243.
4. After you click **Add Host**, the host selection menu opens, as shown in Figure 5-40.

Add Host

Required Fields

Name:

Host connections: ☒ Fibre Channel ☐ iSCSI

Host port (WWPN):

Optional Fields

Host type:

I/O groups:

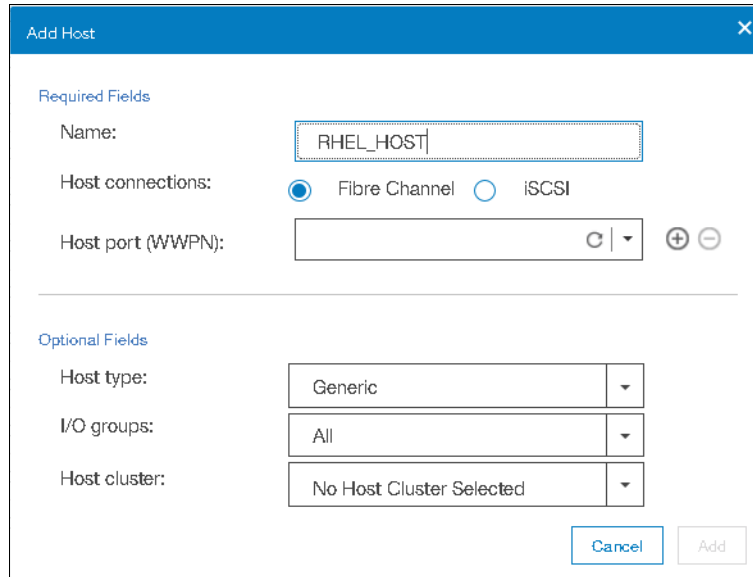
Host cluster:

Figure 5-40 Add Host window

5.5.1 Creating Fibre Channel hosts

To create Fibre Channel hosts, complete the following steps:

1. Click **Fibre Channel** (Figure 5-40 on page 227). The Fibre Channel configuration wizard opens (Figure 5-41).



The 'Add Host' dialog box is shown with a blue header bar containing a close button (X). The dialog is divided into two sections: 'Required Fields' and 'Optional Fields'.

Required Fields:

- Name:** A text input field containing 'RHEL_HOST'.
- Host connections:** Two radio buttons. 'Fibre Channel' is selected (indicated by a blue dot), and 'iSCSI' is unselected.
- Host port (WWPN):** A text input field that is currently empty, followed by a refresh icon (circular arrow) and a dropdown arrow.

Optional Fields:

- Host type:** A dropdown menu with 'Generic' selected.
- I/O groups:** A dropdown menu with 'All' selected.
- Host cluster:** A dropdown menu with 'No Host Cluster Selected' selected.

At the bottom right of the dialog are two buttons: 'Cancel' (highlighted with a blue border) and 'Add' (disabled, shown in grey).

Figure 5-41 Create a Fibre Channel host

2. Enter a host name and click the **Host port** drop-down list to get a list of all known WWPNs (Figure 5-42).

The screenshot shows the 'Add Host' dialog box with the following fields and options:

- Required Fields:**
 - Name:** RHEL_HOST
 - Host connections:** ☒ Fibre Channel ☐ iSCSI
 - Host port (WWPN):** A dropdown menu is open, displaying a list of WWPNs. The list includes:
 - 10000000C9904C33
 - 10000000C9CBA37C
 - 10000000C9CBA37D
 - 10000090FA10C2F6
 - 10000090FA10C2F7
 - 2100000E1E09D1B6
 - 2100000E1E09D1B7
 - 2100001B329A36B4
 - 2100001B329C2479
 - 21000024FF2FAFBC
 - 21000024FF2FAFBD
 - 2101001B32BA36B4
 - 2101001B32BC2479
- Optional Fields:**
 - Host type:**
 - I/O groups:**
 - Host cluster:**

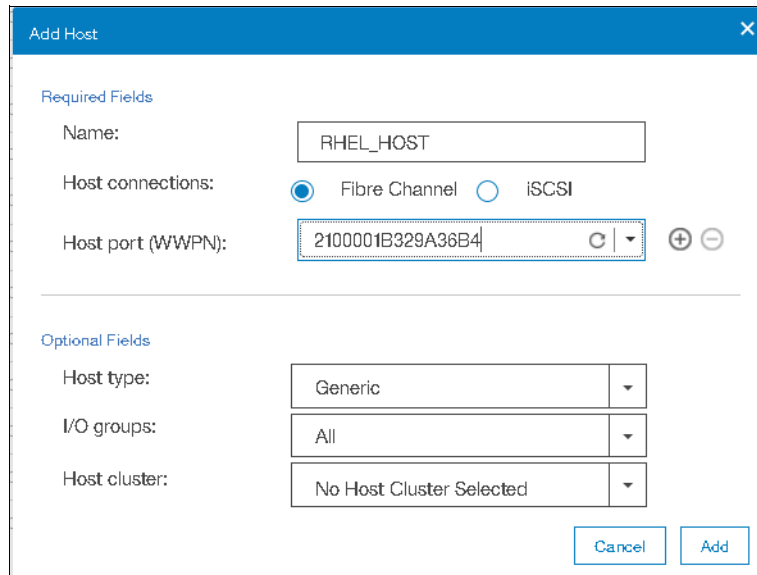
At the bottom of the dialog, there is a status bar showing 'Latency 0 μs' and 'Read Write 0 μs 0 μs'. An 'Add' button is located on the right side of the dialog.

Figure 5-42 Available WWPNs

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 have the host port WWPNs available if you prepared the hosts as described in 5.3, “Preparing the host operating system” on page 191. If they do not appear in the list, scan for new disks in your operating system and click **Rescan** in the configuration wizard. If they still do not appear, check your SAN zoning, correct it, and repeat the scanning.

Note: You can enter WWPNs manually. However, if these WWPNs are not visible to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, the host object appears as offline and it is unusable for I/O operations until the ports are visible.

3. Select the WWPN for your host (Figure 5-43).



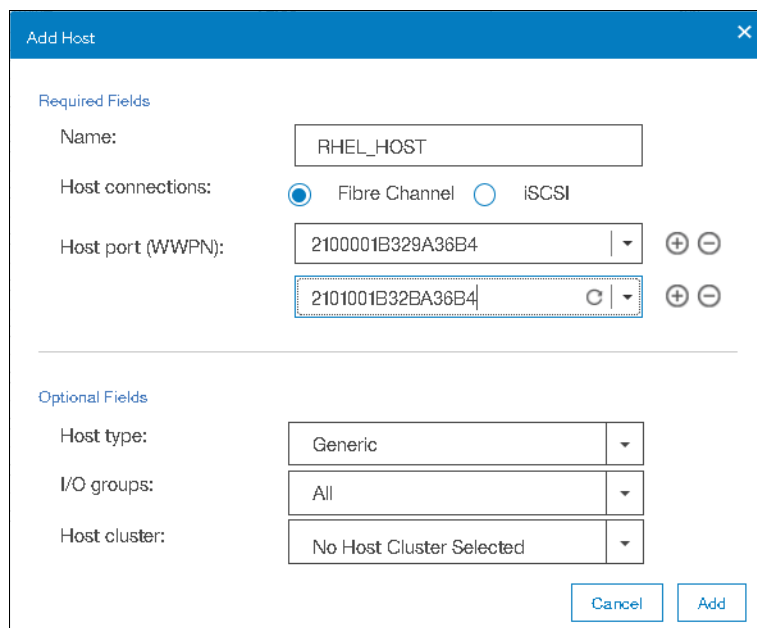
The 'Add Host' dialog box is shown with the following fields:

- Required Fields:**
 - Name: RHEL_HOST
 - Host connections: ☒ Fibre Channel ☐ iSCSI
 - Host port (WWPN): 2100001B329A36B4 (with a refresh icon and a plus/minus icon)
- Optional Fields:**
 - Host type: Generic
 - I/O groups: All
 - Host cluster: No Host Cluster Selected

Buttons: Cancel, Add

Figure 5-43 Add a port to a list

4. If you want to add additional ports to your Host, click the plus sign (+).
5. Add all ports that belong to the host (Figure 5-44).



The 'Add Host' dialog box is shown with the following fields:

- Required Fields:**
 - Name: RHEL_HOST
 - Host connections: ☒ Fibre Channel ☐ iSCSI
 - Host port (WWPN): 2100001B329A36B4 (with a plus/minus icon)
 - Host port (WWPN): 2101001B32BA36B4 (with a refresh icon and a plus/minus icon)
- Optional Fields:**
 - Host type: Generic
 - I/O groups: All
 - Host cluster: No Host Cluster Selected

Buttons: Cancel, Add

Figure 5-44 Add all WWPNs

Creating offline hosts: If you want to create hosts that are offline or not connected currently, you can enter the WWPNs manually. Type them into the Host port (WWPN) field and add them to the list. See Figure 5-45.

The screenshot shows a window titled "Add Host" with a close button (X) in the top right corner. Inside the window, under "Host connections:", there are two radio buttons: "Fibre Channel" (which is selected) and "iSCSI". Below this, there is a "Name:" label followed by a text box containing "ESXi Server". Underneath, there is a "Host port (WWPN):" label followed by two text boxes. The first text box contains "10008C7CFF096F81" and has expand/collapse buttons (+/-) to its right. The second text box contains "200000000000000000" and has a refresh button (circular arrow) and expand/collapse buttons (+/-) to its right. Below these fields is a link that says "Advanced" with a right-pointing triangle icon. At the bottom of the window are two buttons: "Add" and "Cancel".

Figure 5-45 Manually added WWPN

6. If you are creating a Hewlett-Packard UNIX (HP-UX) or Target Port Group Support (TPGS) host, select **Advanced**, and more options appear (Figure 5-46). Select your host type.

The screenshot shows the 'Add Host' dialog box with the 'Advanced Settings' tab selected. Under 'Required Fields', the 'Name' field contains 'RHEL_HOST'. The 'Host connections' section has 'Fibre Channel' selected with a radio button. The 'Host port (WWPN)' field contains '2100001B329A36B4'. Below this, there is another field containing '2101001B32BA36B4'. In the 'Optional Fields' section, the 'Host type' dropdown is set to 'Generic'. The 'I/O groups' dropdown is also set to 'Generic'. The 'Host cluster' dropdown is set to 'Generic'. An 'Add' button is located at the bottom right of the dialog.

Figure 5-46 Add Host: Advanced Settings

7. You can set the **I/O Groups** that your host can access. The host objects must belong to the same I/O groups as the volumes that you want to map. Otherwise, these volumes are not visible to the host. See Figure 5-47.

This screenshot shows the same 'Add Host' dialog box, but with different selections in the 'Optional Fields' section. The 'Host type' remains 'Generic'. The 'I/O groups' dropdown is now set to 'All'. The 'Host cluster' dropdown is also set to 'All'. The 'Add' button remains at the bottom right.

Figure 5-47 Setting I/O groups

Note: The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support a maximum of two nodes for each system. The two nodes are arranged as two I/O groups per cluster. Due to the host object limit per I/O group, for maximum host connectivity, it is best to create hosts that use single I/O groups.

8. Select the desired host cluster if any as shown in Figure 5-48.

The 'Add Host' dialog box is shown with the following fields:

- Required Fields:**
 - Name: RHEL_HOST
 - Host connections: ☒ Fibre Channel ☐ iSCSI
 - Host port (WWPN): 2100001B329A36B4 (with expand/collapse icons)
 - 2101001B32BA36B4 (with expand/collapse icons)
- Optional Fields:**
 - Host type: Generic
 - I/O groups: All
 - Host cluster: No Host Cluster Selected (dropdown menu is open showing:
 - No Host Cluster Selected
 - CETVS0168 2 hosts
 - lsf_CETSS0030 2 hosts

An 'Add' button is visible next to the host cluster dropdown.

Figure 5-48 Selecting Host cluster

9. Click **Add Host**, and the wizard creates the host (Figure 5-49).

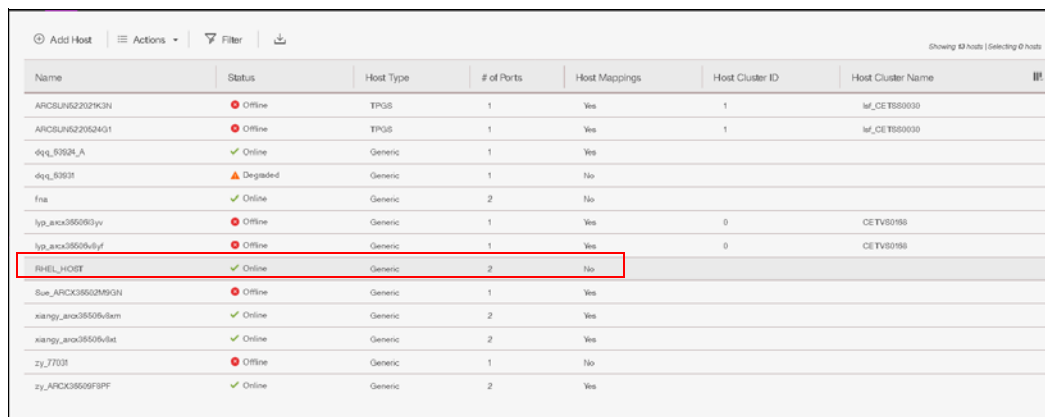
The 'Add Host' dialog box shows the following task completion details:

- Task completed. 100%
- View more details
- Task started. 12:10 PM
- Creating the host RHEL_HOST 12:10 PM
- Running command: 12:10 PM
- `svctask mkhost -fcwwpn 2100001B329A36B4:2101001B32BA36B4 -force -iogrp 0:1 -name RHEL_HOST -type generic` 12:10 PM
- The host (ID 9) was successfully created. 12:10 PM
- Synchronizing memory cache. 12:10 PM
- The task is 100% complete. 12:10 PM
- Task completed. 12:10 PM

Buttons: Cancel, Close

Figure 5-49 Add Host task completes

- Click **Close** to return to the Hosts window. The host that was created shows up in the hosts window as shown in Figure 5-50.



Name	Status	Host Type	# of Ports	Host Mappings	Host Cluster ID	Host Cluster Name	
APCBUN622029K3N	Offline	TPGS	1	Yes	1	Inf_CE1B00000	
APCBUN622025N401	Offline	TPGS	1	Yes	1	Inf_CE1B00000	
qlq_5762N_A	Online	Generic	1	Yes			
qlq_6363I	Degraded	Generic	1	No			
fna	Online	Generic	2	No			
ljp_axc3550503yv	Offline	Generic	1	Yes	0	CETV50198	
ljp_axc3550503yf	Offline	Generic	1	Yes	0	CETV50198	
RHEL_HOST	Online	Generic	2	No			
8ue_APCK35502M503N	Offline	Generic	1	Yes			
xiangy_axc3550503sam	Online	Generic	2	Yes			
xiangy_axc3550503bat	Online	Generic	2	Yes			
zy_7703I	Offline	Generic	1	No			
zy_APCK3550503PF	Online	Generic	2	Yes			

Figure 5-50 Defined hosts

- Repeat steps 1 to 10 for all of your Fibre Channel hosts. After you add all of the Fibre Channel hosts, create volumes and map them to the created hosts. See Chapter 6, “Volume configuration” on page 269.

5.5.2 Configuring the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 for FC connectivity

You can configure the FC ports on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 for use for certain connections only. This capability is referred to as *port masking*. In a system with multiple I/O groups and remote partnerships, port masking is a useful tool for ensuring peak performance and availability.

The following options are available per port:

- **Any:** Allow local and remote communication between nodes.
- **Local:** Allow only local node communication.
- **Remote:** Allow only remote node communication.
- **None:** Do not allow any node communication.

In all cases, host I/O is still permitted, so the None option can be used to exclusively reserve a port for host I/O.

A limit of **16** logins exists per node from another node before an error is logged. A combination of port masking and SAN zoning can help you manage logins and provide optimum I/O performance with local, remote, and host traffic.

To configure FC ports, complete the following steps:

- Go to **Settings** → **Network**, as shown in Figure 5-51 on page 235.




Figure 5-51 Opening the network settings view

2. Select **Fibre Channel Ports** and the Fibre Channel Ports configuration view displays, as shown in Figure 5-52.




Figure 5-52 The Fibre Channel Ports view

3. To configure a port, right-click the port and select **Modify Connection**. The window that is shown in Figure 5-53 opens.

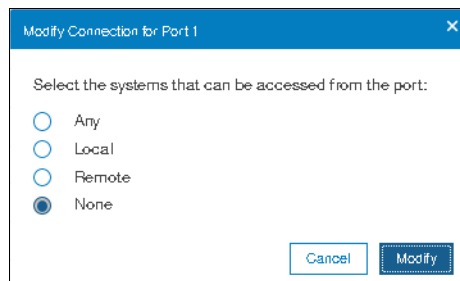


Figure 5-53 Modifying the connection for Port 1

In this example, we select **None** to restrict traffic on this port to host I/O only. Click **Modify** to confirm the selection.

Note: This action configures Port 1 for **all** nodes. You cannot configure FC ports on a per node basis.

- You can view connections between nodes, storage systems, and hosts by selecting **Fibre Channel Connectivity** while you are in the Network settings view. Choose the connections that you want to view and click **Show Results**, as shown in Figure 5-54.

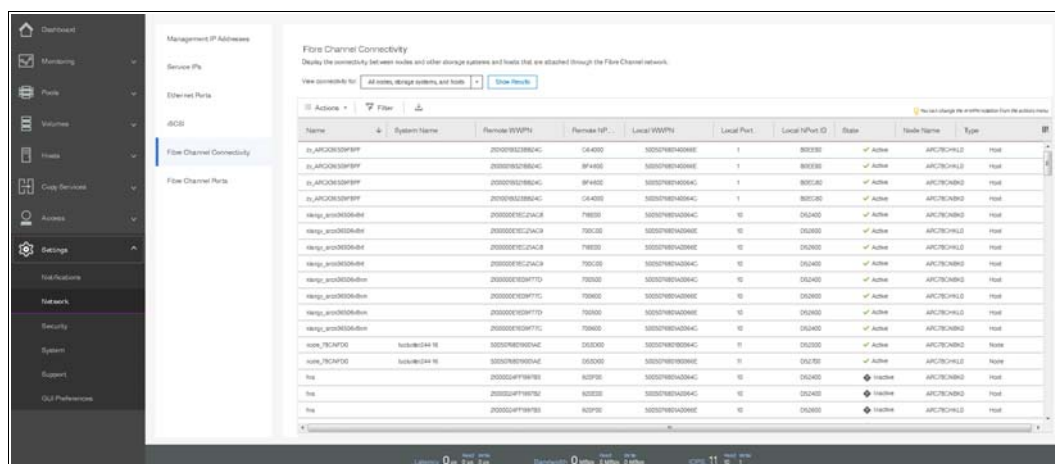


Figure 5-54 Viewing FC connections between nodes, storage systems, and hosts

5.5.3 Creating iSCSI hosts

To create iSCSI hosts, complete the following steps:

- Click **iSCSI** and the iSCSI configuration wizard opens (Figure 5-55).

Figure 5-55 Add an iSCSI host

- Enter a descriptive host name, type the iSCSI initiator name in the name field. Type the iSCSI port information in Figure 5-56 on page 237 the iSCSI port field.. If you want to add several initiator names to one host, repeat this step by clicking the plus sign (+).

3. If you are connecting an HP-UX or TPGS host, select **Advanced** (Figure 5-55 on page 236) and select the correct host type (Figure 5-56). Click **Add**.

The screenshot shows the 'Add Host' dialog box with the following fields and options:

- Required Fields:**
 - Name: ESXi_iSCSI_HOST
 - Host connections: ☐ Fibre Channel ☒ iSCSI
 - iSCSI host IQN: iqn.1998-01.com.vmware.iscsi:itsovn
- Optional Fields:**
 - CHAP authentication: ☐
 - CHAP secret: Enter 1 to 79 characters
 - Host type: Generic (selected in dropdown)
 - I/O groups: Generic (selected in dropdown)
 - Host cluster: HP/UX, OpenVMS, TPGS, VVOL (available in dropdown)
- Add** button

Figure 5-56 Create an iSCSI host: Advanced settings

4. You can set the I/O groups that your host can access.

Important: The host objects must belong to the same I/O groups as the volumes that you want to map. Otherwise, the host cannot see these volumes. For more information, see Chapter 6, “Volume configuration” on page 269.

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support a maximum of four nodes per system. These nodes are arranged as two I/O groups per cluster. Due to the host object limit per I/O group, for maximum host connectivity, it is best to create hosts that use single I/O groups.

5. The wizard completes (Figure 5-57 on page 238). Click **Close**.

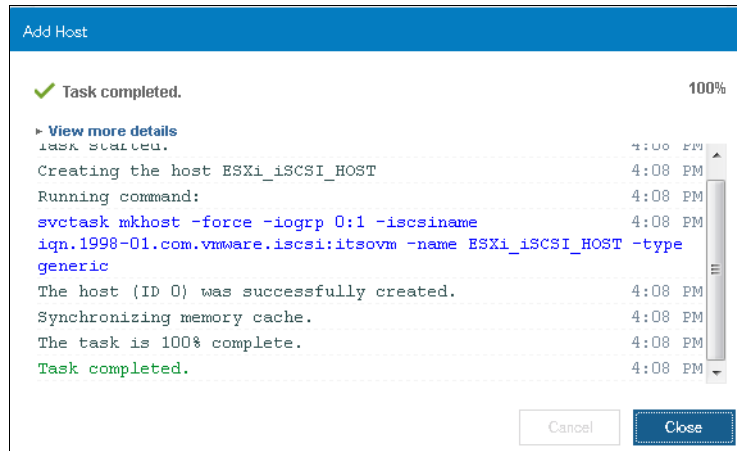


Figure 5-57 Add an iSCSI host: Complete

- Repeat these steps for every iSCSI host that you want to create. Figure 5-58 shows all of the hosts after you create two Fibre Channel hosts and one iSCSI host.

Name	Status	Host Type	# of Ports	Host Mappings	Host Cluster ID	Host Cluster Name
APC-SUN5252N3N	Offline	TPGS	1	Yes	1	M_CET550030
APC-SUN5252N401	Offline	TPGS	1	Yes	1	M_CET550030
esx_639d_A	Online	Generic	1	Yes		
esx_639d	Degraded	Generic	1	No		
ESXi_iSCSI_HOST	Offline	Generic	1	No		
tsa	Online	Generic	2	No		
isp_arco3650d-iv	Offline	Generic	1	Yes	0	CETV5268
isp_arco3650d-yl	Offline	Generic	1	Yes	0	CETV5268
RHEL_HOST	Online	Generic	2	No		
Sas_APC36502M30N	Offline	Generic	1	Yes		
xiangp_arco3650d-bn	Online	Generic	2	Yes		
xiangp_arco3650d-bd	Online	Generic	2	Yes		
zy_770B	Offline	Generic	1	No		
zy_ARC3650F8PF	Online	Generic	2	Yes		

Figure 5-58 All hosts

Note: iSCSI hosts might show a *Degraded* status until the volumes are mapped. This limitation relates to the implementation of iSCSI in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. This status is not necessarily a problem with network connectivity or the host configuration.

The iSCSI host is now configured on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. To provide connectivity, the iSCSI Ethernet ports must also be configured.

5.5.4 Configuring the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 for iSCSI host connectivity

Complete the following steps to enable iSCSI connectivity:

- Switch to the configuration Overview window and select **Network** (Figure 5-59 on page 239).




Figure 5-59 Configuration: Network

2. Select **iSCSI** and the iSCSI Configuration window opens (Figure 5-60).

 A screenshot of the 'iSCSI Configuration' window. On the left is a sidebar with a list of configuration categories: 'Management IP Addresses', 'Service IPs', 'Ethernet Ports', 'iSCSI' (highlighted with a blue bar), 'Fibre Channel Connectivity', and 'Fibre Channel Ports'. The main area is titled 'iSCSI Configuration' with the subtitle 'Configure system properties to connect to iSCSI-attached hosts.' It contains several sections:

- Name**: A 'System Name' field with the value 'mcr-tb5-cluster-29'.
- iSCSI Aliases (optional)**: A table with three columns: 'Node Canister Name', 'iSCSI Alias', and 'iSCSI Name (QN)'. It contains two rows:

Node Canister Name	iSCSI Alias	iSCSI Name (QN)
node1	sv_n1	iqn.1986-03.com.ibm:2145.mcr-tb5-cluster-29.node1
node2	sv_n2	iqn.1986-03.com.ibm:2145.mcr-tb5-cluster-29.node2
- A warning box with a yellow triangle icon stating: 'Changes have not yet been applied to the system for the node alias or name.' with an 'Apply Changes' button.
- iSNS (optional)**: An 'iSNS Address' field.
- A link for 'Modify CHAP Configuration'.

Figure 5-60 iSCSI Configuration window

- The system waits until you apply the changes that you made. Click **Apply Changes**. All changes are applied, as shown in Figure 5-61.

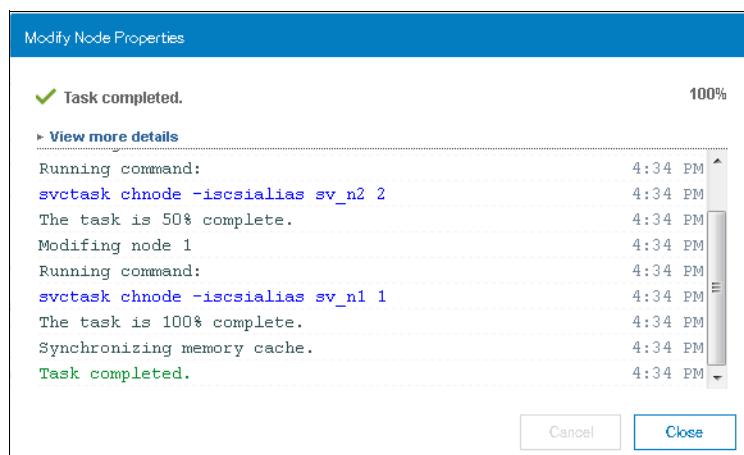


Figure 5-61 Applying all iSCSI changes

- The Configuration window (Figure 5-60 on page 239) shows an overview of all of the iSCSI settings for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. You can configure the iSCSI alias, Internet Storage Name Service (iSNS) addresses, Challenge Handshake Authentication Protocol (CHAP), and the iSCSI IP address, which we also edit in the basic setup.
- Click **Ethernet Ports** to enter the iSCSI IP address (Figure 5-62). Repeat this step for each port that you want to use for iSCSI traffic.

Note: We advise that you reserve at least one port for the management IP address. Typically, reserve Port 1 for the management IP address and configure Port 2 for iSCSI. In our example, we use Port 1 due to limited cabling.

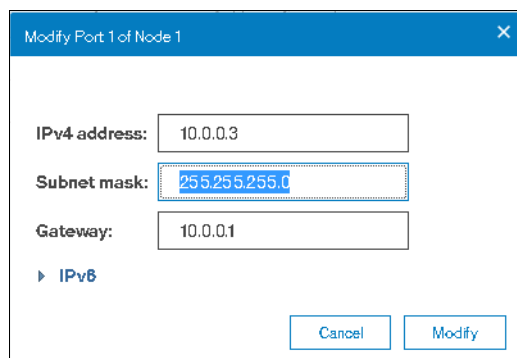


Figure 5-62 Enter an iSCSI IP address

- After you enter the IP address for each port, click **Modify** to enable the configuration.

7. After the changes are successfully applied, click **Close** (Figure 5-63).

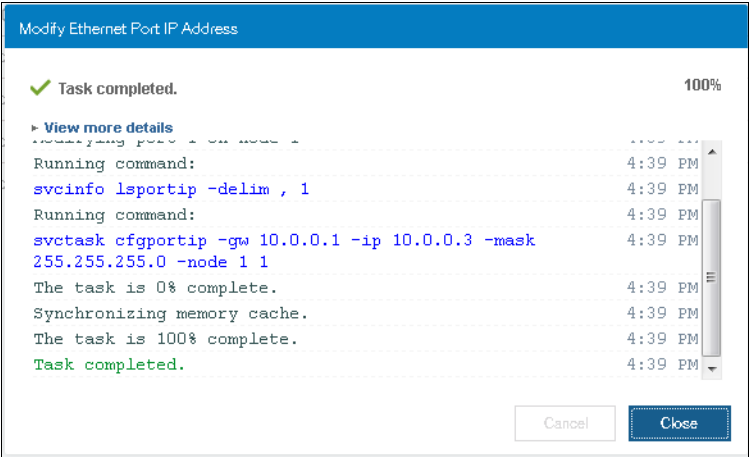


Figure 5-63 iSCSI IP address that was successfully modified

8. Under **Actions**, you can check the hosts that are enabled for iSCSI. See Figure 5-64.

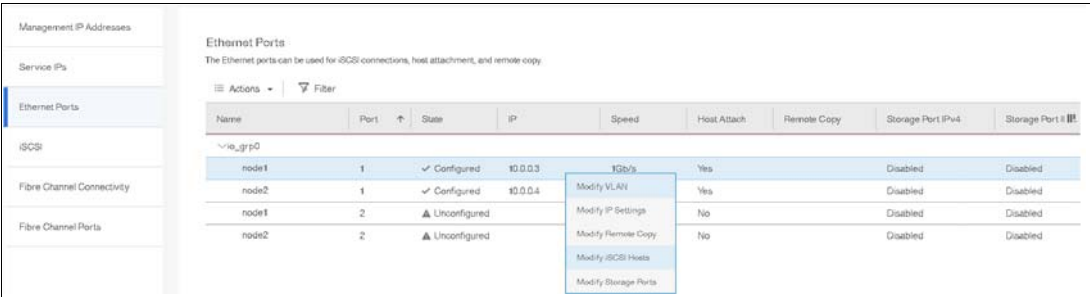


Figure 5-64 Actions menu to modify iSCSI hosts

9. By default, all iSCSI hosts are enabled (Figure 5-65).

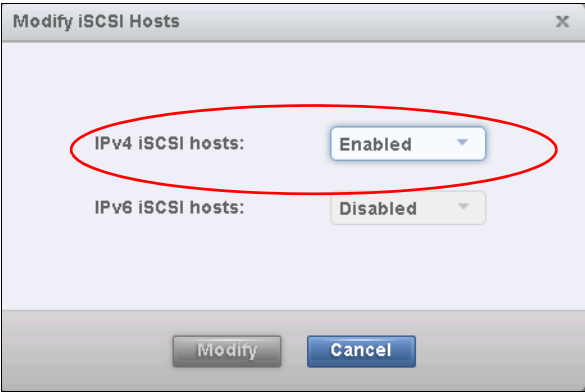


Figure 5-65 Enabled iSCSI hosts

10. To modify virtual LAN (VLAN) settings, right-click the port again and click **Modify VLAN**. The window that is shown in Figure 5-66 opens. Check **Enable** to turn on VLAN tagging and set the tag in the field that is provided. The failover port needs to belong to the same VLAN so leave **Apply change to the failover port too** checked. Click **Modify** to confirm.

Figure 5-66 Modifying VLAN settings for Port 2 of Node 1

11. Repeat the previous steps for all ports that need to be configured.
12. You can also configure iSCSI aliases, an iSNS name, and CHAP authentication. These options are located in the iSCSI Configuration view. To access this view, click **iSCSI** in the Network settings view, as shown in Figure 5-67.

Node Controller Name	iSCSI Alias	iSCSI Name (QNAP)
node1	ev_x1	iqn.1996-03.com.lba:2145.mc-b5-cluster-29,node1
node2	ev_x2	iqn.1996-03.com.lba:2145.mc-b5-cluster-29,node2

Figure 5-67 Advanced iSCSI configuration

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are now configured and ready for iSCSI use. Document the initiator names of your storage canisters (Figure 5-60 on page 239) because you need them later. To create volumes and map them to a host, go to Chapter 6, “Volume configuration” on page 269.

5.5.5 Creating SAS hosts

To create a SAS host, complete the following steps:

1. From the main screen, follow the path **Hosts** → **Hosts** → **Add Host**, and the host configuration wizard opens, as shown in Figure 5-68.

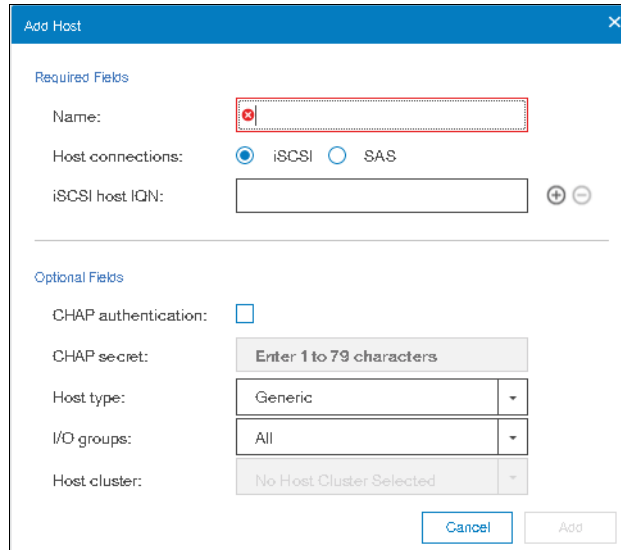


Figure 5-68 The host configuration wizard

2. Enter a descriptive host name and click **SAS** under **Host Connections** option as shown in Figure 5-69.

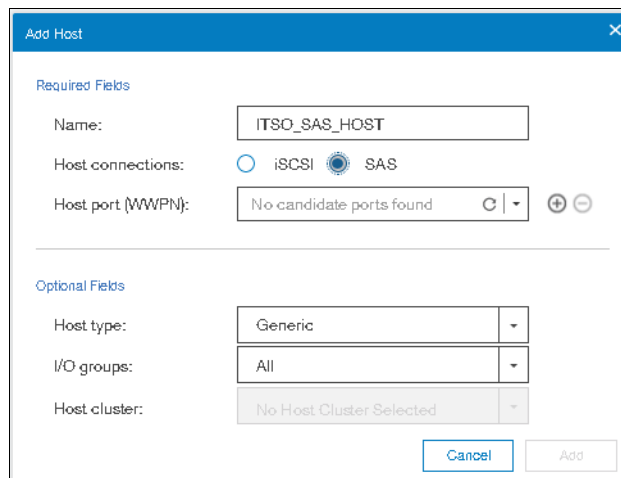


Figure 5-69 SAS WWPNS that are visible to the system

3. Click the **Host port (WWPN)** drop down and select the desired WWPN's belonging to the host from the drop-down as shown in Figure 5-70 on page 244.

Add Host

Required Fields

Name: ITSO_SAS_HOST

Host connections: ☐ iSCSI ☒ SAS

Host port (WWPN): 1234567887654321

Host port (WWPN): 8765432112345678

Optional Fields

Host type: Generic

I/O groups: All

Host cluster: No Host Cluster Selected

Cancel Add

Figure 5-70 SAS WWPNS

- If you previously prepared a SAS host, as described in 5.3, “Preparing the host operating system” on page 191, the WWPNS that you recorded in this section appear. If they do not appear in the list, verify that you completed all of the required steps and check your cabling. Then, click **Rescan** in the configuration wizard. Ensure that the ends of the SAS cables are aligned correctly.

Note: You can enter WWPNS manually. However, if these WWPNS are not visible to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, the host object appears as offline and it is unusable for I/O operations until the ports are visible.

- Under the **Optional Fields** section, you can set host type, the I/O groups that your host can access, and the host cluster the host belongs to if it is defined.

Important: Host objects must belong to the same I/O groups as the volumes that you want to map. Otherwise, the volumes are not visible to the host.

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support a maximum of four nodes per system. These nodes are arranged as two I/O groups per cluster. Due to the host object limit per I/O group, for maximum host connectivity, it is best to create hosts that use single I/O groups.

You can choose the host type. If you use HP/UX, OpenVMS, or TPGS, you must configure the host. Otherwise, the default option (Generic) is acceptable.

- Click **Add Host** and the wizard completes, as shown in Figure 5-71.

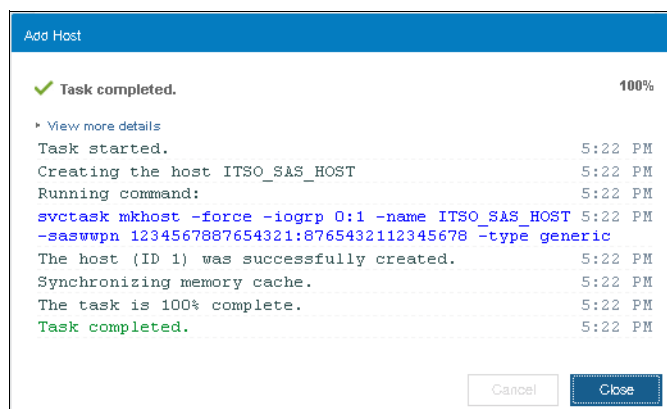


Figure 5-71 Completion of the Add Host wizard

- Click **Close** to return to the host view, which now lists your newly created host object, as shown in Figure 5-72.

Name	Status	Host Type	# of Ports	Host Mappings	Host Cluster ID	Host Cluster Name
curew0013_host	Online	Generic	2	Yes		
ITSO_SAS_HOST	Offline	Generic	2	No		

Figure 5-72 The hosts view lists the newly created host object

- Repeat these steps for all of your SAS hosts.

After all of the host objects are created, see Chapter 6, “Volume configuration” on page 269 to create volumes and map them to the hosts.

5.6 Host Clusters

Host clusters is supported starting with version 7.7.1. A host cluster is a group of logical host objects that can be managed together. For example, you can create a volume mapping that is shared by every host in the host cluster. The systems use internal protocols to manage access to the volumes and ensure consistency of the data. Host objects that represent hosts can be grouped in a host cluster and share access to volumes.

Volume mappings can either be **shared** or **private**. **Shared mappings** are volume mappings that are shared among all the hosts that are in a host cluster. When a host cluster is created, any common volume mappings become shared among all the hosts within the host cluster. If a mapping is not common, it remains a **private mapping** for that host only. Private mappings are mappings that are associated with an individual host.

A host cluster allows a user to create a group of hosts to form a cluster, which is treated as one single entity, thus allowing multiple hosts to have access to the same set of volumes.

Volumes that are mapped to that host cluster, are assigned to all members of the host cluster with the same SCSI ID.

By defining a host cluster, the user can map one or more volumes to the host cluster object. As a result the volume, or set of volumes, in turn gets assigned to each individual host object

that is part of the host cluster and each of the volumes gets mapped with the same SCSI ID to all the hosts that are part of the host cluster with just one command.

A host cluster is made up of individual hosts, and volumes can also be assigned to individual hosts that make up the cluster. Even though a host is part of host cluster, volumes can still be assigned to a particular host in a non-shared manner. A policy can be devised which might pre-assign a standard set of SCSI IDs for volumes to be assigned to the host cluster, and another set of SCSI IDs to be used for individual assignments to hosts.

Note: For example, SCSI IDs 0 - 100 can be used for individual host assignment, and SCSI IDs higher than 100 can be used for host cluster. By employing such a policy, wanted volumes will not be shared, and others can be. For example, the boot volume of each host can be kept private, while data and application volumes can be shared.

A typical use case is to define a host cluster containing all of the WWPNs belonging to the hosts that are participating in a host operating system-based cluster, such as IBM PowerHA, Microsoft Cluster Server (MSCS) and such.

This section describes the following host cluster operations using the GUI:

- ▶ Creating a host cluster
- ▶ Adding a member to the host cluster
- ▶ Listing a host cluster members
- ▶ Assigning a volume to the host cluster
- ▶ Unmapping a volume from the host cluster
- ▶ Removing a host cluster member
- ▶ Removing the host cluster

Note: From V8.1.0 onwards, various **Host Cluster** related operations can also be done using the GUI in addition to CLI.

Note: Host clusters enable you to create individual hosts and add them to a host cluster. Care must be taken to make sure that no loss of access occurs when transitioning to host clusters.

5.6.1 Creating a host cluster

To create a host cluster using the GUI, follow the steps listed.

1. Click **Hosts** → **Hosts** from the main panel as shown in Figure 5-73 on page 247.

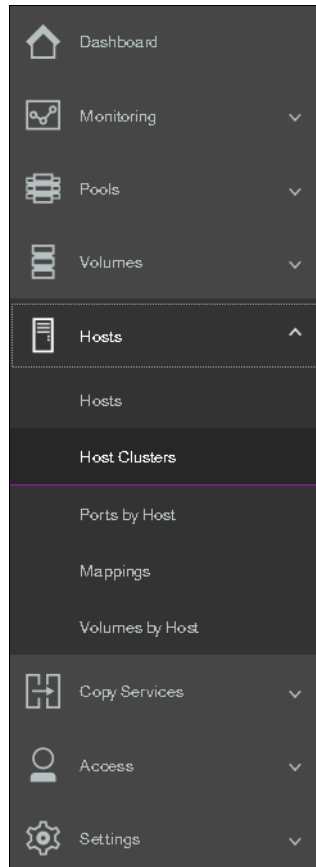


Figure 5-73 Host Clusters

2. Click **Create Host Cluster** as shown in Figure 5-74.

Create Host Cluster							Configure Remote Support	
							Configure	More Info
ID	Name	Status	Host Count	Mappings Count	Ports Count			
0	CETV02088	Offline	2	75	2			
1	sf_CET000030	Offline	2	14	2			

Figure 5-74 Create Host Cluster

3. Provide the name of the Host Cluster that you want to create as shown in Figure 5-75 on page 248.

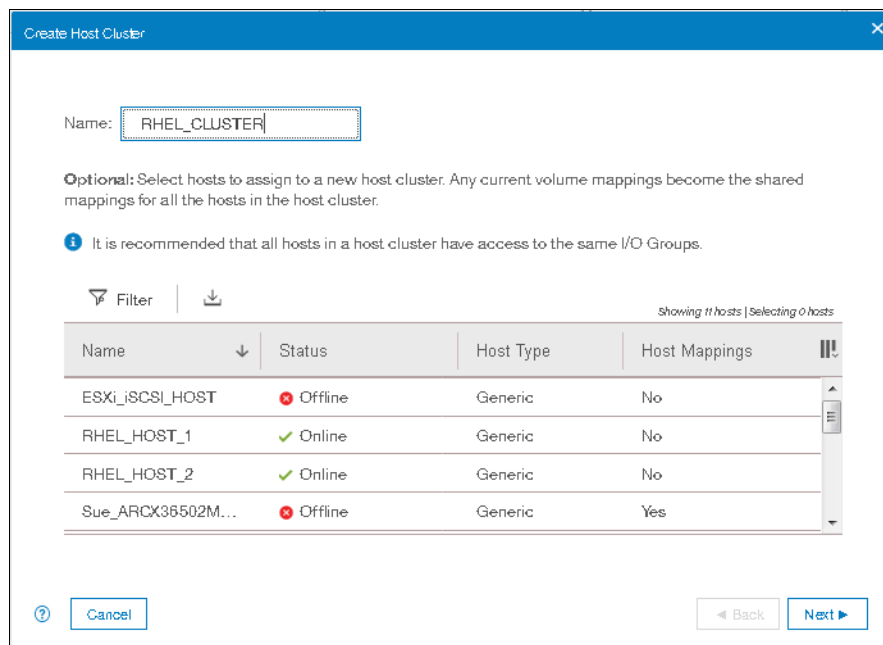


Figure 5-75 Defining Host Cluster

- At this point, you can either select the list of hosts that are going to be the part of the host cluster or you can defer that action. In the given example, we have chosen to defer the action of adding hosts to Host Cluster definition as that part is covered in 5.6.2, "Adding a member to a host cluster" on page 249.
- Click **Next** and you will see host cluster creation completion screen as shown in Figure 5-76.

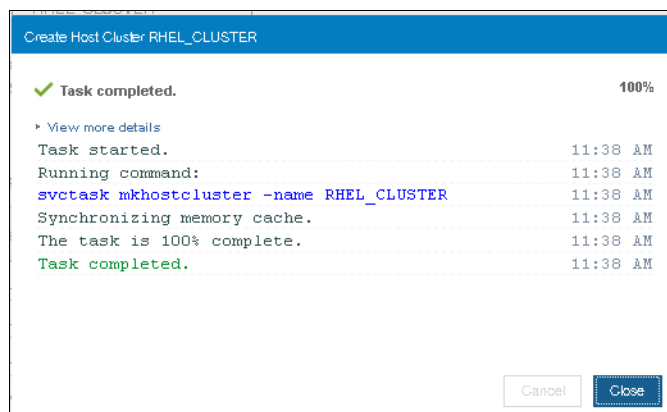


Figure 5-76 Host Cluster creation completed

- Click **Close** and you will be able to see the host cluster you created in the list as shown Figure 5-77.

ID	Name	Status	Host Count	Mappings Count	Ports Count
0	CETV80958	Offline	2	78	2
1	Wf_CET880030	Offline	2	54	2
2	RHEL_CLUSTER	Offline	0	0	0

Figure 5-77 Host Clusters list

Note: A host cluster could be offline either due to no members present OR all the members are offline. In our example it is offline as we have not added any members.

5.6.2 Adding a member to a host cluster

To add a member to an existing host cluster, follow the steps listed.

1. From the **Host Clusters** pane, right click on the desired host cluster to which you wish to add members as shown in Figure 5-78.

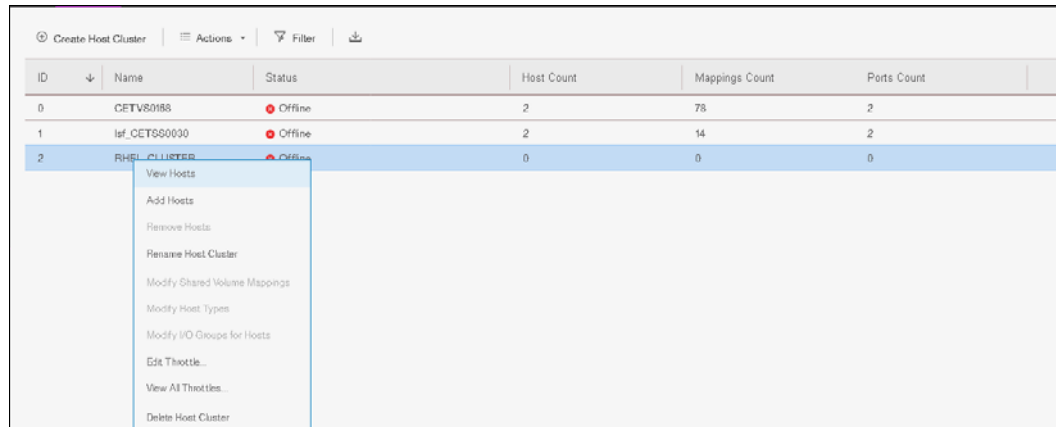


Figure 5-78 Selecting the Host Cluster

2. Click **Add Host** and you will get a selection window as shown in Figure 5-79.

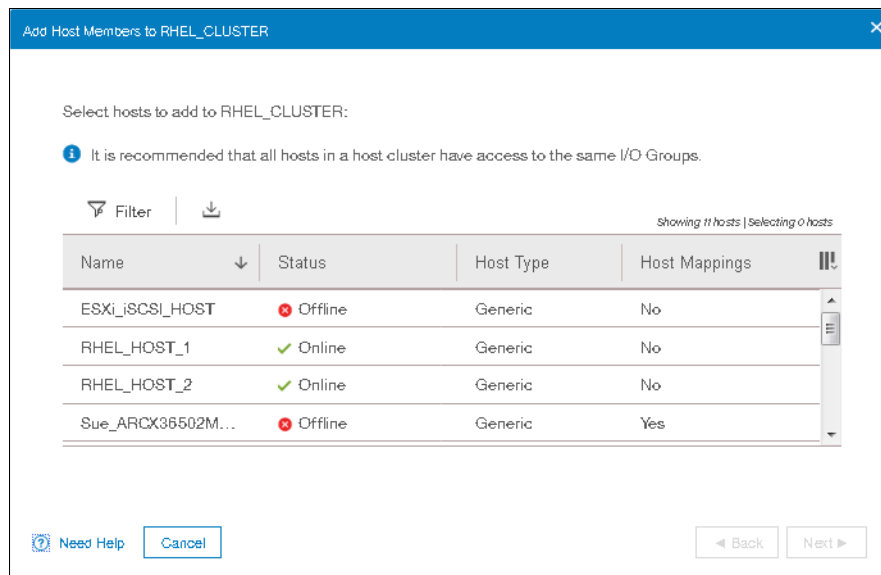


Figure 5-79 Host selection window

3. Select the desired hosts as shown in Figure 5-80 on page 250.




Figure 5-80 Selecting desired hosts

- Click **Next** and a summary of the hosts to be added will be shown as depicted in Figure 5-81.

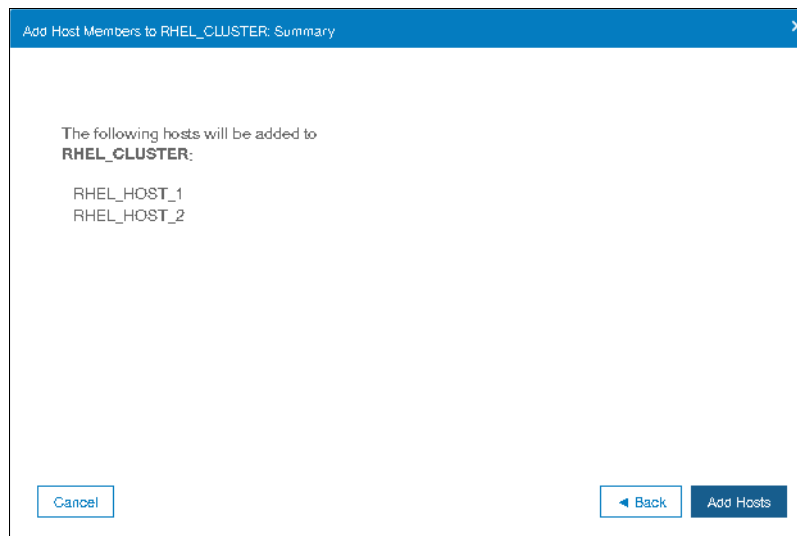


Figure 5-81 Summary of hosts

- Click **Add Hosts** and hosts will be added to the host cluster definition as shown in Figure 5-82 on page 251.

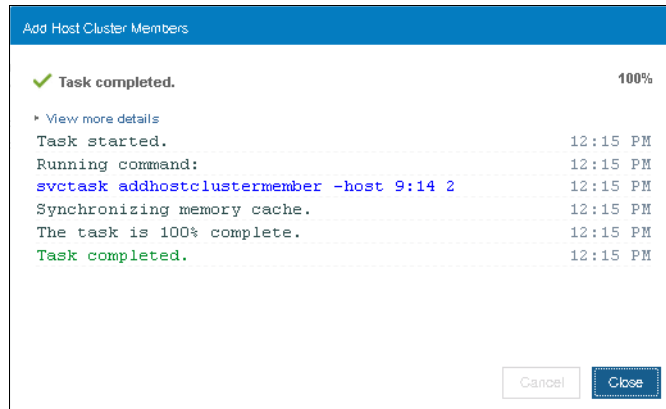


Figure 5-82 Hosts added to host cluster definition

- Click **Close** and you will see host cluster definition will show **online** status as hosts are added to a host cluster as shown in Figure 5-83.

ID	Name	Status	Host Count	Mappings Count	Ports Count
0	CETVS0168	Offline	2	78	2
1	Isf_CETSS0030	Offline	2	14	2
2	RHEL_CLUSTER	Online	2	0	4

Figure 5-83 Host cluster definition online

5.6.3 Listing a host cluster member

To list members of an existing host cluster, follow the steps listed.

- From the **Host Clusters** pane, right click on the desired host cluster for which you wish to view members as shown in Figure 5-84.

ID	Name	Status	Host Count	Mappings Count	Ports Count
0	CETVS0168	Offline	2	78	2
1	Isf_CETSS0030	Offline	2	14	2
2	RHEL_CLUSTER	Online	2	0	4

Figure 5-84 Viewing host cluster members

- Click **View Hosts** will list the host cluster members as shown in Figure 5-85 on page 252.

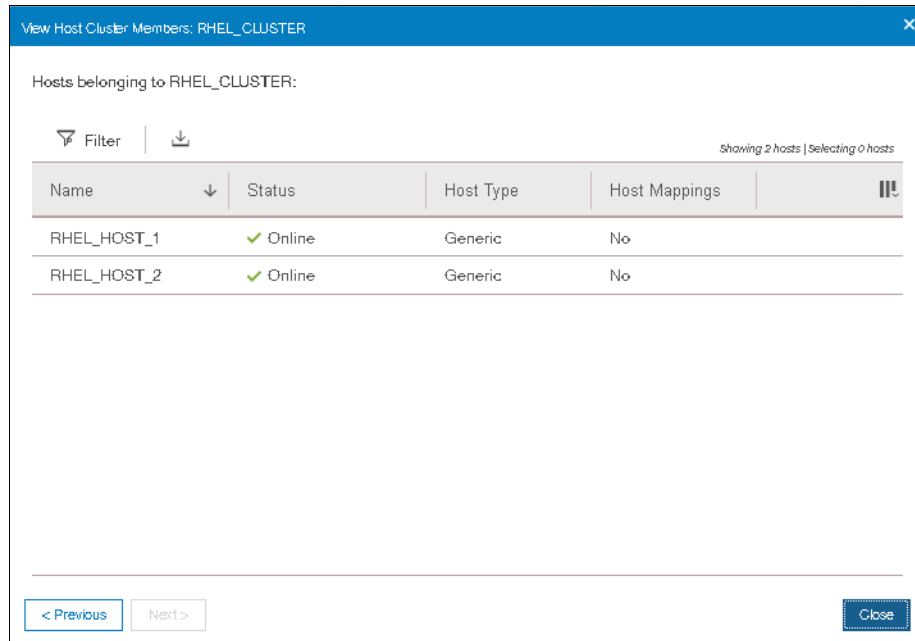


Figure 5-85 Listing host cluster members

- Click **Close** and you will be back on the **Host Clusters** pane as shown in Figure 5-86.

Create Host Cluster

Actions

Filter

Showing 3 host clusters | Selecting 1 host cluster

ID	Name	Status	Host Count	Mappings Count	Ports Count	
0	CETV50166	Offline	2	76	2	
1	lrf_CET660030	Offline	2	14	2	
2	RHEL_CLUSTER	Online	2	0	4	

Figure 5-86 Host Clusters pane

5.6.4 Assigning a volume to a Host Cluster

To assign a volume to a host cluster follow the steps listed.

- From the **Host Clusters** pane, select the desired host cluster as shown in Figure 5-87.

Create Host Cluster

Actions

Filter

Download

Showing 3 host clusters | Selecting 1 host cluster

ID	Name	Status	Host Count	Mappings Count	Ports Count	
0	CETV50166	Offline	2	76	2	
1	lrf_CET660030	Offline	2	14	2	
2	RHEL_CLUSTER	Online	2	0	4	

Figure 5-87 Host Clusters

- Right click on the desired host cluster and select **Modify Shared Volume Mappings** as shown in Figure 5-88 on page 253.

ID	Name	Status	Host Count	Mappings Count	Ports Count
0	CETVS0168	Offline	2	78	2
1	lsf_CETSS0030	Offline	2	14	2
2	RHEL_CLUSTER	Online	2	0	4

Figure 5-88 Modify Shared Volume Mapping for Host Cluster

3. A window showing list of volumes currently mapped to the selected host cluster shows up as shown in Figure 5-89.

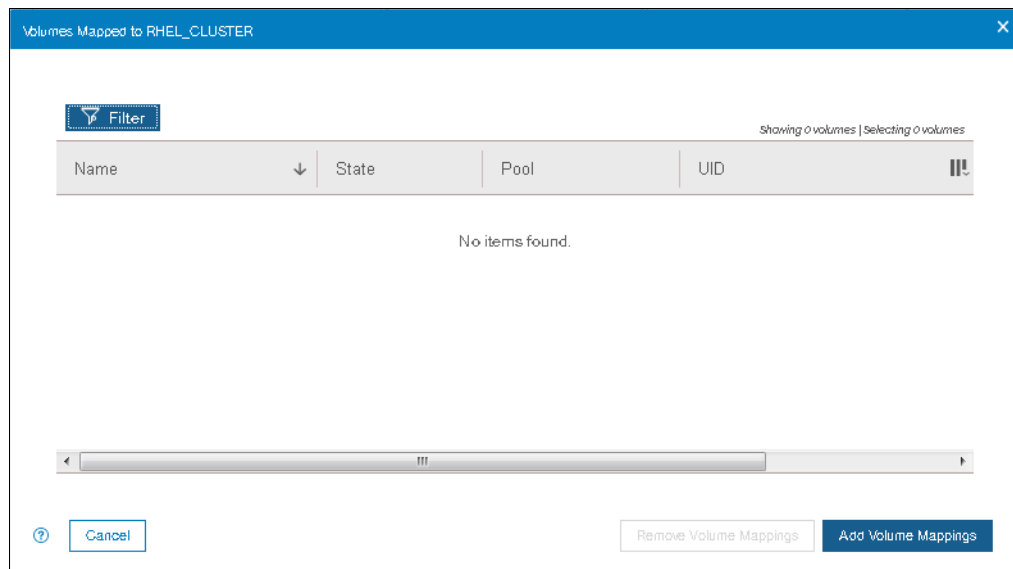


Figure 5-89 Volumes mapped to host cluster

Note: If there are no volumes mapped to the selected host cluster, the list will be empty.

4. Click **Add Volume Mappings** which will bring up a window listing the volumes which you can select to assign to host cluster as shown in Figure 5-90 on page 254.

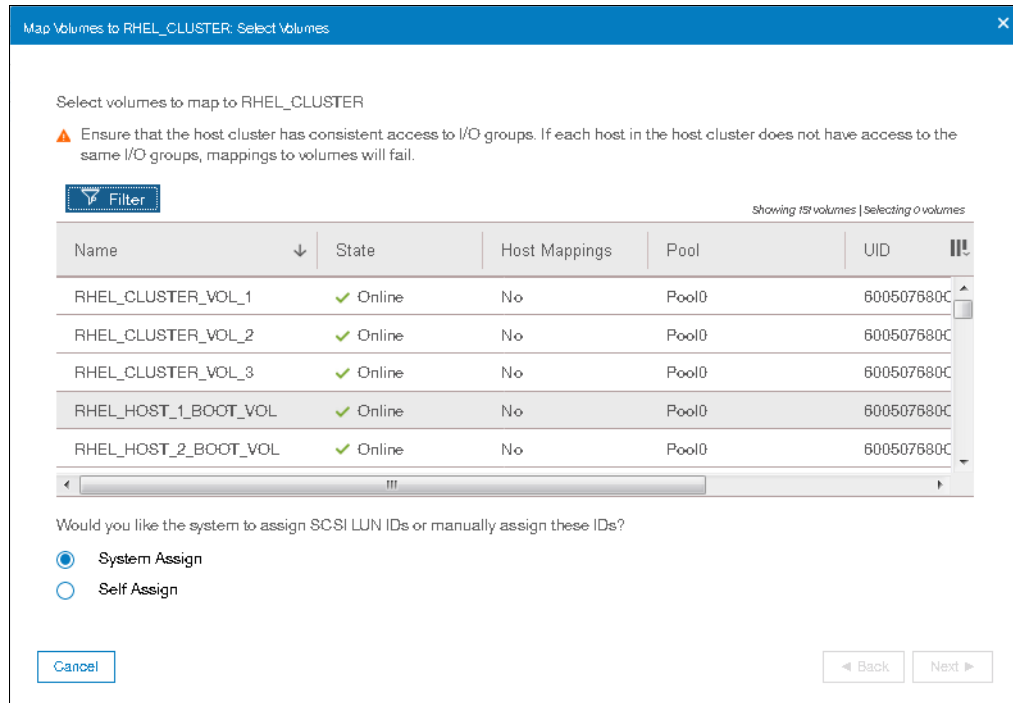


Figure 5-90 Volume list

5. Select the list of volumes to be mapped to the host cluster as shown in Figure 5-91.




Figure 5-91 Selecting the list of volumes

6. You can choose the SCSI LUN ID's be assigned automatically by the system or manually assign them. In this example, we chose the system assigned SCSI LUN ID's.

- Click **Next**. A summary screen with the list of volumes to be mapped will be shown as in Figure 5-92.

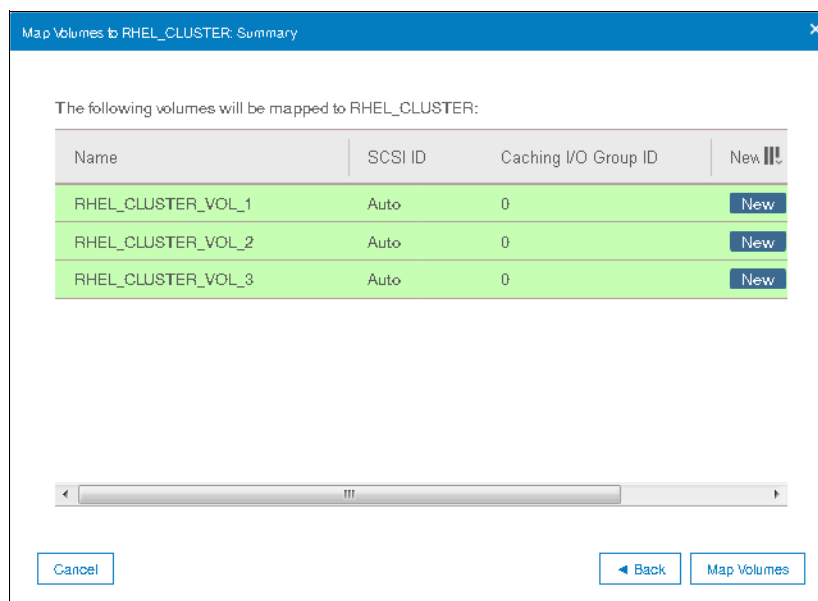


Figure 5-92 Summary of volumes to be mapped to host cluster

- Click **Map Volumes** and then the selected volumes will be mapped to the host cluster as shown in Figure 5-93.

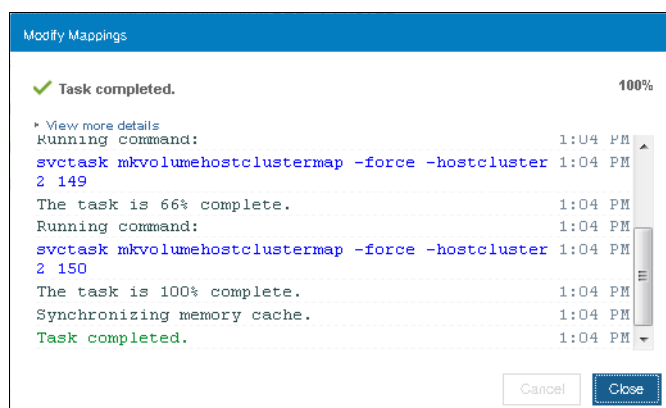


Figure 5-93 Volumes mapped to host cluster

5.6.5 Remove volume mapping from a host cluster

To remove a volume mapping from a host cluster, follow the steps as listed.

- From the **Host Clusters** pane, select the desired host cluster as shown in Figure 5-87 on page 252.

ID	Name	Status	Host Count	Mappings Count	Ports Count
0	CETV50000	Offline	2	75	2
1	w_CETV50000	Offline	2	14	2
2	RHEL_CLUSTER	Online	2	0	4

Figure 5-94 Host Clusters

2. Right click on the desired host cluster and select Modify Shared Volume Mappings as shown in Figure 5-88 on page 253.

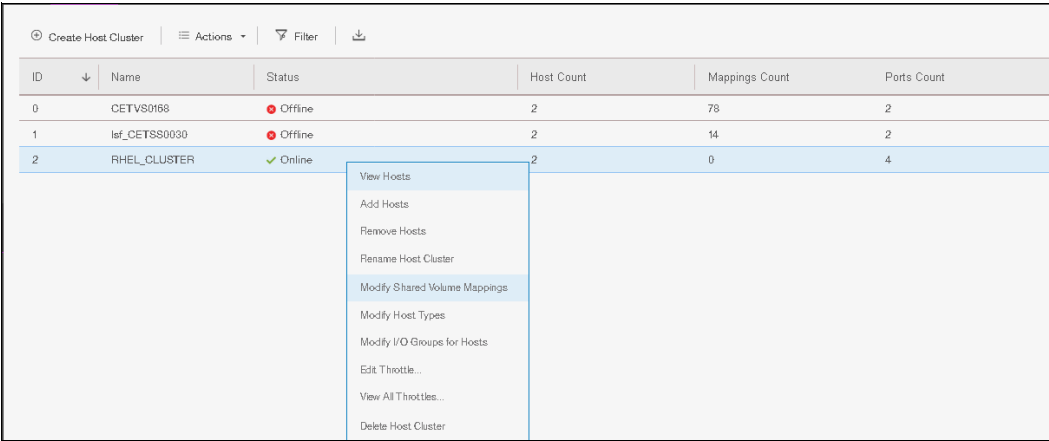


Figure 5-95 Modify Shared Volume Mapping for Host Cluster

3. A window showing list of volumes mapped to the host cluster is shown in Figure 5-96.

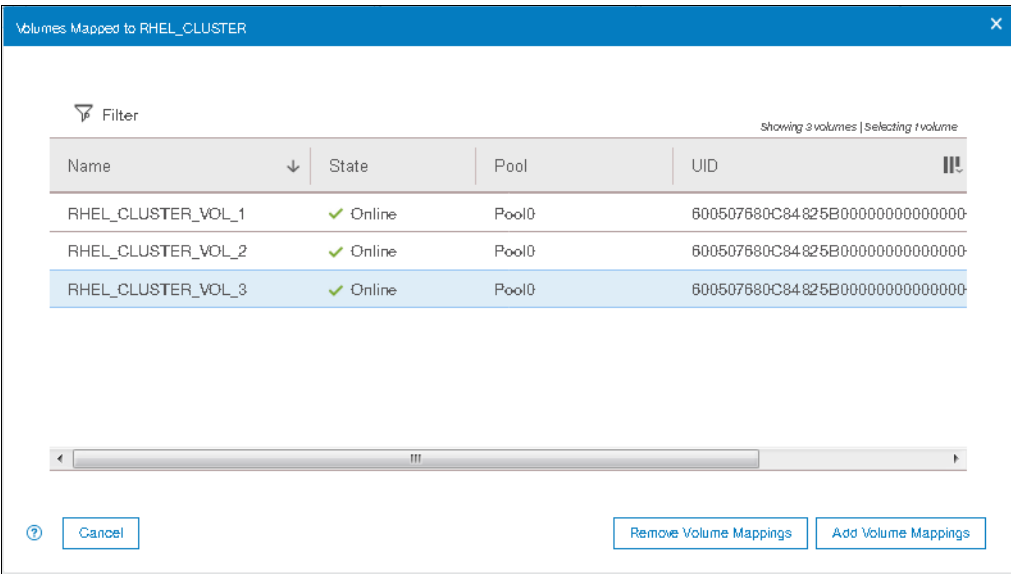


Figure 5-96 List of volumes mapped to a host cluster

4. Select the desired volume to be unmapped as shown in Figure 5-97 on page 257.




Figure 5-97 Selecting the volume to be unmapped

5. Click **Remove Volume Mappings**. A summary window listing the hosts will be shown as in Figure 5-98.

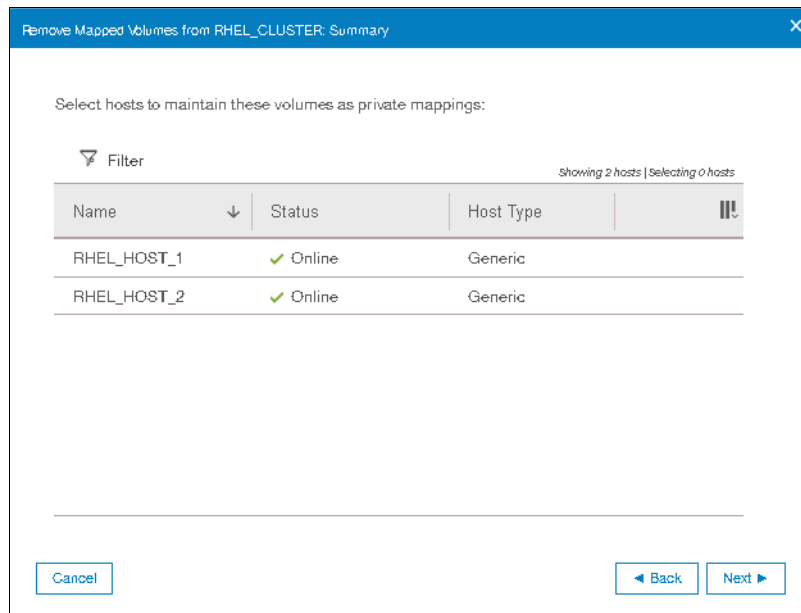


Figure 5-98 Host cluster member list

Note: At this point, you can select any host from the list to keep the private mapping between the selected host and the volume.

6. Click **Next**. A window will appear as shown in Figure 5-99 on page 258.

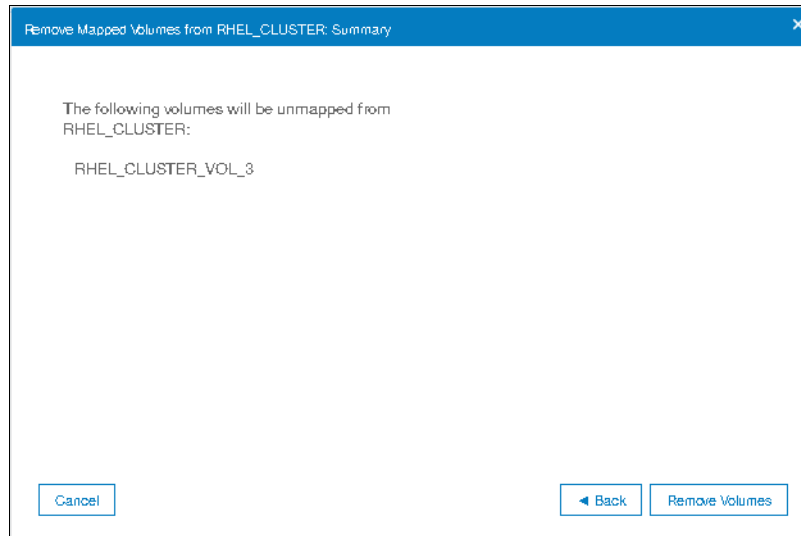


Figure 5-99 Volume unmap from host cluster

- Click **Remove Volumes**. A window with **Task Completed** message will be shown as in Figure 5-100.

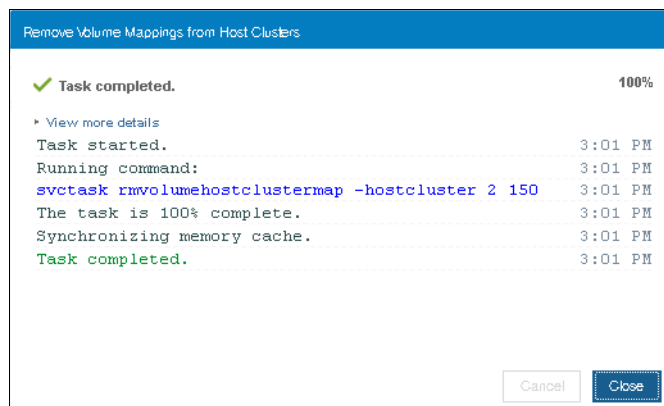


Figure 5-100 Volume unmapping from Host Cluster completed

5.6.6 Removing a host cluster member

To remove a host from the host cluster, follow the steps as listed.

- From the **Host Clusters** pane, select the desired host cluster as shown in Figure 5-101.

Create Host Cluster Actions Filter						
ID	Name	Status	Host Count	Mappings Count	Ports Count	
0	CETV00000	Offline	2	25	2	
1	wp_CETV000000	Offline	2	14	2	
2	RHEL_CLUSTER	Online	2	0	4	

Figure 5-101 Host Clusters

- Right click on the desired host cluster and select **Remove Hosts** as shown in Figure 5-102 on page 259

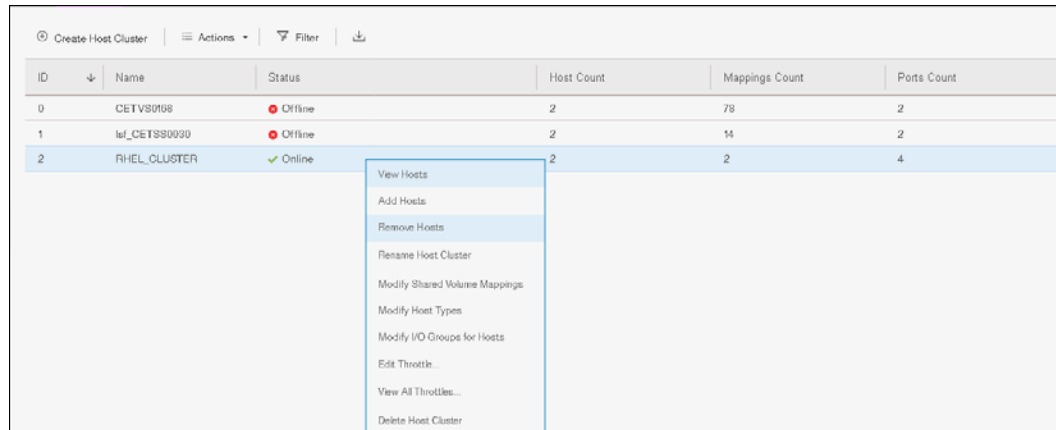


Figure 5-102 Removing hosts from host cluster

3. A window will be shown with list of hosts that are members of the host cluster as shown in Figure 5-103.

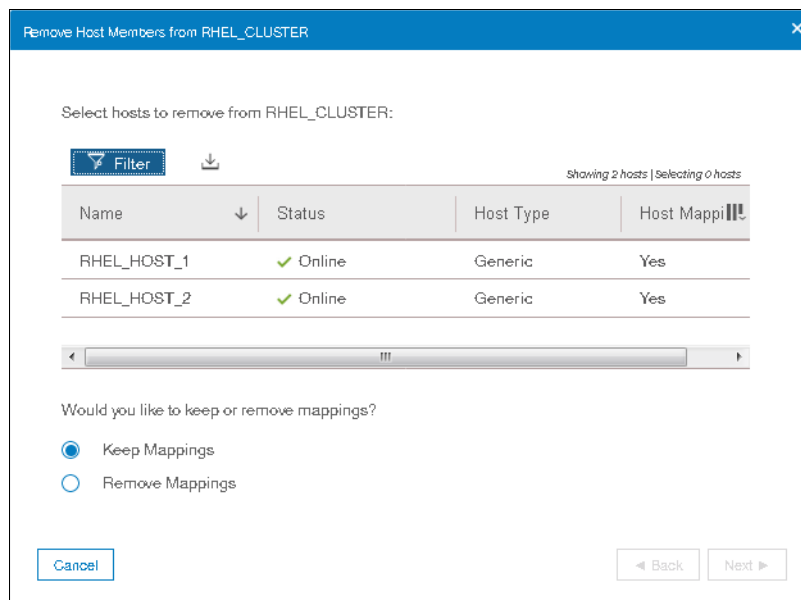


Figure 5-103 Host Cluster member selection

4. Select the host member that needs to be removed as shown in Figure 5-104 on page 260.




Figure 5-104 Host member selected

5. Select appropriate radio button to indicate whether you want to keep the mappings after the host member is removed from the host cluster or to remove those mappings. In our example we chose to **Remove Mappings** as shown in Figure 5-105.




Figure 5-105 Mapping selection during removal of host member from host cluster

Note: Select **Keep Mappings** to retain all the shared mappings in the host cluster as private mappings for the selected hosts. Select **Remove Mappings** to remove all the shared mappings if the host or hosts that is being removed no longer require access to these volumes.

6. Click **Next**. A window informing removal for the selected host member from host cluster will take place as shown in Figure 5-106 on page 261.

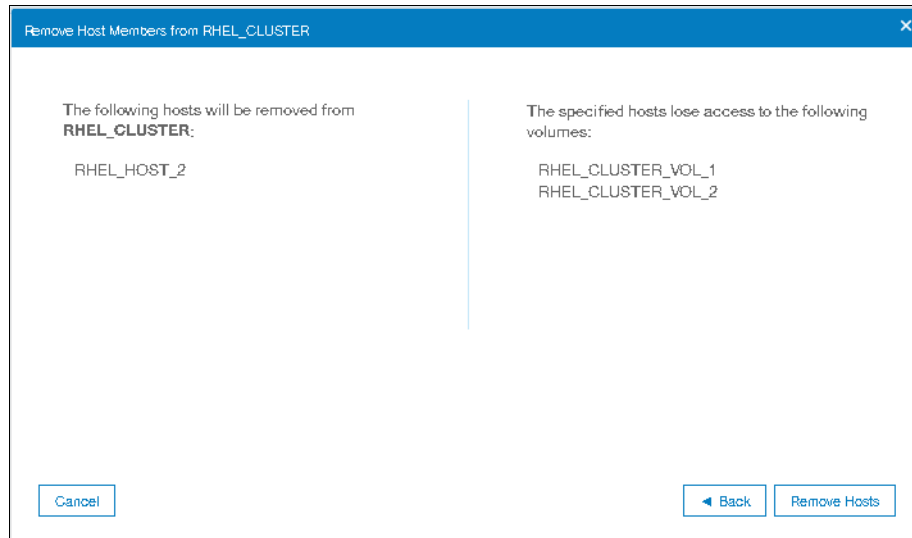


Figure 5-106 Confirmation window

7. Click **Remove Hosts**. A window indicating task completed will be shown as in Figure 5-107.

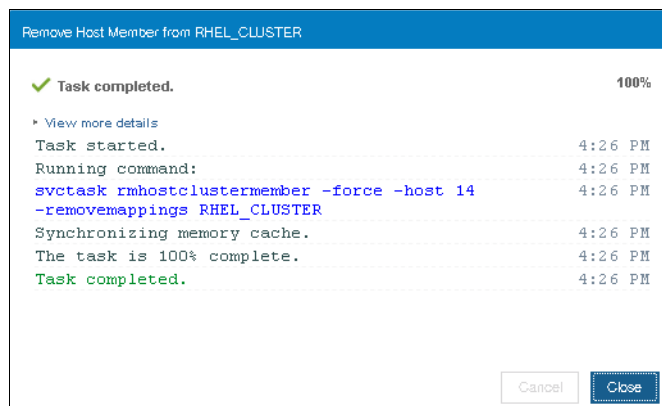


Figure 5-107 Host member removal task from host cluster completed

8. Click **Close**.

5.6.7 Removing a host cluster

To remove an existing host cluster, follow the steps listed.

1. From the **Host Clusters** pane, select the desired host cluster as shown in Figure 5-108.

ID	Name	Status	Host Count	Mappings Count	Ports Count
0	CEFS00000	Offline	2	75	2
1	sw_CEFSS00000	Offline	2	16	2
2	RHEL_CLUSTER	Online	2	0	4

Figure 5-108 Host Clusters

2. Right click on the desired host cluster and select **Delete Host Cluster** as shown in Figure 5-109 on page 262.

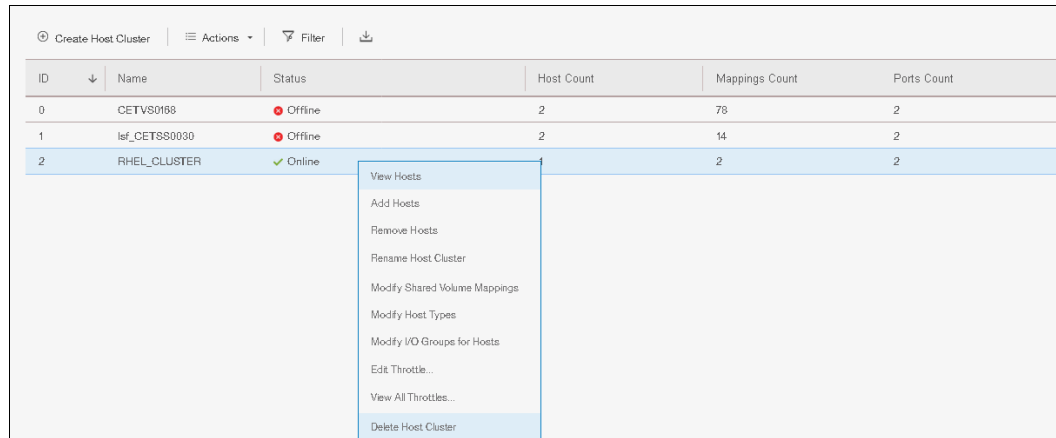


Figure 5-109 Delete Host Cluster selection

3. A dialog box will appear asking you to confirm the deletion of host cluster object along with your selection on mappings as shown in Figure 5-110.

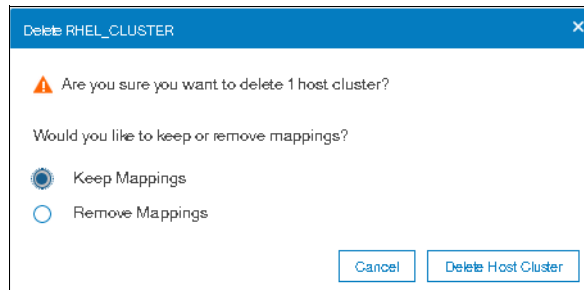


Figure 5-110 Confirm host cluster object deletion

4. Select the desired mappings via the radio button. In our example we chose to **Remove Mappings** as shown in Figure 5-111.




Figure 5-111 Remove mappings after host cluster deletion

Warning: Selecting **Remove Mappings** to remove all the shared mappings if the host or hosts that is being removed no longer require access to these volumes. So exercise this option with caution otherwise the server or servers that are part of the host cluster, will lose access to all the shared volumes.

5. Click **Delete Host Cluster**. You will see a window indicating the task completed successfully as shown in Figure 5-112 on page 263.

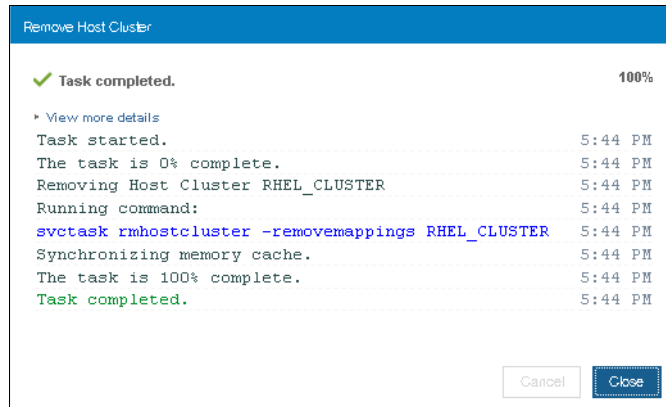


Figure 5-112 Host Cluster deletion task completed

6. Click **Close**.

5.6.8 I/O throttling for hosts and Host Clusters

You can set a limit on the number of I/O operations that are accepted by the storage system. The limit is known as the throttling rate and is set in terms of I/O operations per second (IOPS) or bandwidth. I/O throttling is a way to achieve quality of service (QoS). I/O throttling is a mechanism to limit the volume of I/O processed by the storage system at various levels resulting in better distribution of storage system resources. I/O throttling is also referred to as I/O governing.

In Spectrum Virtualize, by default, no I/O throttling rate is set. However, I/O throttling can be set at any of the following levels:

- ▶ Host
- ▶ Host clusters
- ▶ Volume
- ▶ MDisk group

When I/O throttling is set, the I/O rate is limited by queuing I/Os if it exceeds preset limits. I/O throttling does not guarantee minimum performance. Internal I/Os such as FlashCopy, Metro-Mirror, intra-cluster traffic are not throttled.

The following list illustrates some of the scenarios where I/O throttling can be beneficial:

- ▶ An aggressive host hogging bandwidth of the Spectrum Virtualize system can be limited by a throttle. For example, allow restricted I/Os from a data mining server than an application server.
- ▶ Restrict a group of hosts by their throttles. For example, department A gets more bandwidth than department B.
- ▶ Each volume can have a throttle defined. For example, a backup volume can have less bandwidth than a production volume).

In this section, we illustrate the process of setting I/O throttle on already defined hosts and host clusters.

Setting I/O throttle for host

To set I/O throttle on an already defined host, follow these steps:

1. Select **Hosts** option under the **Hosts** section on the main panel as shown in Figure 5-113 on page 264.




Figure 5-113 Hosts

2. Right click on the desired host and select **Edit Throttle** as shown in Figure 5-114.

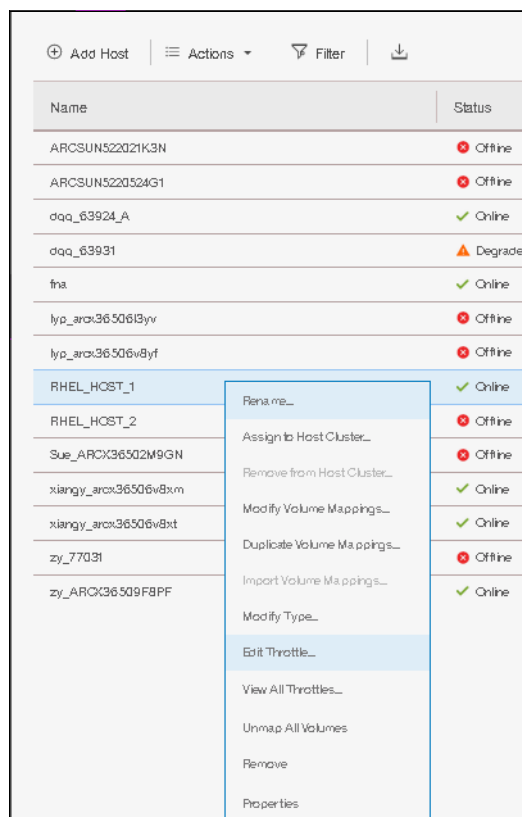


Figure 5-114 Selecting Edit Throttle option

3. Input the desired type of I/O throttle, either IOPS or Bandwidth. In our example we set up IOPs throttle by entering the **IOPS limit** as shown in Figure 5-115 on page 265.




Figure 5-115 Setting IOPS throttle

Note: While defining a throttle for a host, you can either define a throttle in terms of IOPS or bandwidth, but not both at the same time. If you want to have host throttle defined for IOPS and bandwidth both, then you have to define them one after the other.

4. Click **Create**. A window will open indicating that the task of setting up the throttle has completed as shown in Figure 5-116.

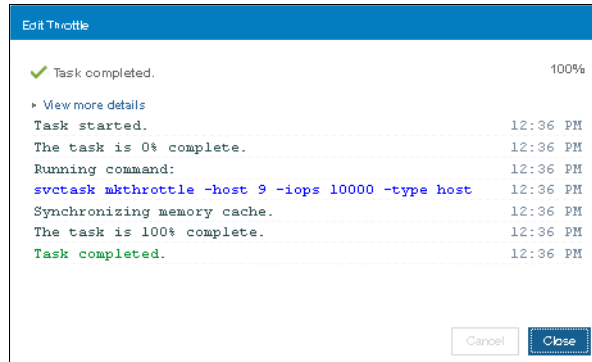


Figure 5-116 IOPS throttle set

5. Click **Close**.

Setting I/O throttle for Host Cluster

To set I/O throttle on an already defined Host Cluster, these steps can be followed:

1. Select **Hosts** option under the **Host Clusters** section on the main panel as shown in Figure 5-117.




Figure 5-117 Host Clusters

2. Right click on the desired host cluster and select **Edit Throttle** as shown in Figure 5-118.

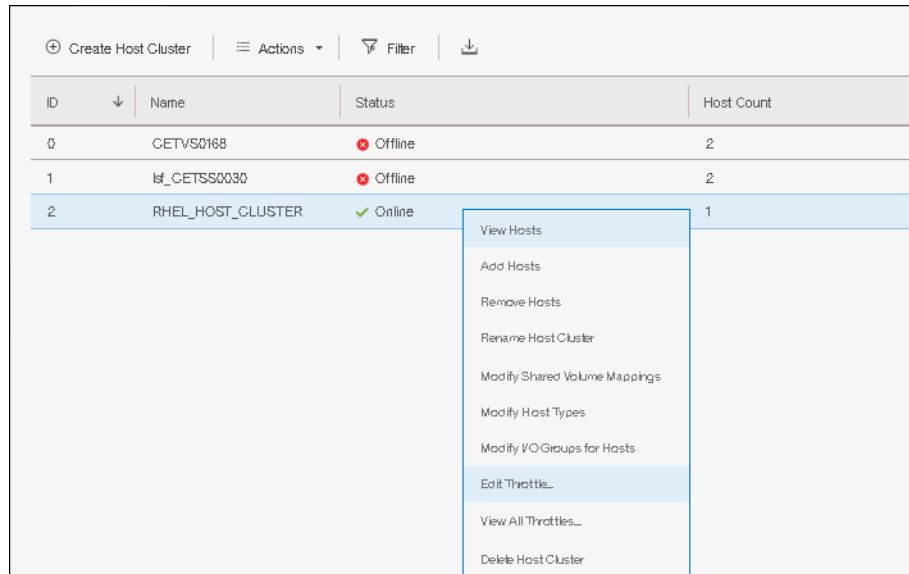


Figure 5-118 Selecting Edit Throttle option

- Input the desired type of I/O throttle, either IOPS or Bandwidth. In our example we set up IOPS throttle by entering the **IOPS limit** as shown in Figure 5-119.




Figure 5-119 Setting IOPS throttle

- Click **Create**. A window will open indicating the task of setting up the throttle completed as shown in Figure 5-120.

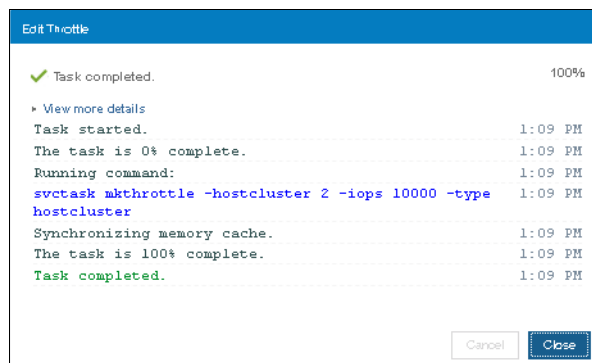


Figure 5-120 IOPS throttle set

- Click **Close**.

The following considerations have to be taken into account when defining the I/O throttle on host or Host Clusters:

- ▶ I/O throttle cannot be defined for host if it is a part of Host Cluster which already has an I/O throttle defined at the Host Cluster level.
- ▶ If the Host Cluster does not have an I/O throttle defined, its member hosts can have their individual I/O throttles defined.
- ▶ The mdiskgrp (storage pool) throttles for child pool and parent pool work independently.
- ▶ If a volume has multiple copies then throttling would be done for the mdiskgrp (storage pool) serving the primary copy. The throttling will not be applicable for the secondary pool for mirrored volumes and stretched cluster implementations.
- ▶ A host cannot be added to a Host Cluster if both of them have their individual throttles defined.
- ▶ A seeding host used for creating a Host Cluster cannot have a host throttle defined for it.

5.7 Proactive Host Failover

During planned maintenance procedures of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, such as firmware upgrade or node canister replacement, at times the host multipathing driver may experience issues recovering paths to volumes. Most of the issues seem to stem from the fact that the path to a volume was removed without warning to the host and the host has to detect that the path is now offline, time-out, and re-drive all I/Os that were in flight.

As a secondary issue, hosts often attempt to restart sending I/O to the preferred node canister as soon as the ports come online but the node canister may still be unpending. During this time, the node canister queues all incoming I/O until it has its configuration data and then volumes start coming online.

From a host point of view, the host would prefer to:

- ▶ Gracefully failover I/O to different node canister if the node canister it is using is about to pend to avoid expensive error recovery as a result of losing active paths that can impact business applications.
- ▶ Send I/O to one node canister even if the preferred node canister is unavailable as round-robinning across multiple node canisters affects application I/O performance.

To minimize the issues outlined here the Proactive Host Failover feature has been added since V7.8.1. Due to the Proactive Host Failover feature, the host multipath driver will get notification for node canister removal or node canister reboot during the planned maintenance procedures of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

Due to the notification received, the host will use the partner node canister for I/O and hence the I/O does not need to be timed-out and retried.

Note: Proactive Host Failover is an internal feature of Lenovo storage V-series software. There are no changes to the CLI or GUI for this feature.

The following points are worth noting in regards to Proactive Host Failover:

- ▶ When a Lenovo storage V-series system knows a node canister is about to go down, it will raise unit attentions to try and trigger host path failovers to surviving node canisters.
- ▶ Requires the host to be tracking the preferred paths - usually requires ALUA support.

- ▶ Delays the failback when a node canister is online until the node canister has confirmed it is ready to service I/O.
- ▶ Works with and without NPIV enabled and for all connection protocols that Lenovo storage V-series system supports.
- ▶ Adjusts preferred paths when node canisters are unavailable so the Lenovo storage V-series system should always be presenting a set of preferred paths.

Volume configuration

A volume is a logical disk provisioned out of a storage pool and is recognized by a host with a unique identifier (UID) field and a parameter list.

The first part of this chapter provides a brief overview of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 volumes, the classes of volumes available, and the topologies that they are associated with. It also provides an overview of advanced customization available.

The second part describes how to create volumes by using the GUI's *Quick* and *Advanced* volume creation menus, and shows you how to map these volumes to defined hosts.

The third part provides an introduction to the new volume manipulation commands, which are designed to facilitate the creation and administration of volumes used for *IBM HyperSwap* topology.

Note: Advanced host and volume administration, such as volume migration and creating volume copies, is described in Chapter 10, "Copy services" on page 451.

This chapter includes the following topics:

- ▶ 6.1, "Introduction to volumes" on page 270
- ▶ 6.2, "Create Volumes menu" on page 280
- ▶ 6.3, "Creating volumes using the Volume Creation" on page 285
- ▶ 6.4, "Mapping a volume to the host" on page 291
- ▶ 6.5, "Creating Custom volumes" on page 293
- ▶ 6.6, "HyperSwap and the mkvolume command" on page 301
- ▶ 6.7, "Mapping Volumes to Host after volume creation" on page 306
- ▶ 6.8, "Migrating a volume to another storage pool" on page 310
- ▶ 6.9, "Migrating volumes using the volume copy feature" on page 313
- ▶ 6.10, "I/O throttling" on page 317

6.1 Introduction to volumes

A volume is a logical disk that the system presents to attached hosts. For The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, the volume presented is from a virtual disk (VDisk). A volume is a discrete area of usable storage that has been virtualized, using Lenovo storage V-series code, from storage area network (SAN) storage that is managed by the Lenovo storage V-series cluster. The term virtual is used because the volume presented does not necessarily exist on a single physical entity.

Volumes have the following characteristics or attributes:

- ▶ Volumes can be created and deleted.
- ▶ Volumes can be resized (expanded or shrunk).
- ▶ Volume extents can be migrated at run time to another MDisk or storage pool.
- ▶ Volumes can be created as fully allocated or thin-provisioned. A conversion from a fully allocated to a thin-provisioned volume and vice versa can be done at run time.
- ▶ Volumes can be stored in multiple storage pools (mirrored) to make them resistant to disk subsystem failures or to improve the read performance.
- ▶ Volumes can be mirrored synchronously or asynchronously for longer distances. A Lenovo storage V-series system can run active volume mirrors to a maximum of three other Lenovo storage V-series systems, but not from the same volume.
- ▶ Volumes can be copied by using FlashCopy. Multiple snapshots and quick restore from snapshots (reverse FlashCopy) are supported.
- ▶ Volumes can be compressed.
- ▶ Volumes can be virtual. The system supports VMware vSphere Virtual Volumes, sometimes referred to as VVols, which allow VMware vCenter to manage system objects, such as volumes and pools. The system administrator can create these objects and assign ownership to VMware administrators to simplify management of these objects.

Note: A managed disk (MDisk) is a logical unit of physical storage. MDisks are either Redundant Arrays of Independent Disks (RAID) from internal storage, or external physical disks that are presented as a single logical disk on the SAN. Each MDisk is divided into several extents, which are numbered, from 0, sequentially from the start to the end of the MDisk. The extent size is a property of the storage pools that the MDisks are added to.

Attention: MDisks are not visible to host systems.

Volumes have two major modes: Managed mode and image mode. Managed mode volumes have two policies: The sequential policy and the striped policy. Policies define how the extents of a volume are allocated from a storage pool.

The *type* attribute of a volume defines the allocation of extents that make up the volume copy:

- ▶ A striped volume contains a volume copy that has one extent allocated in turn from each MDisk that is in the storage pool. This is the default option, but you can also supply a list of MDisks to use as the stripe set as shown in Figure 6-1 on page 271.

Attention: By default, striped volume copies are striped across all MDisks in the storage pool. If some of the MDisks are smaller than others, the extents on the smaller MDisks are used up before the larger MDisks run out of extents. Manually specifying the stripe set in this case might result in the volume copy not being created.

If you are unsure if there is sufficient free space to create a striped volume copy, select one of the following options:

- ▶ Check the free space on each MDisk in the storage pool by using the **lsfreeextents** command.
- ▶ Let the system automatically create the volume copy by not supplying a specific stripe set.




Figure 6-1 Striped extent allocation

- ▶ A *sequential* volume contains a volume copy that has extents that are allocated sequentially on one MDisk.
- ▶ *Image-mode* volumes are a special type of volume that have a direct relationship with one MDisk.

6.1.1 Image mode volumes

Image mode volumes are used to migrate LUNs that were previously mapped directly to host servers over to the control of the Lenovo storage V-series system. Image mode provides a one-to-one mapping between the logical block addresses (LBAs) between a volume and an MDisk. Image mode volumes have a minimum size of one block (512 bytes) and always occupy at least one extent.

An image mode MDisk is mapped to one, and only one, image mode volume.

The volume capacity that is specified must be equal to the size of the image mode MDisk. When you create an image mode volume, the specified MDisk must be in unmanaged mode and must not be a member of a storage pool. The MDisk is made a member of the specified storage pool (Storage Pool_IMG_XXX) as a result of creating the image mode volume.

The reverse process is also supported, in which a managed mode volume can be migrated to an image mode volume. If a volume is migrated to another MDisk, it is represented as being in managed mode during the migration, and is only represented as an image mode volume after it reaches the state where it is a straight-through mapping.

An image mode MDisk is associated with exactly one volume. If the (image mode) MDisk is not a multiple of the MDisk Group's extent size, the last extent is partial (not filled). An image mode volume is a pass-through one-to-one map of its MDisk. It cannot be a quorum disk and it does not have any metadata extents that are assigned to it from the Lenovo storage V-series system. Managed or image mode MDisk are always members of a storage pool.

It is a preferred practice to put image mode MDisk in a dedicated storage pool and use a special name for it (for example, Storage_Pool_IMG_xxx). The extent size that is chosen for this specific storage pool must be the same as the extent size into which you plan to migrate the data. All of the copy services functions can be applied to image mode disks. See Figure 6-2.

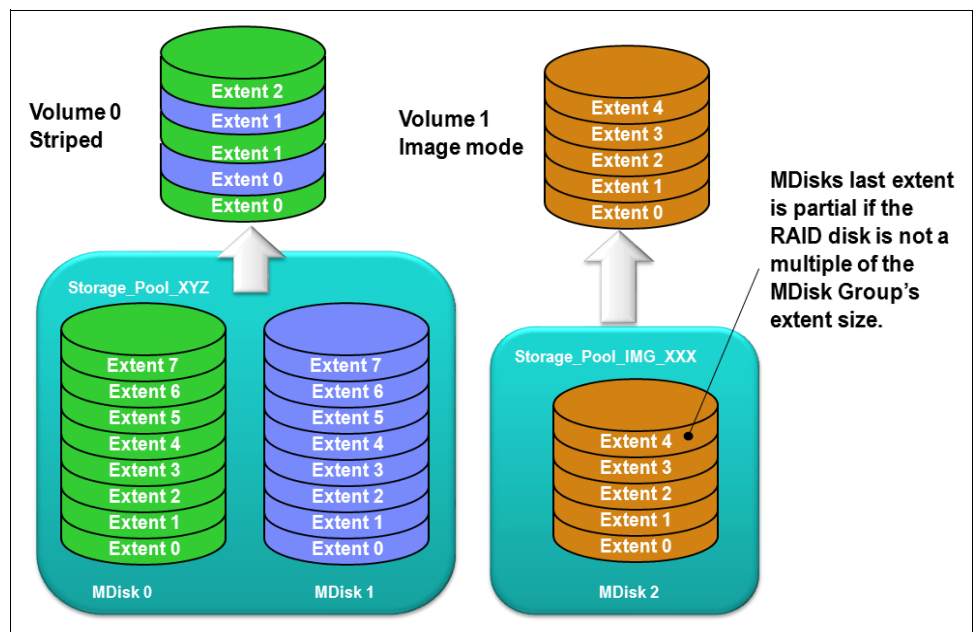


Figure 6-2 Image mode volume versus striped volume

6.1.2 Managed mode volumes

Volumes operating in managed mode provide a full set of virtualization functions. Within a storage pool, an arbitrary relationship between extents on (managed mode) volumes and extents on MDisk is supported. Each volume extent maps to exactly one MDisk extent.

Figure 6-3 on page 273 shows this mapping. It also shows a volume that consists of several extents that are shown as V0 - V7. Each of these extents is mapped to an extent on one of the MDisk: A, B, or C. The mapping table stores the details of this indirection.

Several of the MDisk extents are unused. No volume extent maps to them. These unused extents are available for use in creating volumes, migration, expansion, and so on.




Figure 6-3 Simple view of block virtualization

The allocation of a specific number of extents from a specific set of MDisks is performed by the following algorithm:

- If the set of MDisks from which to allocate extents contains more than one MDisk, extents are allocated from MDisks in a round-robin fashion.
- If an MDisk has no free extents when its turn arrives, its turn is missed and the round-robin moves to the next MDisk in the set that has a free extent.

When a volume is created, the first MDisk from which to allocate an extent is chosen in a pseudo-random way rather than by choosing the next disk in a round-robin fashion. The pseudo-random algorithm avoids the situation where the *striping effect* that is inherent in a round-robin algorithm that places the first extent for many volumes on the same MDisk.

Placing the first extent of several volumes on the same MDisk can lead to poor performance for workloads that place a large I/O load on the first extent of each volume, or that create multiple sequential streams.

6.1.3 Cache mode for volumes

It is also possible to define the cache characteristics of a volume. Under normal conditions, a volume's read and write data is held in the cache of its preferred node, with a mirrored copy of write data that is held in the partner node of the same I/O Group. However, it is possible to create a volume with cache disabled. This setting means that the I/Os are passed directly through to the back-end storage controller rather than being held in the node's cache.

Having cache-disabled volumes makes it possible to use the native copy services in the underlying RAID array controller for MDisks (LUNs) that are used as the image mode volumes. Using Copy Services rather than the underlying disk controller copy services gives better results.

Cache characteristics of a volume can have any of the following settings:

- `readwrite`. All read and write I/O operations that are performed by the volume are stored in cache. This is the default cache mode for all volumes.

- ▶ **readonly.** All read I/O operations that are performed by the volume are stored in cache.
- ▶ **disabled.** All read and write I/O operations that are performed by the volume are not stored in cache. Under normal conditions, a volume's read and write data is held in the cache of its preferred node, with a mirrored copy of write data that is held in the partner node of the same I/O Group. With cache disabled volume, the I/Os are passed directly through to the back-end storage controller rather than being held in the node's cache.

Note: Having cache-disabled volumes makes it possible to use the native copy services in the underlying RAID array controller for MDisk (LUNs) that are used as image mode volumes. Consult Lenovo Support before turning off the cache for volumes in production environment to avoid any performance degradation.

6.1.4 Mirrored volumes

The mirrored volume feature provides a simple RAID 1 function, so a volume has two physical copies of its data. This approach enables the volume to remain online and accessible even if one of the MDisk sustains a failure that causes it to become inaccessible.

The two copies of the volume often are allocated from separate storage pools or by using image-mode copies. The volume can participate in FlashCopy and remote copy relationships. It is serviced by an I/O Group, and has a preferred node.

Each copy is not a separate object and cannot be created or manipulated except in the context of the volume. Copies are identified through the configuration interface with a copy ID of their parent volume. This copy ID can be 0 or 1.

This feature provides a point-in-time copy function that is achieved by “splitting” a copy from the volume. However, the mirrored volume feature does not address other forms of mirroring that are based on remote copy, which is sometimes called *IBM HyperSwap*, that mirrors volumes across I/O Groups or clustered systems. It is also not intended to manage mirroring or remote copy functions in back-end controllers.

Figure 6-4 on page 275 provides an overview of volume mirroring.




Figure 6-4 Volume mirroring overview

A second copy can be added to a volume with a single copy or removed from a volume with two copies. Checks prevent the accidental removal of the only remaining copy of a volume. A newly created, unformatted volume with two copies initially has the two copies in an out-of-synchronization state. The primary copy is defined as “fresh” and the secondary copy is defined as “stale.”

The synchronization process updates the secondary copy until it is fully synchronized. This update is done at the default *synchronization rate* or at a rate that is defined when the volume is created or modified. The synchronization status for mirrored volumes is recorded on the quorum disk.

If a two-copy mirrored volume is created with the **format** parameter, both copies are formatted in parallel, and the volume comes online when both operations are complete with the copies in sync.

If mirrored volumes are expanded or shrunk, all of their copies are also expanded or shrunk.

If it is known that MDisk space (which is used for creating copies) is already formatted or if the user does not require read stability, a *no synchronization* option can be selected that declares the copies as synchronized (even when they are not).

To minimize the time that is required to resynchronize a copy that is out of sync, only the 256 kibibyte (KiB) grains that were written to since the synchronization was lost are copied. This approach is known as an *incremental synchronization*. Only the changed grains must be copied to restore synchronization.

Important: An unmirrored volume can be migrated from one location to another by adding a second copy to the wanted destination, waiting for the two copies to synchronize, and then removing the original copy 0. This operation can be stopped at any time. The two copies can be in separate storage pools with separate extent sizes.

When there are two copies of a volume, one copy is known as the *primary copy*. If the primary is available and synchronized, reads from the volume are directed to it. The user can select the primary when the volume is created or can change it later.

Placing the primary copy on a high-performance controller maximizes the read performance of the volume.

Write I/O operations data flow with a mirrored volume

For write I/O operations to a mirrored volume, the preferred node definition, with the multipathing driver on the host, is used to determine the preferred path. The host routes the I/Os through the preferred path, and the corresponding node is responsible for further destaging written data from cache to both volume copies. Figure 6-5 shows the data flow for write I/O processing when volume mirroring is used.




Figure 6-5 Data flow for write I/O processing in a mirrored volume

As shown in Figure 6-5, all the writes are sent by the host to the preferred node for each volume (1). Then, the data is mirrored to the cache of the partner node in the I/O Group (2), and acknowledgment of the write operation is sent to the host (3). The preferred node then destages the written data to the two volume copies (4).

A volume with copies can be checked to see whether all of the copies are identical or consistent. If a medium error is encountered while it is reading from one copy, it is repaired by using data from the other copy. This consistency check is performed asynchronously with host I/O.

Important: Mirrored volumes can be taken offline if no quorum disk is available. This behavior occurs because the synchronization status for mirrored volumes is recorded on the quorum disk.

Mirrored volumes use bitmap space at a rate of 1 bit per 256 KiB grain, which translates to 1 MiB of bitmap space supporting 2 TiB of mirrored volumes. The default allocation of bitmap space is 20 MiB, which supports 40 TiB of mirrored volumes. If all 512 MiB of variable bitmap space is allocated to mirrored volumes, 1 PiB of mirrored volumes can be supported.

Table 6-1 on page 277 shows you the bitmap space default configuration.

Table 6-1 Bitmap space default configuration

Copy service	Minimum allocated bitmap space	Default allocated bitmap space	Maximum allocated bitmap space	Minimum ^a functionality when using the default values
Remote copy ^b	0	20 MiB	512 MiB	40 TiB of remote mirroring volume capacity
FlashCopy ^c	0	20 MiB	2 GiB	<ul style="list-style-type: none"> ▶ 10 TiB of FlashCopy source volume capacity ▶ 5 TiB of incremental FlashCopy source volume capacity
Volume mirroring	0	20 MiB	512 MiB	40 TiB of mirrored volumes
RAID	0	40 MiB	512 MiB	<ul style="list-style-type: none"> ▶ 80 TiB array capacity using RAID 0, 1, or 10 ▶ 80 TiB array capacity in three-disk RAID 5 array ▶ Slightly less than 120 TiB array capacity in five-disk RAID 6 array

a. The actual amount of functionality might increase based on settings such as grain size and strip size. RAID is subject to a 15% margin or error.

b. Remote copy includes Metro Mirror, Global Mirror, and active-active relationships.

c. FlashCopy includes the FlashCopy function, Global Mirror with change volumes, and active-active relationships.

The sum of all bitmap memory allocation for all functions except FlashCopy must not exceed 552 MiB.

6.1.5 Thin-provisioned volumes

Volumes can be configured to be thin-provisioned or fully allocated. A *thin-provisioned* volume behaves as though application reads and writes were fully allocated. When a thin-provisioned volume is created, the user specifies two capacities:

- ▶ The real physical capacity that is allocated to the volume from the storage pool
- ▶ Its virtual capacity that is available to the host

In a *fully allocated* volume, these two values are the same.

Therefore, the real capacity determines the quantity of MDisk extents that is initially allocated to the volume. The *virtual capacity* is the capacity of the volume that is reported to all other components (for example, FlashCopy, cache, and remote copy), and to the host servers.

The *real capacity* is used to store the user data and the metadata for the thin-provisioned volume. The real capacity can be specified as an absolute value, or as a percentage of the virtual capacity.

Thin-provisioned volumes can be used as volumes that are assigned to the host, by FlashCopy to implement thin-provisioned FlashCopy targets, and with the mirrored volumes feature.

When a thin-provisioned volume is initially created, a small amount of the real capacity is used for initial metadata. I/Os are written to grains of the thin volume that were not previously written, which causes grains of the real capacity to be used to store metadata and the actual user data. I/Os are written to grains that were previously written, which updates the grain where data was previously written.

The grain size is defined when the volume is created. The grain size can be 32 KiB, 64 KiB, 128 KiB, or 256 KiB. The default grain size is 256 KiB, which is the preferred option. If you select 32 KiB for the grain size, the volume size cannot exceed 260 TiB. The grain size cannot be changed after the thin-provisioned volume is created. Generally, smaller grain sizes save space, but they require more metadata access, which can adversely affect performance.

When using thin-provisioned volume as a FlashCopy source or target volume, use 256 KiB to maximize performance. When using thin-provisioned volume as a FlashCopy source or target volume, specify the same grain size for the volume and for the FlashCopy function.

Figure 6-6 shows the thin-provisioning concept.

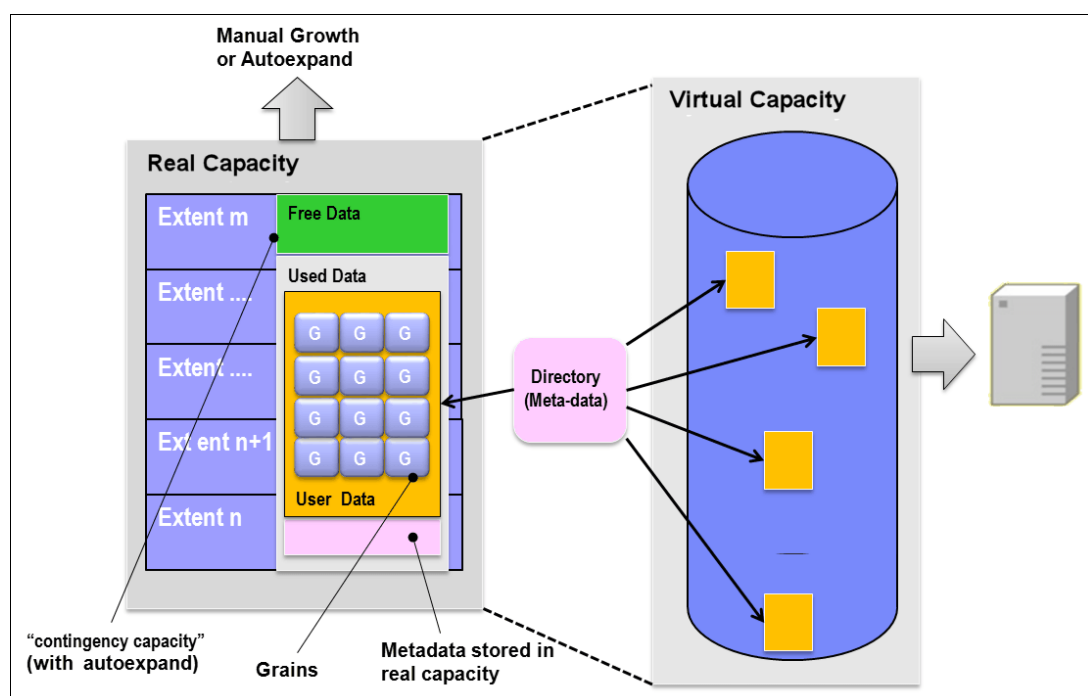


Figure 6-6 Conceptual diagram of thin-provisioned volume

Thin-provisioned volumes store user data and metadata. Each grain of data requires metadata to be stored. Therefore, the I/O rates that are obtained from thin-provisioned volumes are less than the I/O rates that are obtained from fully allocated volumes.

The metadata storage used is never greater than 0.1% of the user data. The resource usage is independent of the virtual capacity of the volume. If you are using the thin-provisioned volume directly with a host system, use a small grain size.

Thin-provisioned volume format: Thin-provisioned volumes do not need formatting. A read I/O that requests data from deallocated data space returns zeros. When a write I/O causes space to be allocated, the grain is “zeroed” before use.

The real capacity of a thin volume can be changed if the volume is not in image mode. Increasing the real capacity enables a larger amount of data and metadata to be stored on the volume. Thin-provisioned volumes use the real capacity that is provided in ascending order as new data is written to the volume. If the user initially assigns too much real capacity to the volume, the real capacity can be reduced to free storage for other uses.

A thin-provisioned volume can be configured to *autoexpand*. This feature causes the controller firmware to automatically add a fixed amount of more real capacity to the thin volume as required. Therefore, autoexpand attempts to maintain a fixed amount of unused real capacity for the volume, which is known as the *contingency capacity*.

The contingency capacity is initially set to the real capacity that is assigned when the volume is created. If the user modifies the real capacity, the contingency capacity is reset to be the difference between the used capacity and real capacity.

A volume that is created without the autoexpand feature, and therefore has a zero contingency capacity, goes offline when the real capacity is used and it must expand.

Autoexpand does not cause the real capacity to grow much beyond the virtual capacity. The real capacity can be manually expanded to more than the maximum that is required by the current virtual capacity, and the contingency capacity is recalculated.

To support the auto expansion of thin-provisioned volumes, the storage pools from which they are allocated have a configurable capacity warning. When the used capacity of the pool exceeds the warning capacity, a warning event is logged. For example, if a warning of 80% is specified, the event is logged when 20% of the free capacity remains.

A thin-provisioned volume can be converted nondisruptively to a fully allocated volume (or vice versa) by using the volume mirroring function. For example, the system allows a user to add a thin-provisioned copy to a fully allocated primary volume, and then remove the fully allocated copy from the volume after they are synchronized.

The fully allocated-to-thin-provisioned migration procedure uses a zero-detection algorithm so that grains that contain all zeros do not cause any real capacity to be used.

6.1.6 Compressed volumes

This is a custom type of volume where data is compressed as it is written to disk, saving additional space. Compression is a separately orderable license that is set on a per enclosure basis. One license is required for each control or expansion enclosure and each enclosure in any external storage systems that use virtualization. To use the compression function, you must obtain the IBM Real-time Compression license.

Note: For Lenovo Storage V3700 V2 and V3700 V2 XP, only Lenovo Storage V5030 model supports compression.

6.1.7 Volumes for various topologies

A *Basic* volume is the simplest form of volume. It consists of a single volume copy, made up of extents *striped* across all MDisks in a storage pool. It services I/O by using *readwrite* cache and is classified as *fully allocated* (reported real capacity and virtual capacity are equal). You can create other forms of volumes, depending on the type of topology that is configured on your system:

- With *standard topology*, which is a single-site configuration, you can create a *basic* volume or a *mirrored* volume.

By using volume mirroring, a volume can have two physical copies. Each volume copy can belong to a different pool, and each copy has the same virtual capacity as the volume. In the management GUI, an asterisk indicates the primary copy of the mirrored volume. The primary copy indicates the preferred volume for read requests.

- With *HyperSwap topology*, which is a three-site HA configuration, you can create a basic volume or a *HyperSwap* volume.

HyperSwap volumes create copies on separate sites for systems that are configured with HyperSwap topology. Data that is written to a HyperSwap volume is automatically sent to both copies so that either site can provide access to the volume if the other site becomes unavailable.

Note: For Lenovo Storage V3700 V2, V3700 V2 XP and V5030, the HyperSwap topology is supported only on Lenovo Storage V5030 system.

- Virtual Volumes (VVols): The controller firmware V7.6.0 release also introduces *Virtual Volumes*. These volumes are available in a system configuration that supports VMware vSphere Virtual Volumes. These volumes allow VMware vCenter to manage system objects, such as volumes and pools. The Spectrum Virtualize system administrators can create volume objects of this class, and assign ownership to VMware administrators to simplify management.

Note: From V7.4.0 onwards, it is possible to prevent accidental deletion of volumes, if they have recently performed any I/O operations. This feature is called *Volume protection*, and it prevents active volumes or host mappings from being deleted inadvertently. This is done by using a global system setting. For more information, see Lenovo Information Center:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_volprotection.html

6.2 Create Volumes menu

The GUI is the simplest means of volume creation, and presents different options in the **Create Volumes** menu depending on the topology of the system.

To start the process of creating a volume, follow the steps listed.

1. Click on the **Volumes** on the main panel as shown in Figure 6-7 on page 281.

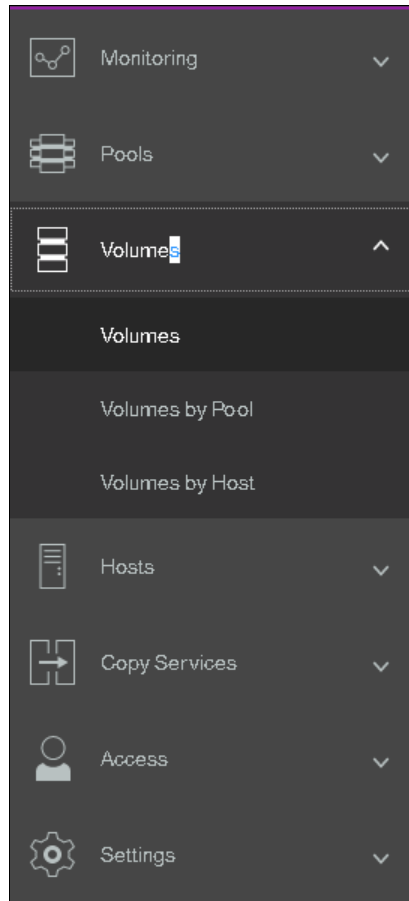


Figure 6-7 Volumes

2. Click **Volumes** in the sub-menu option as shown in Figure 6-8.




Figure 6-8 Volumes sub-menu option

3. On the right pane, you will get a list of existing volumes, if any, as shown in Figure 6-9 on page 282.

Name	State	Synchronized	Pool	UID	Host Mappings	Capacity
test_vol_1	✓ Online		Pool0	6005070300A600000000000000000000	No	1.00 GiB

Figure 6-9 Existing volumes

- Click **Create Volumes** as shown in Figure 6-10.

Name	State	Synchronized	Pool	UID	Host Mappings	Capacity
test_vol_1	✓ Online		Pool0	6005070300A600000000000000000000	No	1.00 GiB

Figure 6-10 Create Volumes

- Depending on the topology of the system, the **Create Volumes** button will provide different options. For standard topology, **Create Volumes** will pop up window with options to create basic, mirrored or custom volume as shown in Figure 6-11.

Create Volumes

Basic Mirrored Custom

Create a preset volume with all the basic features.

Pool: Pool0 Total 2.17 TiB

Volume Details

Quantity: 1 Capacity: x GiB Capacity savings: None Name:

I/O group: Automatic

Summary
Fields Incomplete

Cancel Create and Map Create

Figure 6-11 Create Volume options for standard topology

For HyperSwap topology, **Create Volumes** option will pop up window with options to create basic, HyperSwap or custom volume as shown in Figure 6-12.

Create Volumes

Basic

HyperSwap

Custom

Create a preset volume with all the basic features.

Pool:

Click to select.

Total 0 bytes

Volume Details

Quantity:

1

Capacity:

GiB

Capacity savings:

None

Name:

I/O group:

Automatic

Summary

Fields Incomplete

?

Cancel

Create and Map

Create

Figure 6-12 Create Volume options for HyperSwap topology

Clicking any of the three choices in the **Create Volumes** window opens a drop-down window where volume details can be entered. The example shown in Figure 6-13 on page 284 shows a *Basic* volume to demonstrate this view.

Create Volumes

Basic Mirrored Custom

Create a preset volume with all the basic features.

Pool:

Pool1 Total 2.42 TiB

Volume Details

Quantity: 1 **Capacity:** 1 GiB **Capacity savings:** None **Name:** ITSO_BASIC_VOL

I/O group: Automatic

Summary

1 volume
Volume name: ITSO_BASIC_VOL

1 volume in pool Pool1

Caching I/O group: Automatic
Accessible I/O group: Automatic

Total real capacity: 1.00 GiB
Total virtual capacity: 1.00 GiB

? Cancel Create and Map Create

Figure 6-13 Basic Volume

Notes:

- ▶ A Basic volume is a volume whose data is striped across all available managed disks (MDisks) in one storage pool.
- ▶ A Mirrored volume is a volume with two physical copies, where each volume copy can belong to a different storage pool.
- ▶ A Custom volume, in the context of this menu, is either a Basic or Mirrored volume with customization from the default parameters

Volume Creation also provides, using the **Capacity Savings** parameter, the ability to change the default provisioning of a Basic or Mirrored Volume to Thin-provisioned or Compressed.

Volume migration is described in 6.8, “Migrating a volume to another storage pool” on page 310. Creating volume copies is described in 6.3.2, “Creating Mirrored volumes using Volume Creation” on page 288.

6.3 Creating volumes using the Volume Creation

This section focuses on using the **Volume Creation** operation to create Basic and Mirrored volumes in a system with standard topology. It also covers creating host-to-volume mapping. As previously stated, Volume Creation is available on four different volume classes:

- ▶ Basic
- ▶ Mirrored
- ▶ Custom
- ▶ HyperSwap

Note: The ability to create HyperSwap volumes using the GUI simplifies creation and configuration. This simplification is enhanced by the GUI using the `mkvolume` command.

6.3.1 Creating Basic volumes using Volume Creation

The most commonly used type of volume is the Basic volume. This type of volume is fully provisioned, with the entire size dedicated to the defined volume. The host and the Lenovo storage V-series system see the fully allocated space.

Create a Basic volume by clicking the **Basic** icon as shown in Figure 6-14 on page 286. This action opens an additional input window where you can define the following information:

- ▶ Pool: The pool in which the volume is created (drop-down)
- ▶ Quantity: The number of volumes to be created (numeric up/down)
- ▶ Capacity: Size of the volume in units (drop-down)
- ▶ Capacity Savings (drop-down):
 - None
 - Thin-provisioned
 - Compressed
- ▶ Name: Name of the volume (cannot start with a numeric)
- ▶ I/O group

The Basic volume creation process is shown in Figure 6-14 on page 286.

Create Volumes

Basic Mirrored Custom

Create a preset volume with all the basic features.

Pool: Pool1

Volume Details

Quantity: 1 Capacity: 1 GiB Capacity savings: None Name: ITSO_BASIC_VOL (+)

I/O group: Automatic

Summary

1 volume

Volume name: ITSO_BASIC_VOL

1 volume in pool Pool1

Caching I/O group: Automatic

Accessible I/O group: Automatic

Total real capacity: 1.00 GiB

Total virtual capacity: 1.00 GiB

Cancel Create and Map Create

Figure 6-14 Creating Basic volume

An appropriate naming convention is recommended for volumes for easy identification of their association with the host or host cluster. At a minimum, it should contain the name of the pool or some tag that identifies the underlying storage subsystem. It can also contain the host name that the volume is mapped to, or perhaps the content of this volume, for example, name of applications to be installed.

When all of the characteristics of the Basic volume have been defined, it can be created by selecting one of the following options:

- ▶ Create
- ▶ Create and Map to Host

Note: The Plus sign (+) icon highlighted in green in Figure 6-14, can be used to create more volumes in the same instance of the volume creation wizard.

In this example, the **Create** option has been selected (the volume-to-host mapping can be performed later). At the end of the volume creation, following confirmation window appears as shown in Figure 6-15 on page 287.

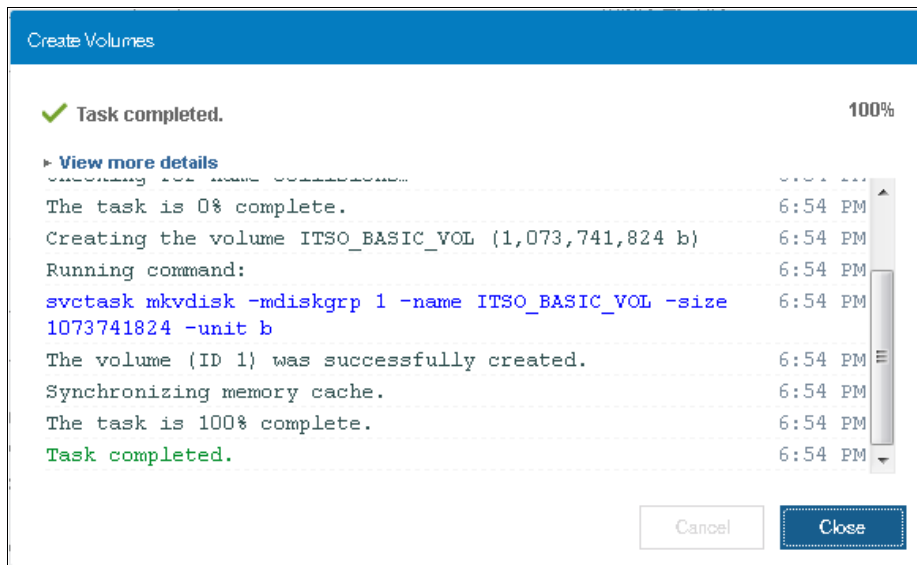


Figure 6-15 Create Volume Task Completion window: Success

Success is also indicated by the state of the Basic volume being reported as formatting in the Volumes pane as shown in Figure 6-16.

Create Volumes							
Actions All Volumes Filter							
Name	State	Synchronized	Pool	UID	Host Mappings	Capacity	
ITSO_BASIC_VOL	✓ Online (formatting)		Pool0	600507E300A600000000000000000002		No	1.00 GB
test_vol_1	✓ Online		Pool0	600507E300A600000000000000000000		No	1.00 GB

Figure 6-16 Basic Volume Fast-Format

Notes:

- Fully allocated volumes are automatically formatted through the quick initialization process after the volume is created. This process makes fully allocated volumes available for use immediately.
- Quick initialization requires a small amount of I/O to complete, and limits the number of volumes that can be initialized at the same time. Some volume actions, such as moving, expanding, shrinking, or adding a volume copy, are disabled when the specified volume is initializing. Those actions are available after the initialization process completes.
- The quick initialization process can be disabled in circumstances where it is not necessary. For example, if the volume is the target of a Copy Services function, the Copy Services operation formats the volume. The quick initialization process can also be disabled for performance testing, so that the measurements of the raw system capabilities can take place without waiting for the process to complete.

For more information, see Lenovo Information Center:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_fullyallocatedvolumes.html

6.3.2 Creating Mirrored volumes using Volume Creation

Lenovo storage V series offer the capability to mirror volumes, which means a single volume, 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 as the volume. When a server writes to a mirrored volume, the system writes the data to both copies. When a server reads a mirrored volume, the system picks one of the copies to read.

Normally this is the primary copy (as indicated in the management GUI 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 are written and resynchronizes these areas when both copies are available.

The use of mirrored volumes results in the following outcomes:

- ▶ Improves availability of volumes by protecting them from a single storage system failure
- ▶ Provides concurrent maintenance of a storage system that does not natively support concurrent maintenance
- ▶ Provides an alternative method of data migration with better availability characteristics
- ▶ Converts between fully allocated volumes and thin-provisioned volumes

Note: Volume mirroring is not a true disaster recovery (DR) solution, because both copies are accessed by the same node pair and addressable by only a single cluster, but it can improve availability.

To create a mirrored volume, complete the following steps:

1. In the Create Volumes window, click **Mirrored** and in the **Mirrored copies** subsection, choose the **Pool** of **Copy1** and **Copy2** by using the drop-down menu. Although the mirrored volume can be created in the same pool, this setup is not typical. Next, in the **Volume Details**, enter **Quantity**, **Capacity**, **Capacity savings**, and **Name**.

Generally, keep mirrored volumes on a separate set of physical disks (Pools). Leave the **I/O group** option at its default setting of **Automatic** (see Figure 6-17 on page 289).

Create Volumes

Basic **Mirrored** Custom

Create preset volumes with copies in multiple pools but at a single site.

Mirrored copies

Pool:

Copy 1: Pool1 Total 2.42 TiB

Copy 2: Pool2 Total 2.36 TiB

Volume Details

Quantity: 1 Capacity: 1 GiB Capacity savings: None Name: ITSO_MIR_VOL-

I/O group: Automatic

Summary

1 volume

Volume name: ITSO_MIR_VOL-

1 volume

2 mirrored copies

1 copy in pool Pool1

1 copy in pool Pool2

Caching I/O group: Automatic

Accessible I/O group: Automatic

Total real capacity: 1.00 GiB

Total virtual capacity: 1.00 GiB

Cancel Create and Map **Create**

Figure 6-17 Mirrored Volume creation

2. Click **Create** (or **Create and Map to Host**)
3. Next, the GUI displays the underlying CLI commands being run to create the mirrored volume and indicates completion as shown in Figure 6-18 on page 290.

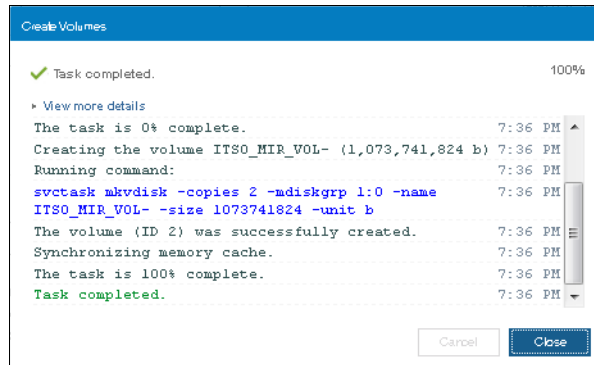


Figure 6-18 Task complete: creating Mirrored volume

Note: When creating a Mirrored volume by using this menu, you are not required to specify the Mirrored Sync rate. It defaults to 2 MBps. Customization of this synchronization rate can be done by using the **Custom** option.

Volume Creation with Capacity Saving options

The Volume Creation operation also provides, using the **Capacity Savings** parameter, the ability to alter the provisioning of a Basic or Mirrored volume into Thin-provisioned or Compressed. Select either **Thin-provisioned** or **Compressed** from the drop-down menu as shown in Figure 6-19.

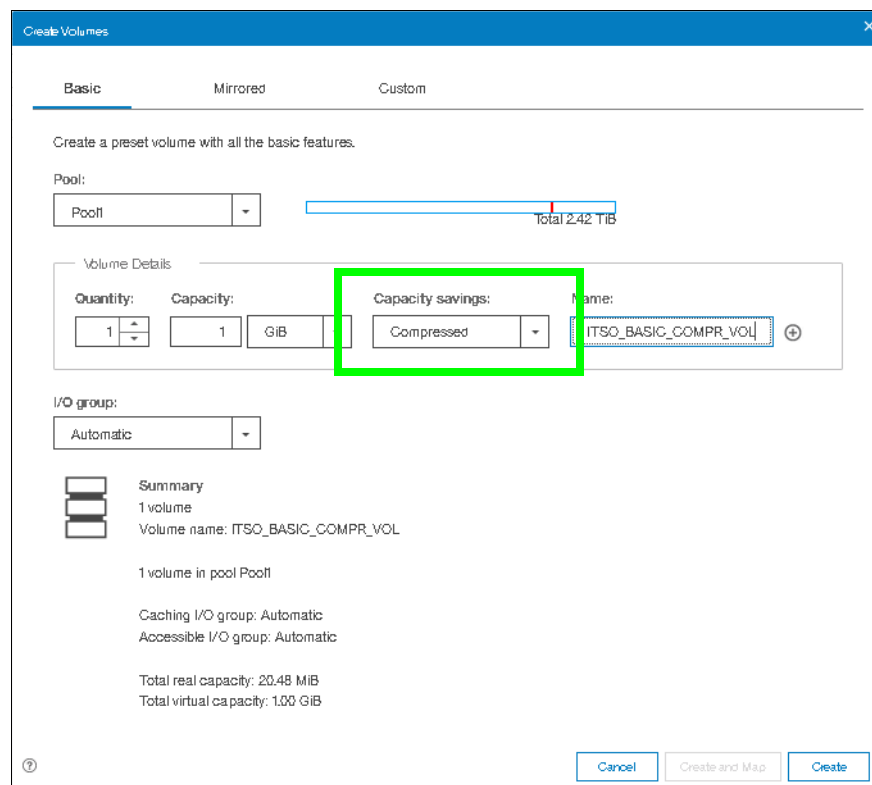


Figure 6-19 Volume Creation with Capacity Saving option set to Compressed

Alternatively, select **Thin-provisioned** from the menu to define a Thin-provisioned volume.

6.4 Mapping a volume to the host

After a volume is created, it can be mapped to a host:

1. From the Volumes menu, highlight the volume that you want to create a mapping for and then select **Actions** from the menu bar.

Tip: An alternative way of opening the **Actions** menu is to highlight (select) a volume and use the right mouse button.

2. From the Actions menu, select the **Map to Host or Host Cluster** option as shown in Figure 6-20.

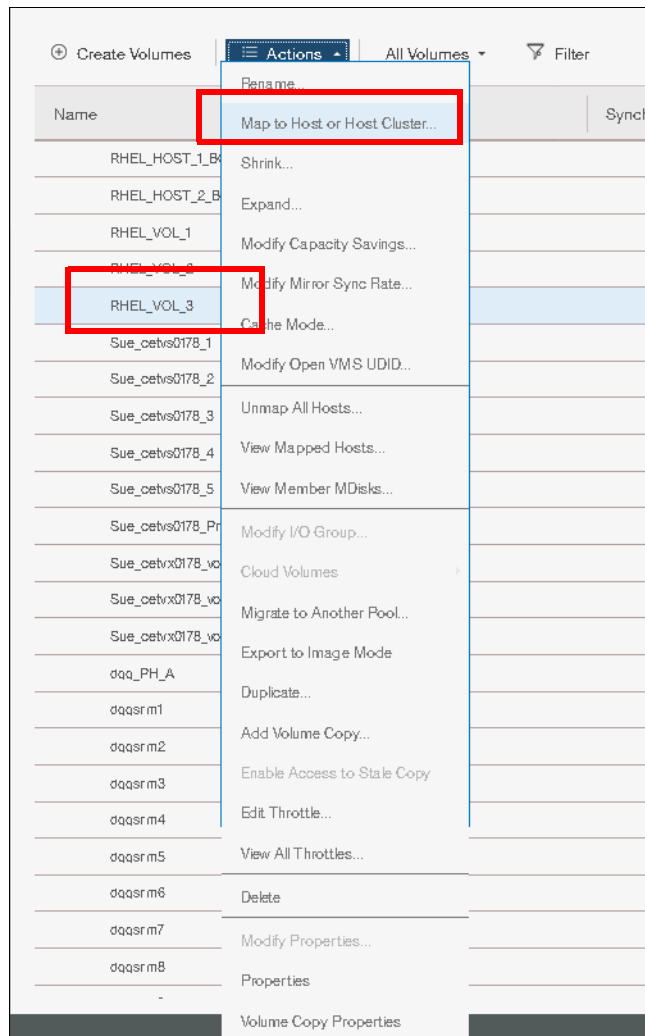


Figure 6-20 Map to Host

3. This action opens a **Create Mapping** window. In this window, indicate whether the volume needs to be mapped to a host or host cluster, select the desired host or host cluster, and whether you want system assigned SCSI ID or self assigned as shown in Figure 6-21 on page 292.

Create Mapping

Create Mappings to:

☒ Hosts

☐ Host Clusters

Select hosts to map to RHEL_VOL_3

Filter

Showing 6 hosts | Selecting 1 host

Name	Status	Host Type	Host Mappings
ARCSUN522021K3N	Offline	TPGS	Yes
ARCSUN5220524G1	Offline	TPGS	Yes
RHEL_HOST_1	Online	Generic	No
RHEL_HOST_2	Online	Generic	Yes
Sue_ARCX36502M...	Offline	Generic	Yes
dqq_63924_A	Online	Generic	Yes

Would you like the system to assign SCSI LUN IDs or manually assign these IDs?

☒ System Assign

☐ Self Assign

Cancel Back Next

Figure 6-21 Mapping a Volume to Host

- Click **Next**. A window will open up listing already mapped volumes to that host along with the new volume will be mapped, as shown in Figure 6-22 with blue rectangle.

Map Volumes to RHEL_HOST_2: Summary

The following volumes will be mapped to RHEL_HOST_2:

Name	SCSI ID	Caching I/O Group ID	New
RHEL_VOL_3	Auto	0	New
RHEL_VOL_1	0	0	
RHEL_VOL_2	1	0	

Cancel Back Map Volumes

Figure 6-22 Map volume to host

5. Click **Map Volumes** and the Modify Mappings window displays the command details, and then a Task completed message as shown in Figure 6-23.

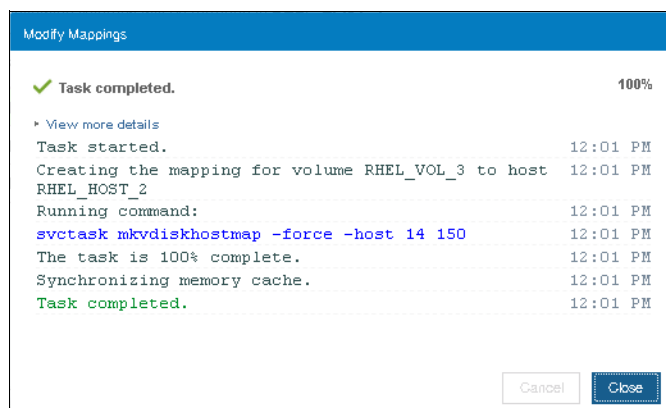


Figure 6-23 Volume mapped

6.5 Creating Custom volumes

The Create Volumes window enables Custom volume creation. It provides an alternative method of defining Capacity savings options, such as Thin-provisioning and Compression, but also expands on the base level default options for available Basic and Mirrored volumes. A Custom volume can be customized regarding Mirror sync rate, Cache mode, and Fast-Format.

The **Custom** volume creation operation consists of several options:

- Volume Location (Mandatory, defines the Pools to be used)
- Volume Details (Mandatory, defines the *Capacity savings* option such as Thin Provisioning or Compression or none)
- General (for changing default options for Cache mode and Formatting)
- Summary

Work through these options to customize your *Custom* volume as wanted, and then commit these changes by using **Create** as shown in Figure 6-24 on page 294.

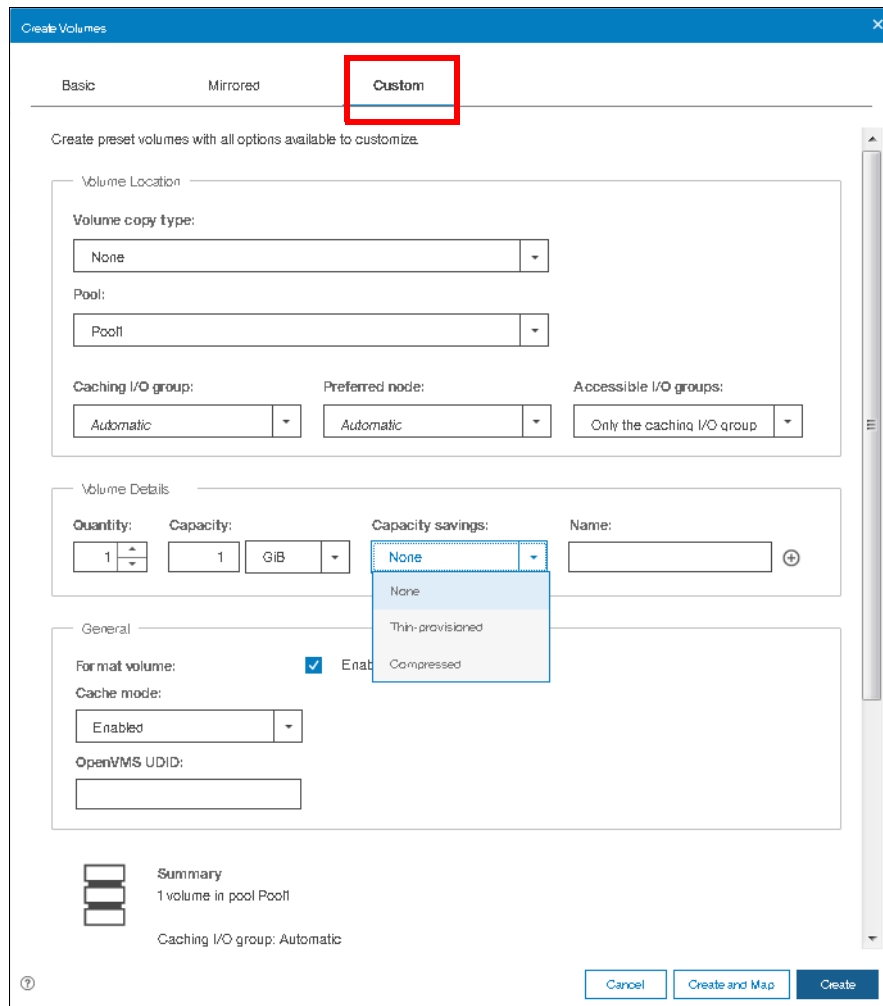


Figure 6-24 Customization submenus

6.5.1 Creating a custom thin-provisioned volume

A thin-provisioned volume can be defined and created by using the **Custom** option. Regarding application reads and writes, thin-provisioned volumes behave as though they were fully allocated. When creating a thin-provisioned volume, you can specify two capacities:

- ▶ The real physical capacity that is allocated to the volume from the storage pool. The real capacity determines the quantity of extents that are initially allocated to the volume.
- ▶ Its virtual capacity available to the host. The virtual capacity is the capacity of the volume that is reported to all other components (for example, FlashCopy, cache, and remote copy) and to the hosts.

To create a thin-provisioned volume, complete the following steps:

1. From the Create Volumes window, select the **Custom** option. In the **Volume Location** subsection define the pool in which the volume is created. Use the drop-down menu in the **Pool** option to choose the pool. All other options, such as **Volume copy type**, **Caching I/O group**, **Preferred node**, and **Accessible I/O groups**, can be left with their default options as shown in Figure 6-25 on page 295.

Create Volumes

Basic Mirrored **Custom**

Create preset volumes with all options available to customize.

Volume Location

Volume copy type:
None

Pool:
Click to select

Caching I/O group: Automatic Preferred node: Automatic Accessible I/O groups: Only the caching I/O group

Volume Details

Quantity: 1 Capacity: 1 GIB Capacity savings: Thin-provisioned Name:

Thin Provisioning

Real capacity: 2 % of Virtual capacity

Automatically expand: ☒ Enabled

Warning threshold: ☒ Enabled

Thin-Provisioned Grain Size: 256 KIB

General

Cancel Create and Map Create

Figure 6-25 Volume Location for thin-provisioned volume

- Next, in the **Volume Details** subsection you can input the **Quantity**, **Capacity** (virtual), **Capacity Savings** (choose **Thin-provisioned** from the drop-down menu), and **Name** of the volume being created as shown in Figure 6-26.

Volume Details

Quantity: 1 Capacity: 1 GIB Capacity savings: Thin-provisioned Name: ITSO_TP_VOL

Figure 6-26 Volume Details

3. Next, in the **Thin Provisioning** subsection enter the real and virtual capacity, expansion criteria, and grain size as shown in Figure 6-27.

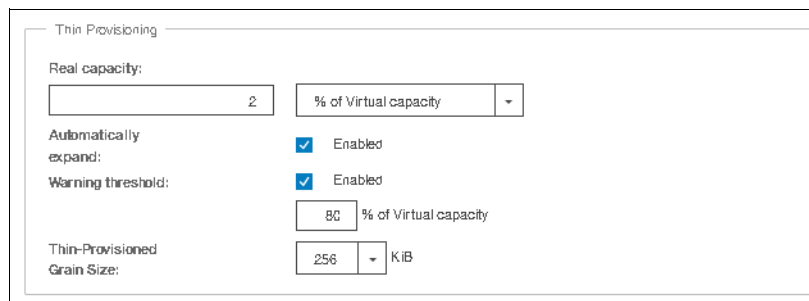
A screenshot of the 'Thin Provisioning' configuration window. It contains several settings: 'Real capacity' is set to '2' with a dropdown menu showing '% of Virtual capacity'; 'Automatically expand' is checked and set to 'Enabled'; 'Warning threshold' is checked and set to 'Enabled' with a dropdown showing '80 % of Virtual capacity'; and 'Thin-Provisioned Grain Size' is set to '256' with a dropdown showing 'KiB'.

Figure 6-27 Thin Provisioning

The Thin Provisioning options are as follows (defaults are displayed in parentheses):

- **Real capacity:** (2%). Specify the size of the real capacity space used during creation.
- **Automatically Expand:** (Enabled). This option enables the automatic expansion of real capacity, if more capacity is to be allocated.
- **Warning threshold:** (Enabled). Enter a threshold for receiving capacity alerts.
- **Grain Size:** (256 kibibytes (KiB)). Specify the grain size for real capacity. This option describes the size of the chunk of storage to be added to used capacity. For example, when the host writes 1 megabyte (MB) of new data, the capacity is increased by adding four chunks of 256 KiB each.

Important: If you do not use the **autoexpand** feature, the volume will go offline after reaching its real capacity.

The default grain size is 256 KiB. The optimum choice of grain size depends on volume use type. For more information, see the “Performance Problem When Using EasyTier With Thin Provisioned Volumes” topic:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_spaceefficientvdisks_3r7ayd.html

- ▶ If you are *not* going to use the thin-provisioned volume as a FlashCopy source or target volume, use 256 KiB to maximize performance.
- ▶ If you *are* going to use the thin-provisioned volume as a FlashCopy source or target volume, specify the same grain size for the volume and for the FlashCopy function.

4. Next, in the General sub-section enter the caching mode as Enabled, Read-only or Disabled as shown in Figure 6-28. Also enter unit device identifier (UDID) if this volume is going to be mapped to an OpenVMS host.

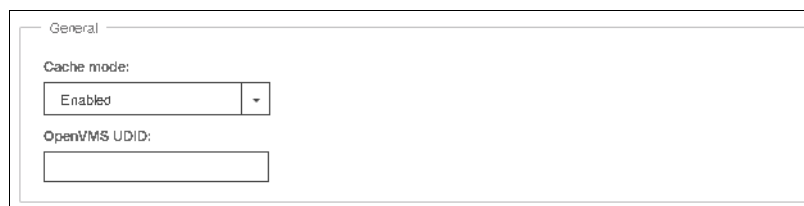
A screenshot of the 'General' configuration window. It contains two settings: 'Cache mode' is set to 'Enabled' with a dropdown menu; and 'OpenVMS UDID' is an empty text input field.

Figure 6-28 General details

5. Click **Create** to define the volume as shown in Figure 6-29.

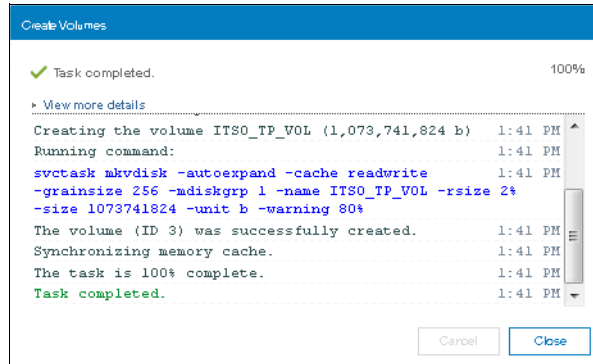


Figure 6-29 Task completed, the thin-provisioned volume is created

6.5.2 Creating Custom Compressed volumes

The configuration of compressed volumes is similar to thin-provisioned volumes. To create a Compressed volume, complete the following steps:

1. From the **Create Volumes** window, select the **Custom** option. In the **Volume Location** subsection define the pool in which the volume is created. Use the drop-down menu in the **Pool** option to choose the pool. All other options, such as **Volume copy type**, **Caching I/O group**, **Preferred node**, and **Accessible I/O groups**, can be left with their default options as shown in Figure 6-30 on page 298.

Create Volumes

Basic Mirrored **Custom**

Create preset volumes with all options available to customize.

Volume Location

Volume copy type:
None

Pool:
Pool1

Caching I/O group: Automatic Preferred node: Automatic Accessible I/O groups: Only the caching I/O group

Volume Details

Quantity: 1 Capacity: 1 GiB Capacity savings: Compressed Name: ITSO_COMPR_VOL

Compressed

Real capacity: 2 % of Virtual capacity

Automatically expand: ☒ Enabled

Warning threshold: ☒ Enabled 80 % of Virtual capacity

⚠ This volume is the first compressed volume on the system. The performance of existing volumes might be affected. For details on implementing Real-time Compression in pre-existing systems, [click here](#).

Cancel Create and Map Create

Figure 6-30 Defining a volume as compressed using the Capacity savings option

- Next, in the **Volume Details** subsection you can input the **Quantity**, **Capacity** (virtual), **Capacity Savings** (choose **Compressed** from the drop-down menu), and **Name** of the volume being created as shown in Figure 6-31.

Volume Details

Quantity: 1 Capacity: 1 GiB Capacity savings: Compressed Name: ITSO_COMPR_VOL

None
Thin-provisioned
Compressed

Compressed

Real capacity:

Figure 6-31 Volume Details

- Next, in the **Compressed** subsection enter the real either in terms of % of virtual capacity or in GiB, expansion criteria, warning threshold and desired % of the virtual capacity at which you should receive a warning as shown in Figure 6-32 on page 299.

Compressed

Real capacity: % of Virtual capacity

Automatically expand: ☒ Enabled

Warning threshold: ☒ Enabled

% of Virtual capacity

⚠ This volume is the first compressed volume on the system. The performance of existing volumes might be affected. For details on implementing Real-time Compression in pre-existing systems, [click here](#).

Figure 6-32 Compressed details

- Next, in the **General** sub-section enter the **caching mode** as Enabled, Read-only or Disabled as shown in Figure 6-33. Also enter unit device identifier (UDID) if this volume is going to be mapped to an OpenVMS host.

General

Cache mode:

OpenVMS UDID:

Figure 6-33 General details

- Click **Create** to define the volume as shown in Figure 6-34.

Create Volumes

✓ Task completed. 100%

► View more details

The task is 0% complete. 2:52 PM

Creating the volume ITSO_COMPR_VOL (1,073,741,824 b) 2:52 PM

Running command: 2:52 PM

`svctask mkvdisk -autoexpand -cache readwrite -compressed 2:52 PM`

`-mdiskgrp 1 -name ITSO_COMPR_VOL -rsize 2% -size 1073741824`

`-unit b -warning 80%`

The volume (ID 4) was successfully created. 2:52 PM

Synchronizing memory cache. 2:52 PM

The task is 100% complete. 2:52 PM

Task completed. 2:52 PM

Cancel Close

Figure 6-34 Compressed volume created

6.5.3 Custom Mirrored Volumes

The **Custom** option in the **Create Volumes** window is used to customize volume creation. Using this feature, the default options can be overridden and volume creation can be tailored to the specifics of the clients environment.

Modifying the Mirror sync rate

The **Mirror sync rate** can be changed from the default setting by using the **Custom** option, subsection **Volume Location**, of the Create Volumes window. This option sets the priority of copy synchronization progress, enabling a preferential rate to be set for more important volumes (Figure 6-35).

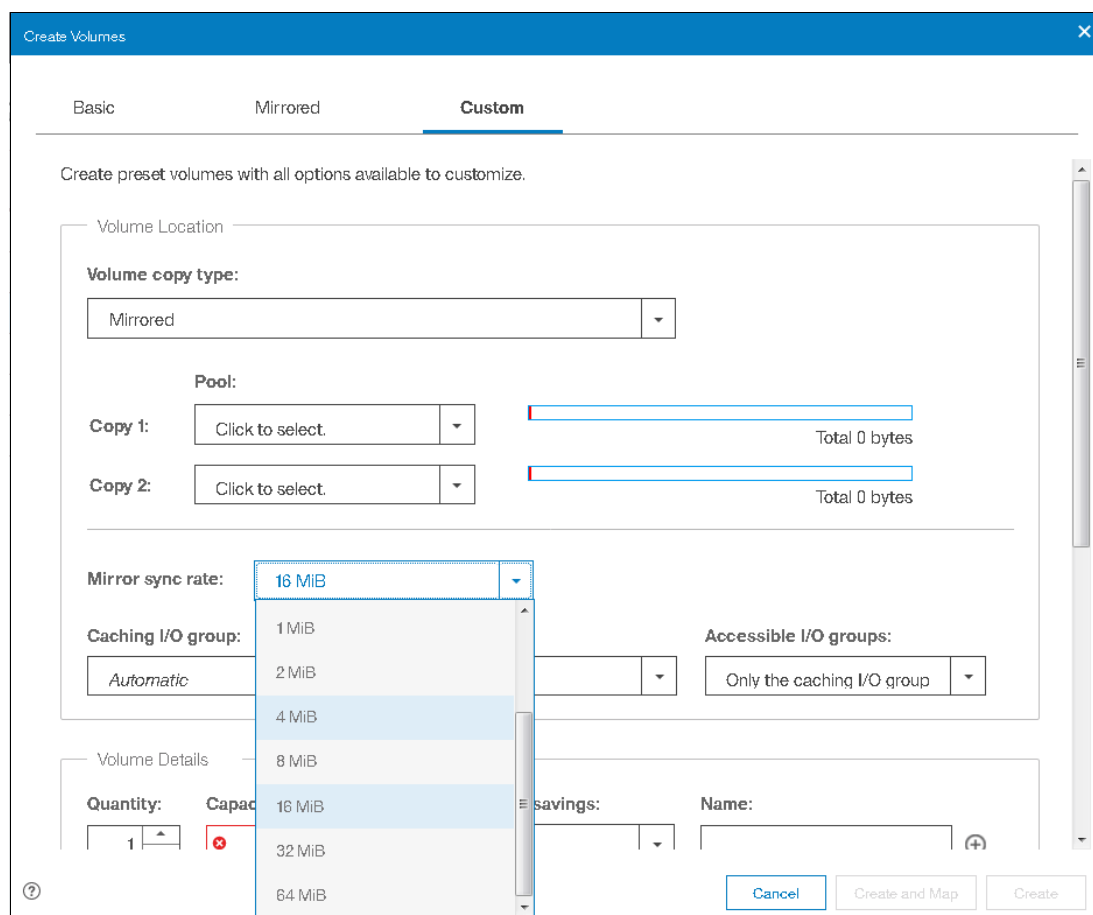


Figure 6-35 Customization of Mirrored sync rate

The progress of formatting and synchronization of a newly created Mirrored Volume can be checked from the **Running Tasks** menu. This menu reports the progress of all currently running tasks, including **Volume Format** and **Volume Synchronization** (Figure 6-36).




Figure 6-36 Progress of all currently running tasks

Creating a Custom Thin-provisioned Mirrored volume

The **Custom** option in the **Create Volumes** window is used to customize volume creation. Using this feature, the default options can be overridden and volume creation can be tailored to the specifics of the clients environment.

The **Mirror Sync rate** can be changed from the default setting under the **Volume Location** subsection of the **Create Volume** window. This option sets the priority of copy

synchronization progress, enabling a preferential rate to be set for more important mirrored volumes.

The summary shows you the capacity information and the allocated space. You can customize the thin-provision settings or the mirror synchronization rate. After you create the volume, the task completed window opens as shown in Figure 6-37.

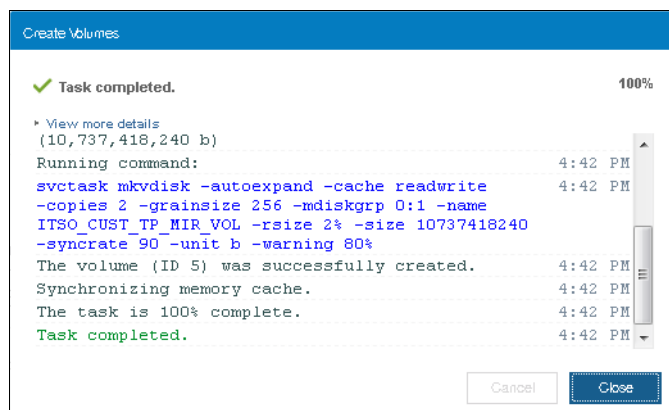


Figure 6-37 Task completed window

The initial synchronization of thin-mirrored volumes is fast when a small amount of real and virtual capacity is used.

6.6 HyperSwap and the mkvolume command

HyperSwap volume configuration is not possible until site awareness has been configured.

When the HyperSwap topology is configured, the GUI uses **mkvolume** command to create volumes instead of traditional **mkvdisk** command. This section describes the **mkvolume** command that is used in HyperSwap topology. The GUI continues to use **mkvdisk** when all other classes of volumes are created.

In this section, the new **mkvolume** command is described, and how the GUI uses this command, when HyperSwap topology has been configured, rather than the “traditional” **mkvdisk** command.

HyperSwap volumes are a new type of HA volumes that are supported by controller firmware. They are built off two existing technologies:

- ▶ Metro Mirror
- ▶ Volume Mirroring

These technologies have been combined in an active-active configuration deployed by using Change Volumes (as used in the Global Mirror with Change Volumes) to create a Single Volume (from a host perspective) in an HA form. The volume presented is a combination of four “traditional” volumes but is a single entity from a host (and administrative) perspective as shown in Figure 6-38 on page 302.

What does a HyperSwap volume look like?

- Composite object, comprised of:
 - 4 VDisks
 - 1 “active-active” relationship
 - 4 FlashCopy maps (for Change Volumes)
 - Additional access I/O group
- Remote Copy relationship and FlashCopy maps are automatically controlled
- Can be thick/thin/compressed
- Can be encrypted

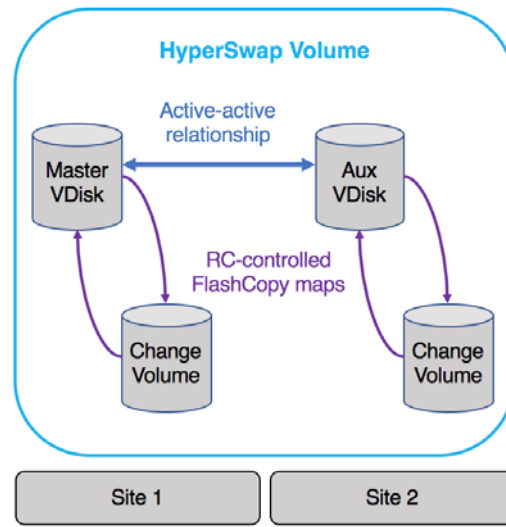


Figure 6-38 What makes up a HyperSwap Volume

The GUI simplifies the complexity of HyperSwap volume creation by only presenting the volume class of HyperSwap as a **Volume Creation** option after HyperSwap topology has been configured.

In the following example, HyperSwap topology has been configured and the **Volume Creation** window is being used to define a HyperSwap Volume as shown in Figure 6-39 on page 303.

The capacity and name characteristics are defined as for a Basic volume (highlighted in blue in the example) and the mirroring characteristics are defined by the Site parameters (highlighted in red).

Figure 6-39 HyperSwap Volume creation with Summary of actions

The drop-downs help during creation, and the Summary (lower left of the creation window) indicates the actions that are carried out when the **Create** option is selected. As shown in Figure 6-39, a single volume is created, with volume copies in `site1` and `site2`. This volume is in an active-active (Metro Mirror) relationship with extra resilience provided by two change volumes.

The command that is issued to create this volume is shown in Figure 6-40 on page 304, and can be summarized as follows:

```
svctask mkvolume -name <name_of_volume> -pool <X:Y> -size <Size_of_volume> -unit <units>
```

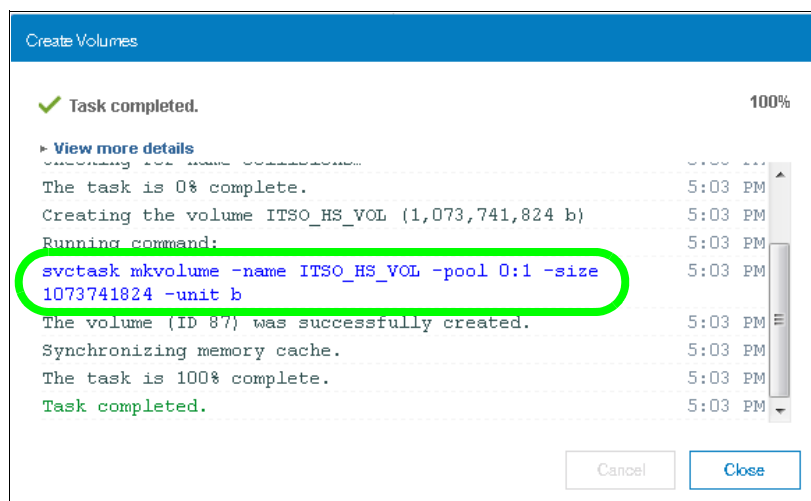


Figure 6-40 Example mkvolume command

6.6.1 Volume manipulation commands

Five CLI commands for administering volumes were released in V7.6.0. However, the GUI continues to use existing commands for all volume administration, except for HyperSwap volume creation (**mkvolume**) and deletion (**rmvolume**). The following CLI commands are available for administering volumes:

- ▶ **mkvolume**
- ▶ **mkimagevolume**
- ▶ **addvolumecopy**
- ▶ **rmvolumecopy**
- ▶ **rmvolume**

In addition, the **lsdisk** and GUI functionality is available. The **lsdisk** command now includes **volume_id**, **volume_name**, and **function** fields to easily identify the individual VDisk that make up a HyperSwap volume. These views are “rolled-up” in the GUI to provide views that reflect the client’s view of the HyperSwap volume and its site-dependent copies, as opposed to the “low-level” VDisks and VDisk-change-volumes.

As shown in the Figure 6-41, **Volumes** → **Volumes** shows the HyperSwap Volume ITSO_HS_VOL with an expanded view opened by using the twisty (V) to reveal two volume copies: ITSO_HS_VOL (site1) (Master VDisk) and ITSO_HS_VOL (site2) (Auxiliary VDisk). We do not show the VDisk-Change-Volumes.

Create Volumes Actions All Volumes Filter							Showing 31 volumes Selecting 1 volume (1.00 GB)
Name	State	Synchronized	Pool	UID	Host Mappings	Capacity	
ITSO_HS_VOL	Online		Multiple	6005076300888307880000000000009F	No	1.00 GB	
ITSO_HS_VOL (site1)	Online	Yes	Site1_P	6005076300888307880000000000009F	No	1.00 GB	
ITSO_HS_VOL (site2)	Online	Yes	Site2_P	600507630088830788000000000000A0	No	1.00 GB	

Figure 6-41 Hidden Change Volumes

Likewise, the status of the HyperSwap volume is reported at a “parent” level. If one of the copies is synchronized or not or offline, the HyperSwap volume reflects this state as shown in Figure 6-42 on page 305.

Name	State	Synchronized	Pool	Volume Group	UID	Host Mappings	Capacity
IT80_H8_VOL	Online (formatting)		ARCDS8KFCLE0		600507680C868245F000000000000009	No	1.00 GB
Copy 0*	Online (formatting)	Yes	ARCDS8KFCLE0		600507680C868245F000000000000009	No	1.00 GB
Copy 1	Online (formatting)	No	ARCDS8KX0700		600507680C868245F000000000000009	No	1.00 GB

Figure 6-42 Parent volume reflects state of copy volume

HyperSwap-related individual commands are described briefly here:

► **mkvolume**

Create an empty volume by using storage from existing storage pools. The type of volume created is determined by the system topology and the number of storage pools specified. Volume is always formatted (zeroed). This command can be used to create these items:

- Basic volume: Any topology
- Mirrored volume: Standard topology
- HyperSwap volume: HyperSwap topology

► **rmvolume**

Remove a volume. For a HyperSwap volume, this process includes deleting the active-active relationship and the change volumes.

The **-force** parameter with **rmvdisk** is replaced by individual override parameters, making it clearer to the user exactly what protection they are bypassing.

► **mkimagevolume**

Create an image mode volume. This command can be used to import a volume, preserving existing data. Implemented as a separate command to provide greater differentiation between the action of creating an empty volume and creating a volume by importing data on an existing MDisk.

► **addvolumecopy**

Add a copy to an existing volume. The new copy is always synchronized from the existing copy. For HyperSwap topology systems, this process creates a highly available volume. This command can be used to create the following volume types:

- Mirrored volume: Standard topology
- HyperSwap volume: HyperSwap topology

► **rmvolumecopy**

Remove a copy of a volume. Leaves the volume intact. Converts a Mirrored or HyperSwap volume into a basic volume. For a HyperSwap volume, this process includes deleting the active-active relationship and the change volumes.

This command enables a copy to be identified simply by its site.

The **-force** parameter with **rmvdiskcopy** is replaced by individual override parameters, making it clearer to the user exactly what protection they are bypassing.

See Lenovo Information Center for more details:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_hyperswap_configuration.html

6.7 Mapping Volumes to Host after volume creation

Newly created volume can be mapped to the host at creation time, or later. If the volume was not mapped to a host during creation, then to map it to a host, follow the steps in 6.7.1, “Mapping newly created volumes to the host using the wizard” on page 306.

6.7.1 Mapping newly created volumes to the host using the wizard

This section involves continue to map the volume that was created in 6.3, “Creating volumes using the Volume Creation” on page 285. We assume that you followed that procedure and are on the Volumes pane showing list of volumes as shown in Figure 6-43.

Create Volumes Actions All Volumes Filter									
Showing 132 volumes (selecting 1 volume (1/132))									
Name	State	Synchronized	Pool	Volume Group	UUID	Host Mappings	Capacity		
ITSO_NEW_VOL	Online		Pvc0		6055f760c84825800000000000000000	No	1.00 GiB		
RHEL_HOST_1_BOOT_VOL	Online		Pvc0		6055f760c84825800000000000000000	Yes	10.00 GiB		
RHEL_HOST_2_BOOT_VOL	Online		Pvc0		6055f760c84825800000000000000000	No	10.00 GiB		
RHEL_VOL_1	Online		Pvc0		6055f760c84825800000000000000000	Yes	1.00 GiB		
RHEL_VOL_2	Online		Pvc0		6055f760c84825800000000000000000	Yes	1.00 GiB		
RHEL_VOL_3	Online		Pvc0		6055f760c84825800000000000000000	Yes	1.00 GiB		
Sus_cetiv078_1	Online		Pvc0		6055f760c84825800000000000000000	Yes	50.00 GiB		
Sus_cetiv078_2	Online		Pvc0		6055f760c84825800000000000000000	Yes	50.00 GiB		
Sus_cetiv078_3	Online		Pvc0		6055f760c84825800000000000000000	Yes	50.00 GiB		
Sus_cetiv078_4	Online		Pvc0		6055f760c84825800000000000000000	Yes	50.00 GiB		
Sus_cetiv078_5	Online		Pvc0		6055f760c84825800000000000000000	Yes	50.00 GiB		
Sus_cetiv078_ProdPH	Online		Pvc0		6055f760c84825800000000000000000	Yes	10.00 GiB		
Sus_cetiv078_v06	Online		Pvc0		6055f760c84825800000000000000000	Yes	50.00 GiB		
Sus_cetiv078_v07	Online		Pvc0		6055f760c84825800000000000000000	Yes	50.00 GiB		
Sus_cetiv078_v08	Online		Pvc0		6055f760c84825800000000000000000	Yes	50.00 GiB		
oaz_PH_A	Online		Pvc0		6055f760c84825800000000000000000	Yes	10.00 GiB		
oqar m1	Online		xiang_arm		6055f760c84825800000000000000000	Yes	10.00 GiB		
oqar m2	Online		xiang_arm		6055f760c84825800000000000000000	Yes	10.00 GiB		
oqar m3	Online		xiang_arm		6055f760c84825800000000000000000	Yes	10.00 GiB		
oqar m4	Online		xiang_arm		6055f760c84825800000000000000000	Yes	10.00 GiB		
oqar m5	Online		xiang_arm		6055f760c84825800000000000000000	Yes	10.00 GiB		
oqar m6	Online		xiang_arm		6055f760c84825800000000000000000	Yes	10.00 GiB		
oqar m7	Online		xiang_arm		6055f760c84825800000000000000000	Yes	10.00 GiB		

Figure 6-43 Volume list

To map a volume, complete the following steps:

1. Right click on the volume name to be mapped and select **Map to Host or Host Cluster** menu option as shown in Figure 6-44 on page 307.

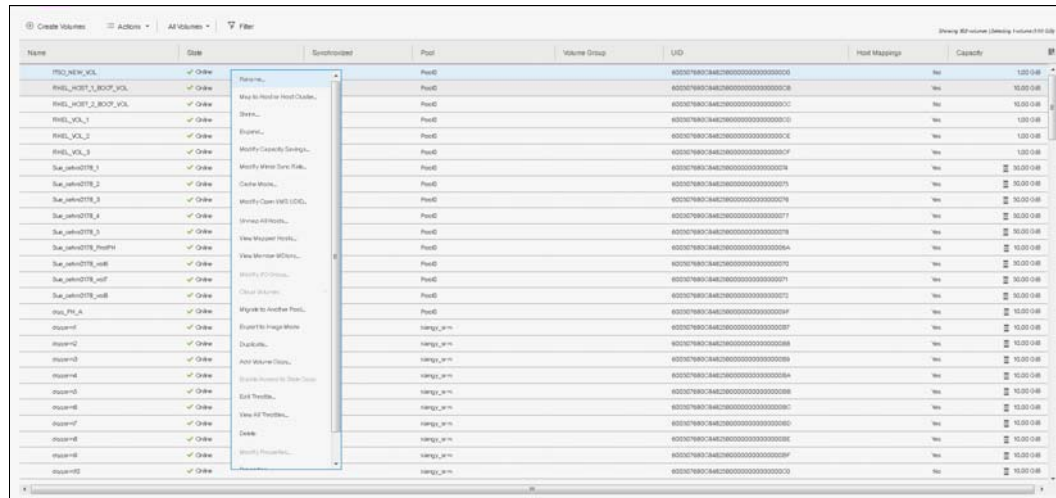


Figure 6-44 Select Map to Host or Host Cluster

2. Select a host or a host cluster to which the new volume should be attached as shown in Figure 6-45.




Figure 6-45 Select a host or a host cluster

Note: At this point you can let system assign a SCSI ID or choose it to assign it manually via Self Assign radio button. In this example, we chose to let system assign a SCSI ID.

3. Click **Next**. You will be shown a summary window with the volume to be mapped along with already existing volumes mapped to the host as shown in Figure 6-46 on page 308.

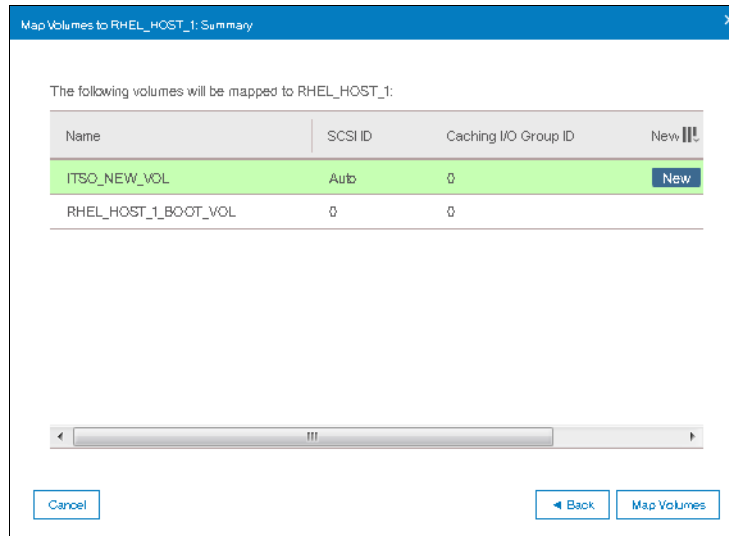


Figure 6-46 Map Volume summary

4. Click **Map Volumes**. A window indicating task of volume mapping completed will be shown as in Figure 6-47.

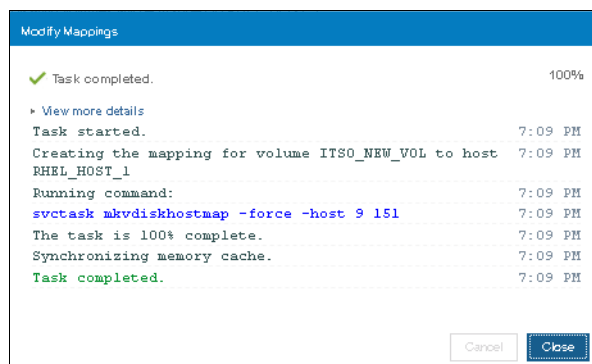


Figure 6-47 Volume mapping task completed

5. After the task completes, the wizard returns to the Volumes window.
6. You can verify the host mapping by clicking on **Hosts** from the main panel and selecting **Hosts** as shown in Figure 6-48 on page 309.




Figure 6-48 Hosts from main panel

7. Right click on the host to which the volume was mapped and select **Modify Volume Mappings** as shown in Figure 6-49

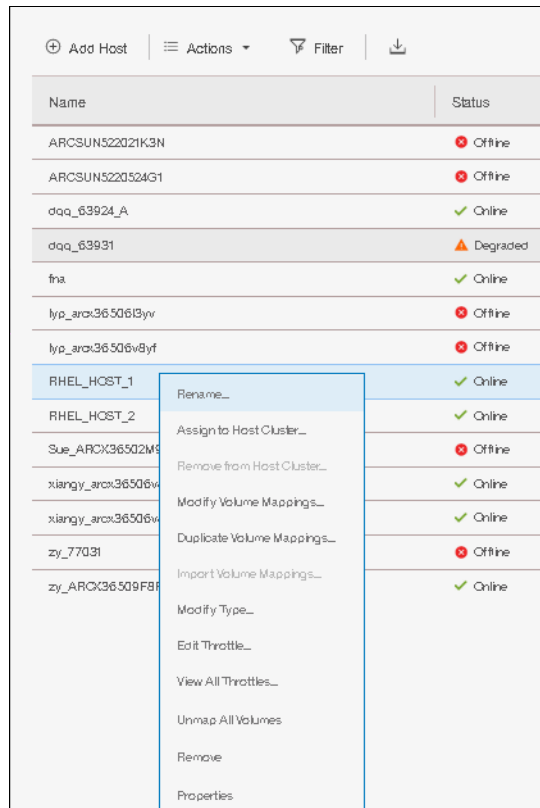


Figure 6-49 Modify host mappings

8. A window will pop up showing volumes currently mapped to the selected host as shown in Figure 6-50 on page 310.

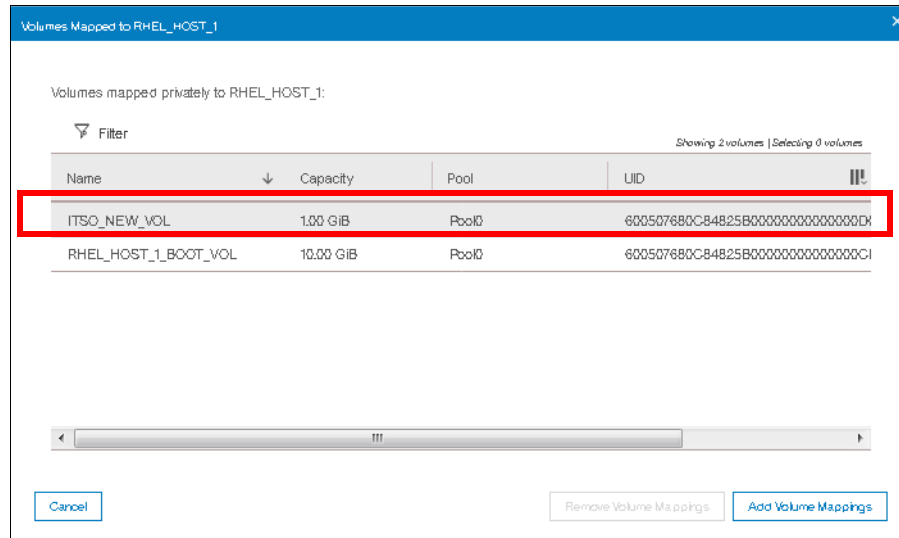


Figure 6-50 Volumes mapped to a host

The host is now able to access the volumes and store data on them. See 6.8, “Migrating a volume to another storage pool” on page 310 for information about discovering the volumes on the host and making additional host settings, if required.

Multiple volumes can also be created in preparation for discovering them later, and customize mappings.

6.8 Migrating a volume to another storage pool

The controller firmware provides online volume migration while applications are running. Storage pools are managed disk groups, as described in Chapter 4, “Storage pools” on page 139. With volume migration, data can be moved between storage pools, regardless of whether the pool is an internal pool, or a pool on another external storage system. This migration is done without the server and application knowing that it even occurred.

The migration process itself is a low priority process that does not affect the performance of the Lenovo storage V-series systems. However, it moves one extent after another to the new storage pool, so the performance of the volume is affected by the performance of the new storage pool after the migration process.

To migrate a volume to another storage pool, complete the following steps:

1. Click **Volumes** on the main panel and select **Volumes** sub-menu option as shown in Figure 6-51.




Figure 6-51 Volumes submenu option from main panel

2. Right click on the desired volume and select **Migrate to Another Pool** as shown in Figure 6-52.

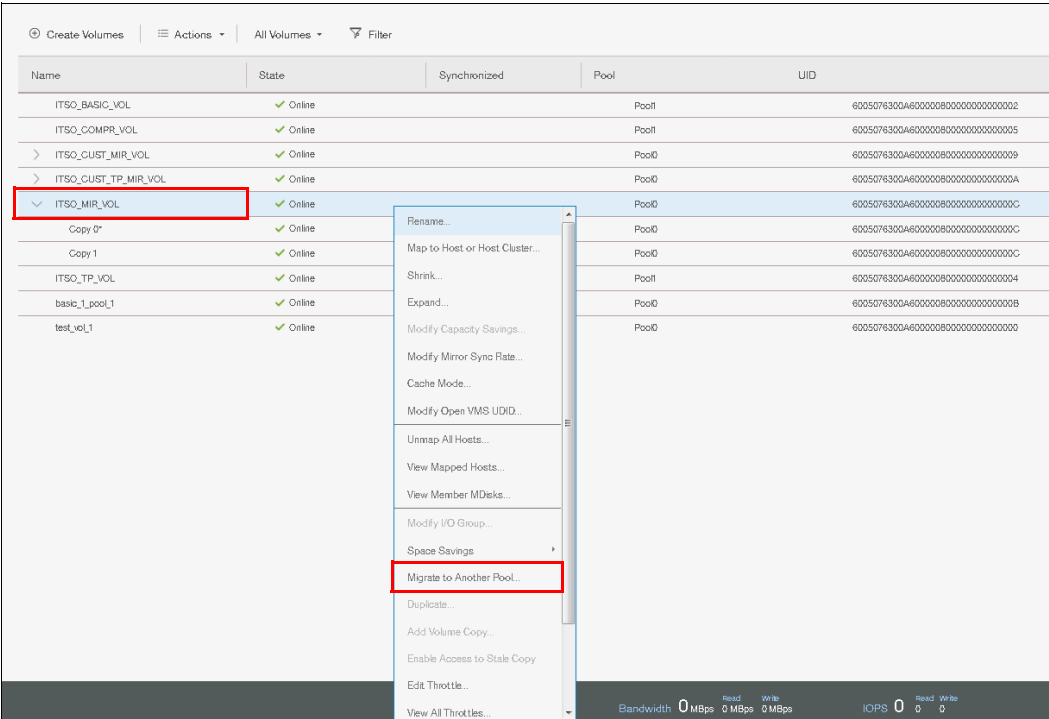


Figure 6-52 Selecting Migrate to Another Pool option

3. The Migrate Volume Copy window opens. If your volume consists of more than one copy, select the copy (from the menu shown in Figure 6-53) that you want to migrate to another storage pool. If the selected volume consists of one copy, this selection menu is not available.

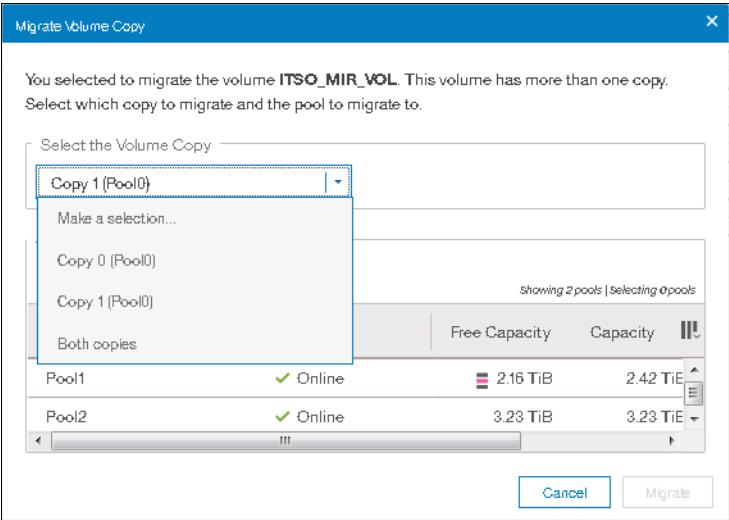


Figure 6-53 Migrate Volume Copy window: Select copy

4. Select the new target storage pool as shown in Figure 6-54 on page 312.




Figure 6-54 Selecting the new pool for volume migration

- Click **Migrate** and the volume copy migration starts as shown in Figure 6-55. Click **Close** to return to the Volumes pane.

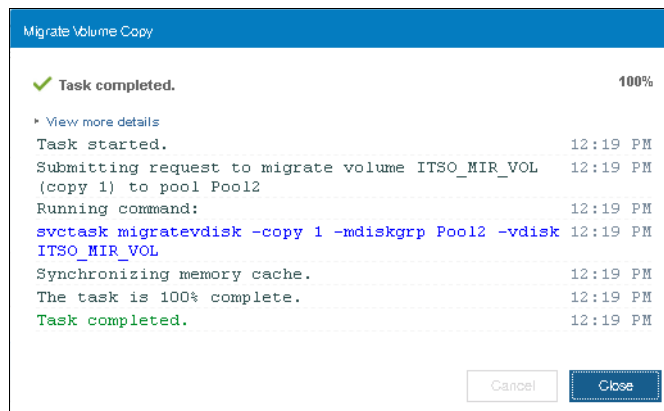


Figure 6-55 Migrate Volume Copy started

Depending on the size of the volume, the migration process takes some time, but you can monitor the status of the migration in the running tasks bar at the bottom of the GUI as shown in Figure 6-56.

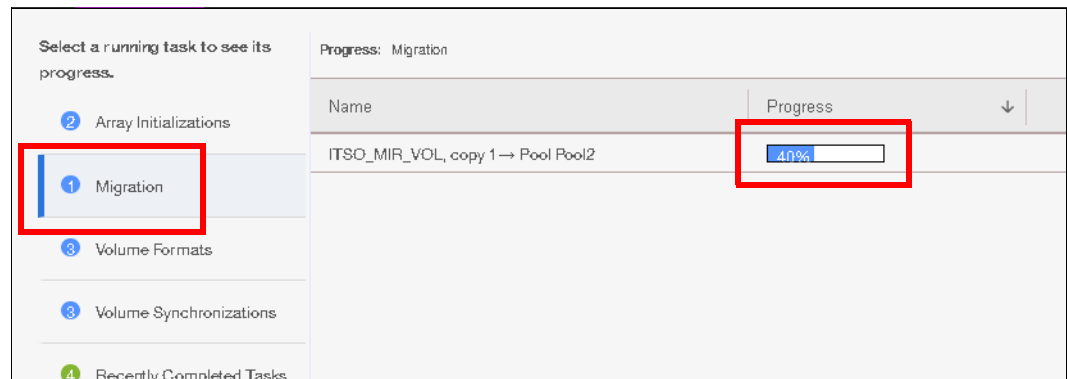
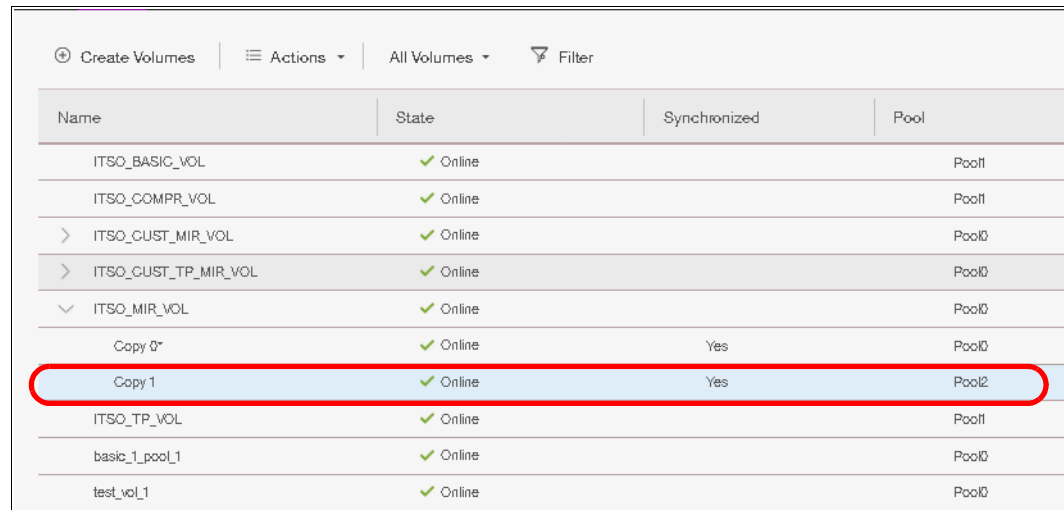


Figure 6-56 Migration progress

After the migration is completed, the volume is shown in the new storage pool. Figure 6-57 shows that it was moved from the Poo10 to the Poo12.



Name	State	Synchronized	Pool
ITSO_BASIC_VOL	✓ Online		Pool1
ITSO_COMPR_VOL	✓ Online		Pool1
> ITSO_CUST_MIR_VOL	✓ Online		Pool10
> ITSO_CUST_TP_MIR_VOL	✓ Online		Pool10
▽ ITSO_MIR_VOL	✓ Online		Pool10
Copy 0*	✓ Online	Yes	Pool10
Copy 1	✓ Online	Yes	Pool12
ITSO_TP_VOL	✓ Online		Pool1
baske_1_pool_1	✓ Online		Pool10
test_vol_1	✓ Online		Pool10

Figure 6-57 Migration complete

The volume copy has now been migrated without any host or application downtime to the new storage pool. It is also possible to migrate both volume copies to other pools online.

Another way to migrate volumes to another pool is by performing the migration using the volume copies, as described in 6.9, “Migrating volumes using the volume copy feature” on page 313.

Note: Migrating a volume between storage pools with different extent sizes is *not* supported. If you need to migrate a volume to a storage pool with a different extent size, use volume copy features instead.

6.9 Migrating volumes using the volume copy feature

The controller firmware supports creation, synchronizing, splitting, and deleting volume copies. A combination of these tasks can be used to migrate volumes to other storage pools. The easiest way to migrate volume copies is to use the migration feature described in 6.8, “Migrating a volume to another storage pool” on page 310. If you use this feature, one extent after another is migrated to the new storage pool. However, using volume copies provides another possibility to migrate volumes.

To migrate a volume, complete the following steps:

1. Select the volume and select **Add Volume Copy** operation as shown in Figure 6-58 on page 314.




Figure 6-58 Adding the volume copy to another pool

2. Select the desired pool into which a new copy will be created as shown in Figure 6-59




Figure 6-59 Add Volume Copy

3. You will get a confirmation window as shown in Figure 6-60.

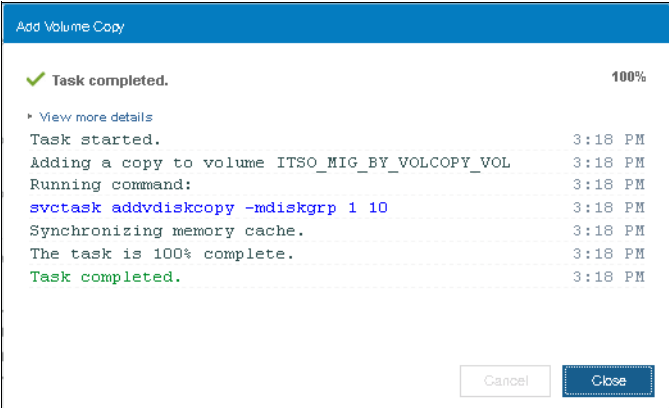


Figure 6-60 Volume Copy task completed

4. Wait until the copies are synchronized. Then, change the role of the copies and make the new copy of the primary copy as shown in Figure 6-61.




Figure 6-61 Making the new copy in a different storage pool as primary

5. A task completion window will be shown as in Figure 6-62 on page 316.

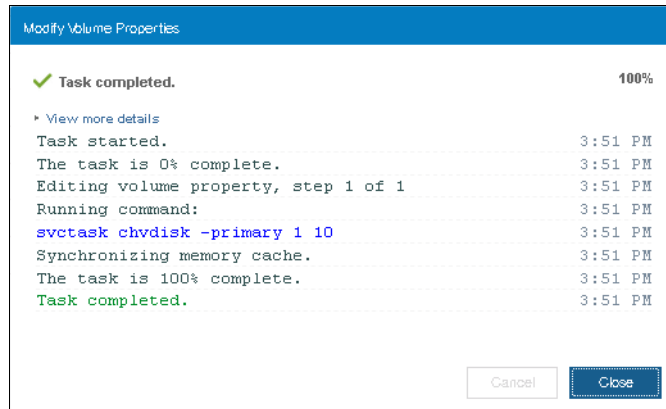


Figure 6-62 Primary copy change task completed

6. Split or delete the old copy from the volume as shown in Figure 6-63.




Figure 6-63 Deleting the older copy

7. Ensure that the new copy is in the target storage pool as shown in Figure 6-64 on page 317.




Figure 6-64 Verifying the new copy in the target storage pool

This migration process requires more user interaction, but it offers some benefits, for example, if you migrate a volume from a tier 1 storage pool to a lower performance tier 2 storage pool. In step 1, you create the copy on the tier 2 pool. All reads are still performed in the tier 1 pool to the primary copy. After the synchronization, all writes are destaged to both pools, but the reads are still only done from the primary copy.

Now you can switch the role of the copies online (step 3), and test the performance of the new pool. If you are done testing your lower performance pool, you can split or delete the old copy in tier 1, or switch back to tier 1 in seconds, in case tier 2 pool did not meet your performance requirements.

6.10 I/O throttling

You can set a limit on the number of I/O operations that are accepted for a volume. The limit is set in terms of I/O operations per second (IOPS) or MBps. By default, no I/O throttling rate is set when a volume is created. I/O throttling is also referred to as I/O governing.

Base the choice between I/O and MB as the I/O governing throttle on the disk access profile of the application. Database applications generally issue large amounts of I/O, but they transfer only a relatively small amount of data. In this case, setting an I/O governing throttle that is based on MBps does not achieve much. It is better to use an IOPS as a second throttle.

At the other extreme, a streaming video application generally issues a small amount of I/O, but it transfers large amounts of data. In contrast to the database example, setting an I/O governing throttle that is based on IOPS does not achieve much, so it is better to use an MBps throttle.

An I/O governing rate of 0 does not mean that zero IOPS (or MBps) can be achieved. It means that no throttle is set.

Note:

- ▶ I/O governing does not affect FlashCopy and data migration I/O rates.
- ▶ I/O governing on a Metro Mirror and Global Mirror secondary volume does not affect the rate of data copy from the primary volume.

6.10.1 Define throttle on a volume

To set a throttle on a volume, follow the steps as listed.

1. From the Volumes window, select the desired volume to throttle as shown in Figure 6-65 on page 319.

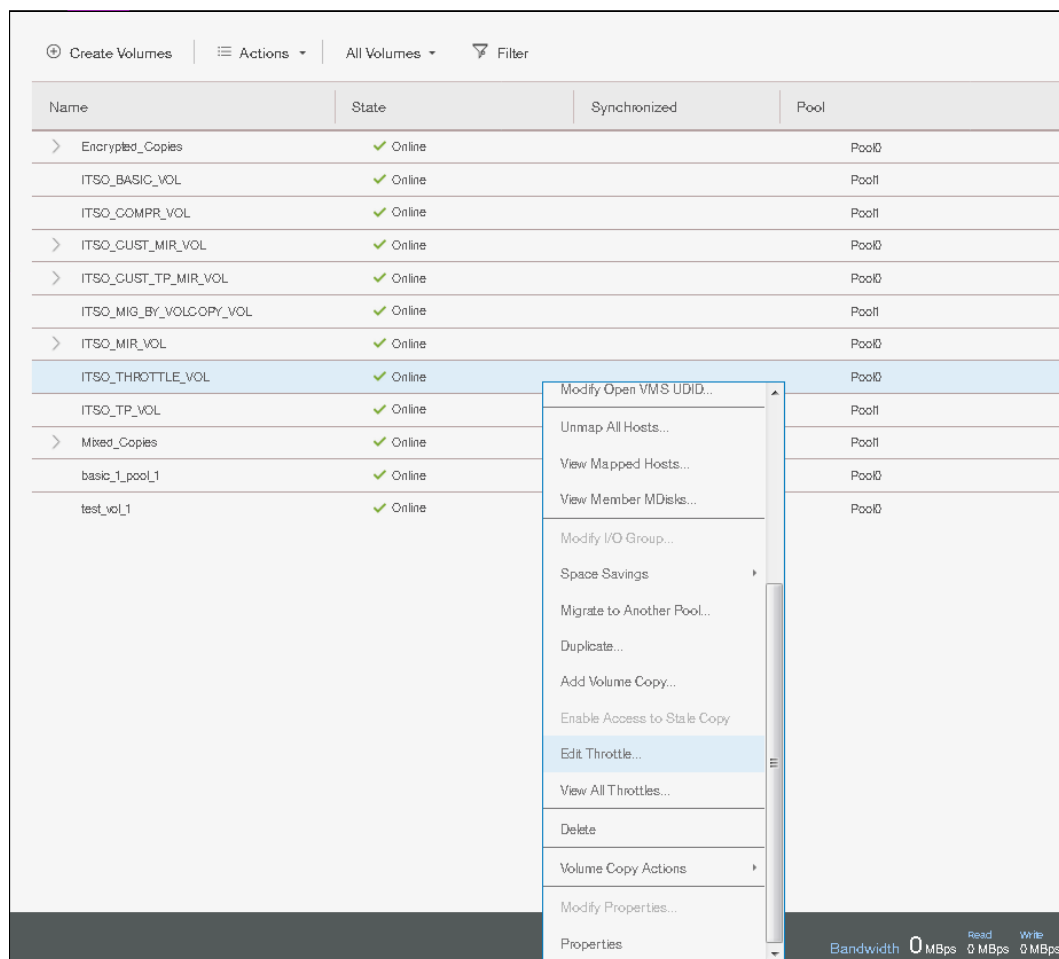


Figure 6-65 Edit Throttle

2. A window will come up where you can set the throttle either in terms of IOPS or bandwidth (MBps) or both. In our example, we set the throttle on IOPS as shown in Figure 6-66.

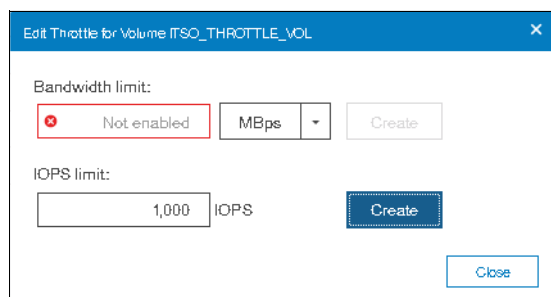


Figure 6-66 IOPS throttle on a volume

3. Click **Create** and you will get a task completed window as shown in Figure 6-67 on page 320

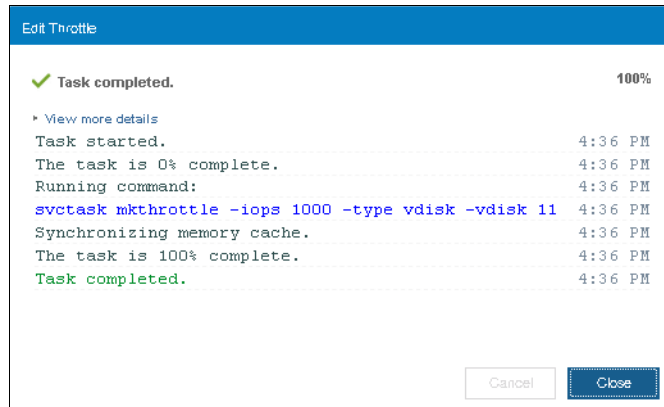


Figure 6-67 Volume throttle completed

6.10.2 Remove a throttle from a volume

To set a throttle on a volume, follow the steps as listed

1. From the Volumes window, select the desired volume to throttle as shown in Figure 6-65 on page 319.

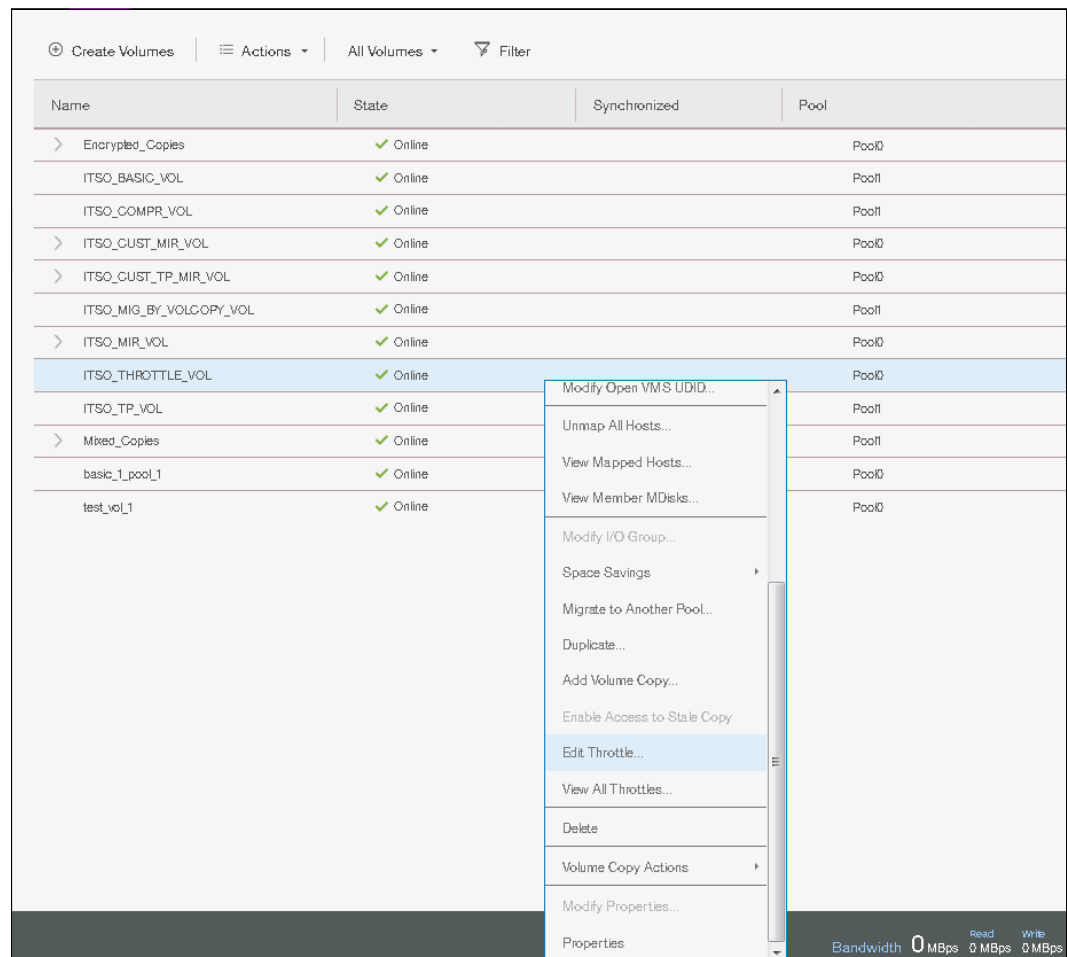


Figure 6-68 Edit Throttle

2. A window will come up where you can remove the throttle that was defined. In our example, we remove the throttle on IOPS by clicking on **Remove** as shown in Figure 6-69.

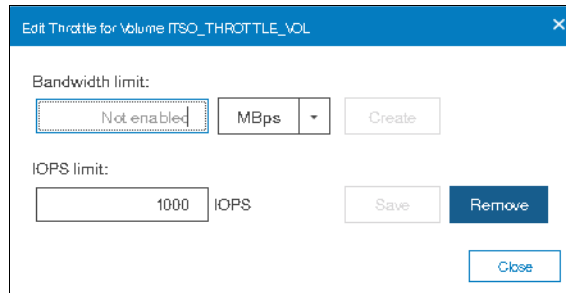


Figure 6-69 Remove throttle

3. A task completion window will be shown as in Figure 6-70

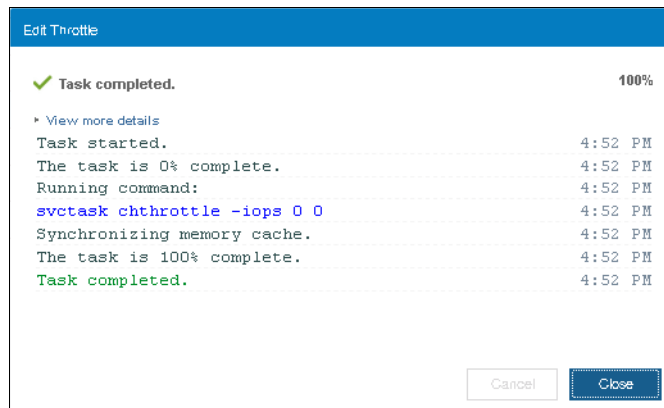


Figure 6-70 Volume throttle removal task completed

Storage migration

This chapter provides a detailed walkthrough of the storage migration wizard. The storage migration wizard migrates data from external storage systems to the internal capacity of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

The data migration from other storage systems to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 consolidate storage and enable the benefits of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 functionalities across all the volumes, such as the intuitive GUI, internal virtualization, thin provisioning and FlashCopy.

This chapter includes the following topics:

- ▶ 7.1, “Storage migration wizard overview” on page 324
- ▶ 7.2, “Interoperation and compatibility” on page 324
- ▶ 7.3, “Storage migration wizard” on page 325

7.1 Storage migration wizard overview

The storage migration wizard uses the Volume Mirroring functionality to allow reads and writes during the migration, eliminating disruption and downtime. After the end of the migration, the external storage can be retired. The storage migration wizard can also be used to migrate data from other Lenovo Storage systems, such as the IBM Storwize V5000 for Lenovo or IBM Storwize V3700 for Lenovo.

There are multiple reasons to use the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 migration features:

- ▶ To redistribute workload within a clustered system across the disk subsystem
- ▶ To move workload onto newly installed storage
- ▶ To move workload off old or failing storage, ahead of decommissioning it
- ▶ To migrate data from an older disk subsystem to Lenovo Storage V3700 V2, V3700 V2 XP, and V5030
- ▶ To migrate data from one disk subsystem to another disk subsystem

Specifically, this chapter provides information about the following topics:

- ▶ Interoperation and compatibility
- ▶ Storage migration wizard

Command-line interface (CLI): For more information about the command-line interface setup, see Appendix A, “CLI setup and SAN Boot” on page 761.

Manually migrating data: For more information about migrating data manually, see Chapter 11, “External storage virtualization” on page 607.

7.2 Interoperation and compatibility

Interoperation is an important consideration when a new storage system is deployed into an environment that already has an existing storage infrastructure. This section describes how to check that the existing storage environment, the storage system and the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are ready for the data migration process.

To ensure interoperation and compatibility across all the elements that connect to the Storage Area Network (SAN) fabric, check the proposed configuration with the Lenovo interoperability matrix. The interoperability matrix can confirm whether a solution is supported and provide recommendations for hardware and software levels. Interoperability matrix validates the components within a single storage solution. To confirm the interoperation between multiple storage solutions, it is necessary to request separate validation for each of them.

For more information about interoperability, refer to this web pages:

https://download.lenovo.com/storage/lenovo_storage_v5030_8_1_x.xls

https://download.lenovo.com/storage/lenovo_storage_v3700v2_8_1_x.xls

7.3 Storage migration wizard

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage migration wizard simplifies the migration task. The wizard features intuitive panels that guides users through the entire process.

7.3.1 External virtualization capability

All of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models can migrate data from external storage controllers, including migrating from any other previous IBM Storwize for Lenovo systems generation. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 use the functionality that is provided by its external virtualization capability to perform the migration. This capability places external Fibre Channel (FC)-connected logical units (LUs) under the control of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. After the volumes are virtualized, hosts continue to access them through the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, which act as a proxy.

The difference between Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models is that the Lenovo Storage V3700 V2 and Lenovo Storage V3700 V2 XP models can perform data migration only, external storage controllers cannot be virtualized on them. Lenovo Storage V5030 can be used to migrate data and externally virtualize storage from external storage controllers. For more information about external virtualization, see Chapter 11, “External storage virtualization” on page 607.

7.3.2 Model and adapter card considerations

Storage migration can use either the Fibre Channel SAN or direct-attach serial-attached Small Computer System Interface (SCSI) (SAS) connections to perform the data transfer.

Fibre Channel migration to a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 require the purchase of a pair of the optional 16 Gb FC adapter cards.

SAS migration on the Lenovo Storage V3700 V2 and Lenovo Storage V5030 systems require the purchase of a pair of optional SAS adapter cards. SAS migration on Lenovo Storage V3700 V2 can be performed without an adapter card by using the onboard SAS host attachment ports.

Table 7-1 summarizes the requirements for each model.

Table 7-1 Comparison of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models for storage migration

	Lenovo Storage V3700 V2	Lenovo Storage V3700 V2 XP	Lenovo Storage V5030
External virtualization	Not supported	Not supported	Licensed feature
FC SAN migration	With 16 G FC cards	With 16 G FC cards	With 16 G FC cards
SAS device adapter (DA) migration	With 12 G SAS cards	Yes (onboard ports)	With 12 G SAS cards

7.3.3 Overview of the storage migration wizard

An overview of the storage migration wizard process includes the following considerations:

1. Typically, storage systems segregate storage into many Small Computer System Interface (SCSI) LUs that are presented to hosts through a Fibre Channel SAN. Storage can also be presented through direct SAS attachment to a host. In general, the steps to migrate either one of these storage systems are the same.
2. Input/output (I/O) to the Logical Unit Numbers (LUN) must be stopped, changes must be made to the mapping of the storage system LUs and to the SAN fabric zoning or SAS connectivity, so that the original LUs are presented directly to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and not to the hosts. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 discover the external LUs as unmanaged (not a member of any storage pools) managed disks (MDisks).
3. The unmanaged MDisks are then imported to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 as image mode MDisks and placed in a storage pool. This storage pool is now a logical container for the externally attached LUs.
4. Each volume has a one-to-one mapping with an image mode MDisk. From a data perspective, the image mode volume represents the SAN-attached LUs exactly as they were before the import operation. The image mode volumes are on the same physical drives of the existing storage system and the data remains unchanged. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 present active images of the SAN-attached LUs and it acts as a proxy.
5. You need to remove the existing storage system multipath device driver from the hosts. Then, configure the hosts for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 attachment. If you are migrating over Fibre Channel, further zoning changes are made for host-to-Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 SAN connections. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 hosts are defined with worldwide port names (WWPNs) and the volumes are mapped to the hosts. After the volumes are mapped, the hosts discover the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 volumes through a host rescan or reboot operation.
6. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 volume mirror operations are then initiated. The image mode volumes are mirrored to generic volumes. The mirrors are online migration tasks, which means that a defined host can still access and use the volumes during the mirror synchronization process.
7. After the mirror operations are complete, the volume mirror relationships and the image mode volumes are removed. The other storage system LUs are migrated and the now redundant existing storage can be retired or reused elsewhere.

Important: If you are migrating volumes from another IBM Storwize for Lenovo products, be aware that the target Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems need to be configured at the replication layer for the source to discover the target system as a host. The default layer setting for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 is storage.

Ensure that the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems are configured as a replication layer system. Enter the following command:

chsystem -layer replication

If you do not enter this command, you cannot add the target system as a host on the source storage system or see source volumes on the target system.

To change the source IBM Storwize for Lenovo system to the storage layer, enter the following command:

chsystem -layer storage

For more information about layers and how to change them, see Chapter 10, “Copy services” on page 451.

7.3.4 Storage migration wizard tasks

The storage migration wizard is designed for an easy migration of data from other storage systems to the internal capacity of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

This section describes the following storage migration wizard tasks:

- ▶ Avoiding data loss
- ▶ Verifying prerequisites for Fibre Channel connections
- ▶ Verifying prerequisites for SAS connections
- ▶ Verifying prerequisites for iSCSI connections
- ▶ Identifying restrictions and prerequisites for the wizard
- ▶ Preparing the environment for migration
- ▶ Mapping storage
- ▶ Migrating MDisk
- ▶ Configuring hosts
- ▶ Mapping volumes to hosts
- ▶ Selecting a storage pool
- ▶ Finishing the storage migration wizard
- ▶ Finalizing migrated volumes

Avoiding data loss

It is prudent to avoid any potential data loss by creating a backup of all of the data that is stored on the hosts, the existing storage systems and the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 before you use the storage migration wizard.

Verifying prerequisites for Fibre Channel connections

Cable this system into the SAN of the external storage that you want to migrate. Ensure that your system is cabled and zoned into the same storage area network (SAN) as the external storage system that you are migrating. If you are using Fibre Channel, connect the Fibre Channel cables to the Fibre Channel ports in both canisters of your system, and then to the

Fibre Channel network. If you are using Fibre Channel over Ethernet (FCoE), connect Ethernet cables to the 10 Gbps Ethernet ports.

Ensure that all systems are running a software level that enables them to recognize the other nodes in the cluster. Also, ensure that the systems use Fibre Channel adapters at the same speed. To avoid performance bottlenecks, do not use a combination of 8 Gbps and 16 Gbps links.

Examples of how to connect an IBM Storwize V5000 for Lenovo to a Lenovo Storage V5030 system are shown in Figure 7-1 and Figure 7-2 on page 329.




Figure 7-1 Direct Fibre Channel connections between systems




Figure 7-2 Fibre Channel connections using switches between the systems

Verifying prerequisites for SAS connections

For migrating from an IBM Storwize V3500 for Lenovo, IBM Storwize V3700 for Lenovo, or IBM Storwize V5000 for Lenovo, ensure that all systems are running a software level that can support SAS migration.

Cable the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 directly to the external storage system that you want to migrate. Depending on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 models, the cabling differs slightly. The Lenovo Storage V3700 V2 and Lenovo Storage V5030 need four SAS cables (two SAS cables per node canister) that are connected to the optional SAS card. The Lenovo Storage V3700 V2 needs four SAS cables (two SAS cables per node canister) that are connected to SAS port 2 and SAS port 3.

The IBM Storwize V3500 for Lenovo or IBM Storwize V3700 for Lenovo source systems require two cables per node canister. Each canister must be connected to each Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 canister. On the IBM Storwize V3500 or V3700 for Lenovo, you can use SAS ports 1, 2, or 3. Do not use SAS port 4.

Examples of how to connect a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 to the IBM Storwize V3500/V3700 for Lenovo are shown in Figure 7-3 on page 330, Figure 7-4 on page 330, and Figure 7-5 on page 331.




Figure 7-3 Connecting SAS cables from an IBM Storwize V3500 or V3700 for Lenovo to a Lenovo Storage V3700 V2 system




Figure 7-4 Connecting SAS cables from an IBM Storwize V3500 or V3700 for Lenovo to a Lenovo Storage V3700 V2 XP system




Figure 7-5 Connecting SAS cables from an IBM Storwize V3500 or V3700 for Lenovo to a Lenovo Storage V5030 system

IBM Storwize V5000 for Lenovo source systems require two cables per node canister. Each canister must be connected to each Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 canister. On the IBM Storwize V5000 for Lenovo, you must use SAS port 1 or 2. Do not use SAS port 3 or 4.

Examples of how to connect a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 to an IBM Storwize V5000 for Lenovo are shown in Figure 7-6, Figure 7-7 on page 332, and Figure 7-8 on page 332.




Figure 7-6 Connecting SAS cables from an IBM Storwize V5000 for Lenovo system to a Lenovo Storage V3700 V2 system




Figure 7-7 Connecting SAS cables from an IBM Storwize V5000 for Lenovo system to a Storwize V3700 V2 XP system




Figure 7-8 Connecting SAS cables from an IBM Storwize V5000 for Lenovo system to a Lenovo Storage V5030 system

Verifying prerequisites for iSCSI connections

On Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, you can use iSCSI connections to migrate data from different systems (Lenovo Storage V3700 V2, V3700 V2 XP and V5030) and to virtualize external storage systems (Lenovo Storage V5030).

Migration considerations and configurations can vary depending on the type of system to be migrated or virtualized. You can use an iSCSI attachment to migrate data from an IBM Storwize for Lenovo and Lenovo Storage V series system to Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 do not support iSCSI connections to migrate data from IBM Storwize V3500 for Lenovo.

You can use any available Ethernet port to establish iSCSI connectivity between the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems and the backend storage controller.

Note: If you are using onboard Ethernet ports on a Lenovo Storage V3700 V2 or Lenovo Storage V3700 V2 XP system, ensure that the onboard Ethernet port 2 on the system is not configured to be used as the technician port.

To avoid performance bottlenecks, the iSCSI initiator and target systems must use Ethernet ports at the same speed. Do not use a combination of Ethernet links that run at different speeds.

For full redundancy and increased throughput, use two or more Ethernet switches. Similarly numbered Ethernet ports on each node of each system must be connected to the same switch. They must also be configured on the same subnet or VLAN.

Figure 7-9 shows iSCSI connections between a Lenovo Storage V5030 system (iSCSI initiator) and an IBM Storwize V3700 for Lenovo system (iSCSI target).




Figure 7-9 iSCSI connections between a Lenovo Storage V5030 system (iSCSI initiator) and an IBM Storwize V3700 for Lenovo system (iSCSI target)

Accessing the storage migration wizard

Select **System Migration** in the Pools menu (Figure 7-10) to open the System Migration panel.

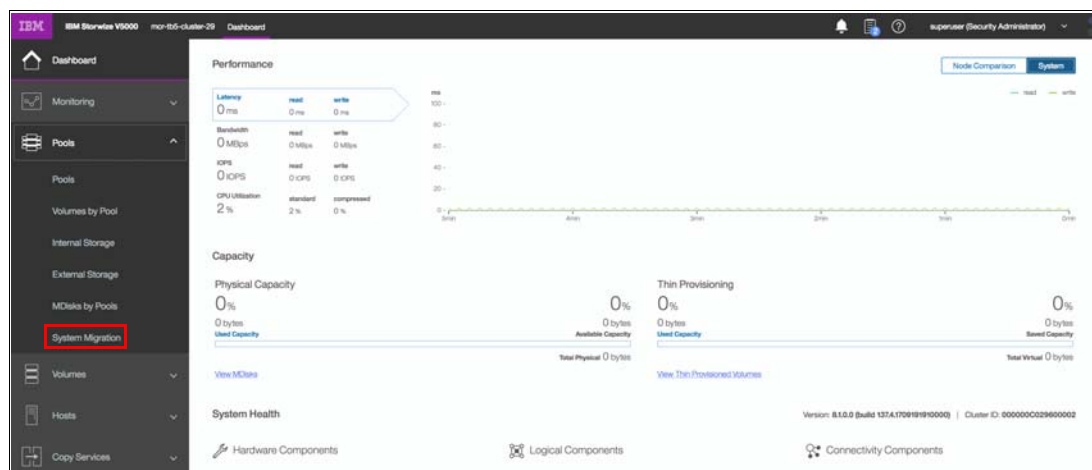


Figure 7-10 Pools menu

The System Migration panel provides access to the storage migration wizard and displays the migration progress information. Click **Start New Migration** to begin the storage migration wizard. Figure 7-11 shows the System Migration panel.




Figure 7-11 System Migration panel

Important:

- ▶ You might receive a warning message as shown in Figure 7-12 that indicates that no externally attached storage controllers were found if you did not configure your zoning correctly (or if the layer was incorrectly set if another IBM Storwize for Lenovo system is attached). Click **Close** and correct the problem before you start the migration wizard again.
- ▶ The subsequent panels in the migration wizard, as shown in Figure 7-14 on page 337, direct you to remove the host zoning to the external storage and create zones between the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and the external storage. However, these steps must be performed *before* you start the wizard. For the list of instructions to complete before you start the data migration wizard, see “Preparing the environment for migration” on page 337 and “Mapping storage” on page 337.



Figure 7-12 Error message that is displayed when no external storage is found

Identifying restrictions and prerequisites for the wizard

This panel of the storage migration wizard describes the restrictions and prerequisites for the wizard, as shown in Figure 7-13 on page 336.

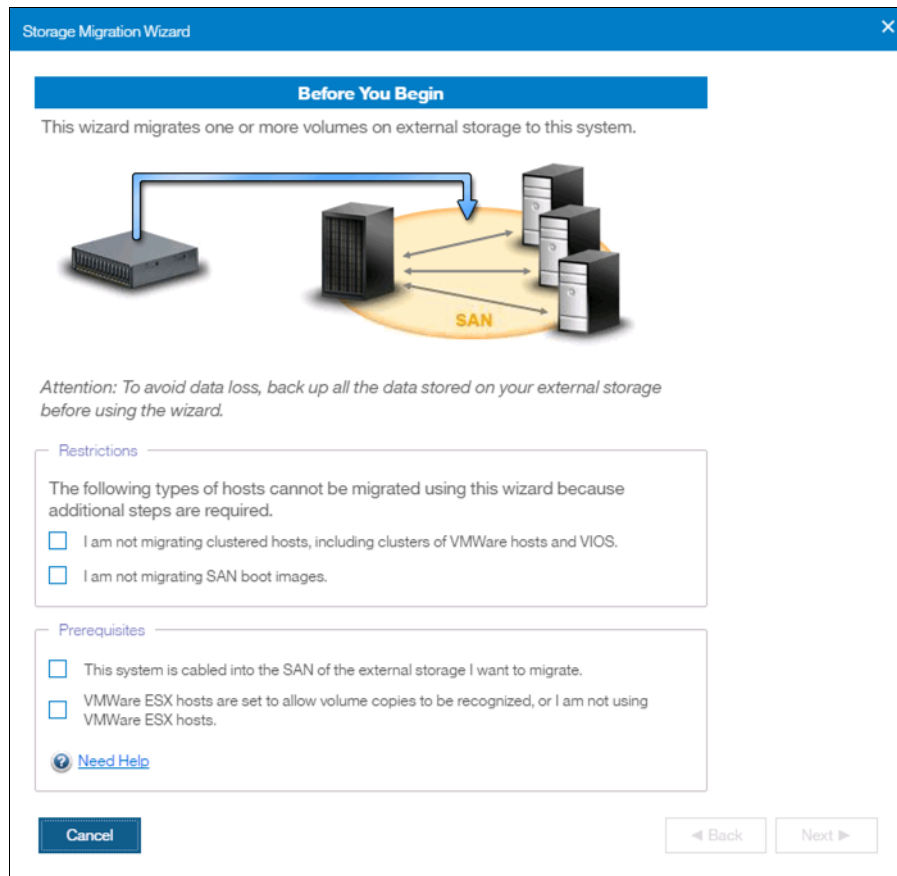


Figure 7-13 Before you begin the storage migration wizard

Restrictions

Confirm that the following conditions are met:

- ▶ You are not using the storage migration wizard to migrate cluster hosts, including clusters of VMware hosts and Virtual I/O Servers (VIOS).
- ▶ You are not using the storage migration wizard to migrate SAN Boot images.

If you identify that any of the restriction options cannot be selected, the migration must be performed outside of this wizard because more steps are required. For more information, see the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Lenovo Information Center at this web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/lenovo_vseries.html

The VMware vSphere Storage vMotion feature might be an alternative for migrating VMware clusters. For more information, see this web page:

<http://www.vmware.com/products/vsphere/features/storage-vmotion.html>

For more information about migrating SAN Boot images, see Appendix A, “CLI setup and SAN Boot” on page 761.

Prerequisites

Confirm that the following prerequisites apply:

- Ensure that the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, existing storage system, hosts, and Fibre Channel ports are physically connected to the SAN fabrics.
- If VMware ESX hosts are involved in the data migration, ensure that the VMware ESX hosts are set to allow volume copies to be recognized. For more information, see the VMware ESX product documentation at this web page:

<http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html?>

If all options can be selected, click **Next**. In all other cases, the button will not get available and the data must be migrated without the use of this wizard.

Preparing the environment for migration

Figure 7-14 shows the Prepare Environment for Migration panel. Follow the instructions that are provided in this panel carefully. When all of the required tasks are complete, click **Next**.

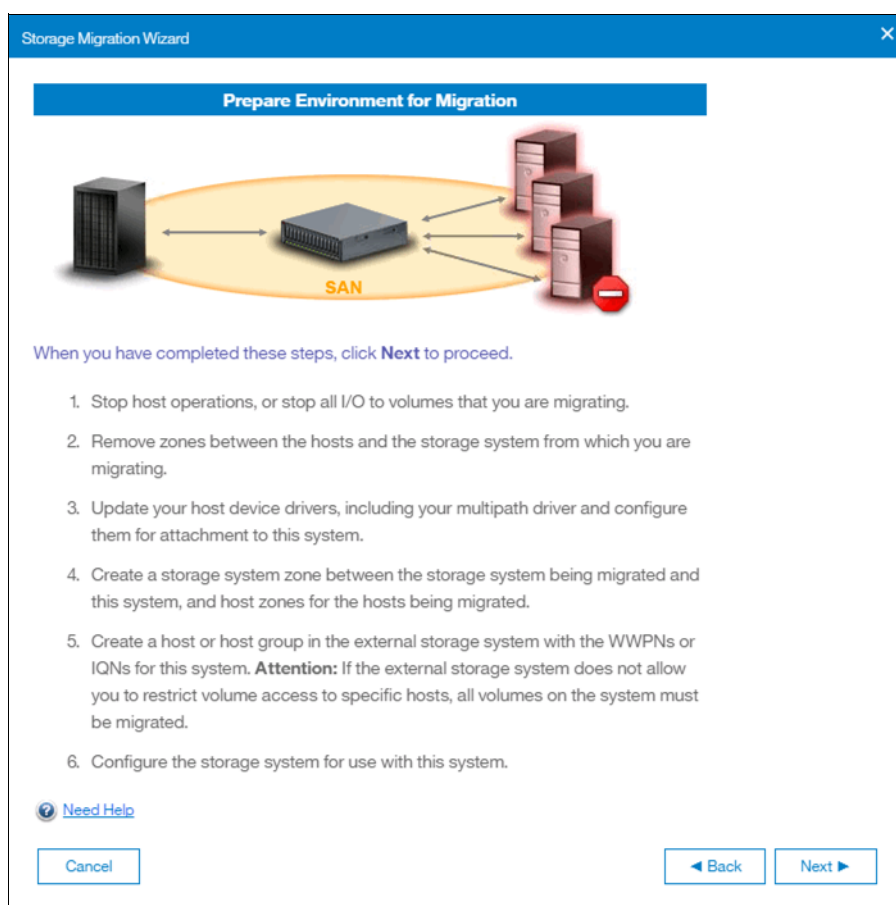


Figure 7-14 Prepare the migration environment

Mapping storage

Follow the instructions that are shown in the Map Storage panel that is shown in Figure 7-15 on page 338 and click **Next**. Record all of the details carefully because the information is used in later panels. Table 7-2 on page 338 shows an example table for you to capture the information that relates to the external storage system LUs.

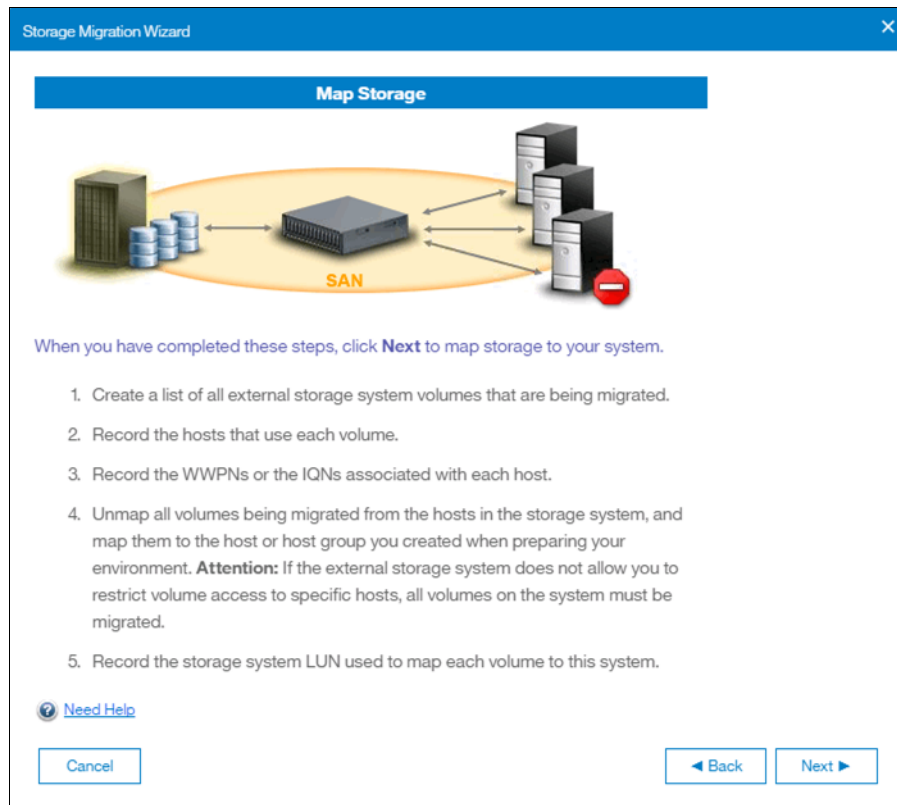


Figure 7-15 Directions to record migration data

Table 7-2 Example table to capture the external LU information

LU name	Controller	Array	SCSI ID	Host name	Capacity
V3700external0	Node2	V3700	0		50 GiB
V3700external1	Node1	V3700	1		50 GiB
V3700external2	Node2	V3700	2		50 GiB
V3700external3	Node1	V3700	3		50 GiB
V3700external4	Node2	V3700	4		50 GiB
V3700external5	Node1	V3700	5		50 GiB

SCSI ID: Record the SCSI ID of the LUs to which the host is originally mapped. Certain operating systems do not support the change of the SCSI ID during the migration.

Table 7-3 shows an example table to capture host information.

Table 7-3 Example table to capture host information

Host name	Adapter/Slot/Port	Worldwide port name (WWPN)	Host bus adapter (HBA) firmware	HBA device driver	Operating system	V5000 multipath software
mcr-host-153	QLE2562/2/1	21000024FF2D076C	2.10	9.1.9.25	Red Hat Enterprise Linux 5 (RHEL5)	Device Mapper
mcr-host-153	QLE2562/2/2	21000024FF2D076D	2.10	9.1.9.25	RHEL5	Device Mapper

After all of the data is collected and the tasks are performed in the Map Storage section, click **Next**. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 run the discover devices task and sequentially shows the Migrating MDisks panel.

Migrating MDisks

Select the MDisks from the existing storage system to migrate and click **Next**. Figure 7-16 shows the Migrating MDisks panel.

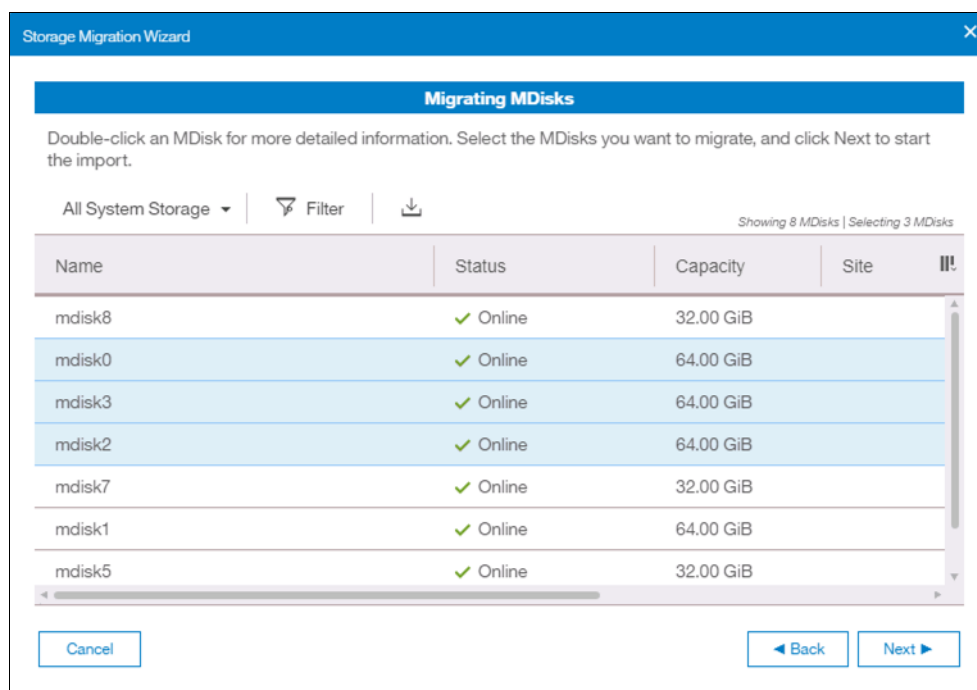


Figure 7-16 Migrating MDisks panel

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 run the Import MDisks task and sequentially shows the Configuring Hosts panel.

MDisk selection: Select only the MDisks that are applicable to the current migration plan. After the current migration completes, you can start another migration to migrate any remaining MDisks.

Configuring hosts

Note: This step is optional. You can bypass it by selecting **Next** and moving to “Mapping volumes to hosts” on page 341.

Follow this step of the wizard to select or configure new hosts as required. Figure 7-17 shows the Configure Hosts (optional) panel. If hosts are defined, they are listed in the panel as shown in Figure 7-19 on page 341. If no hosts are defined, they can be created by selecting the **Add Host** option.

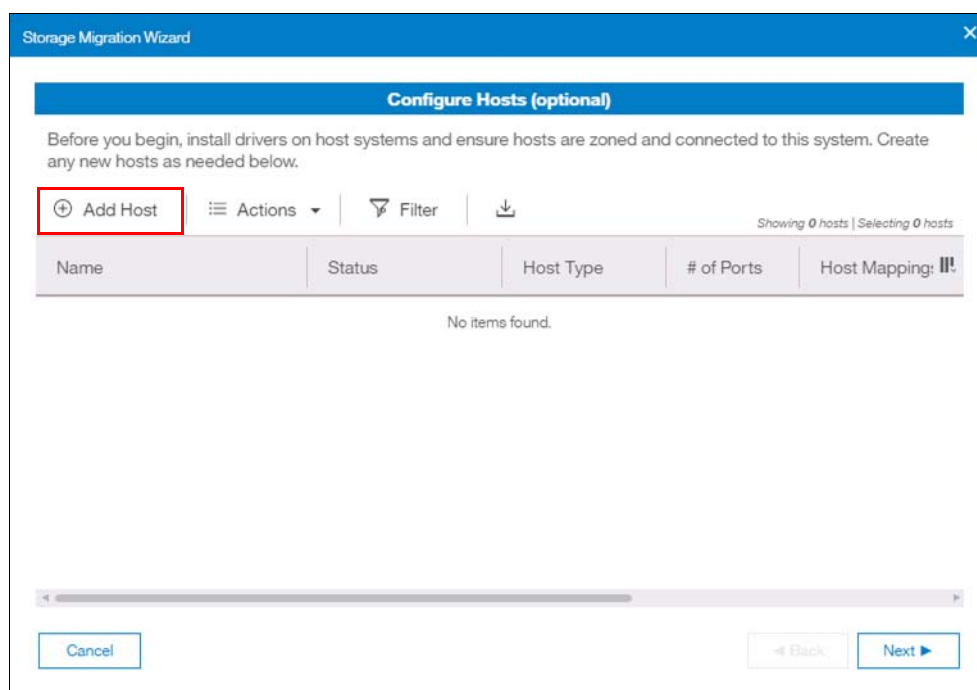
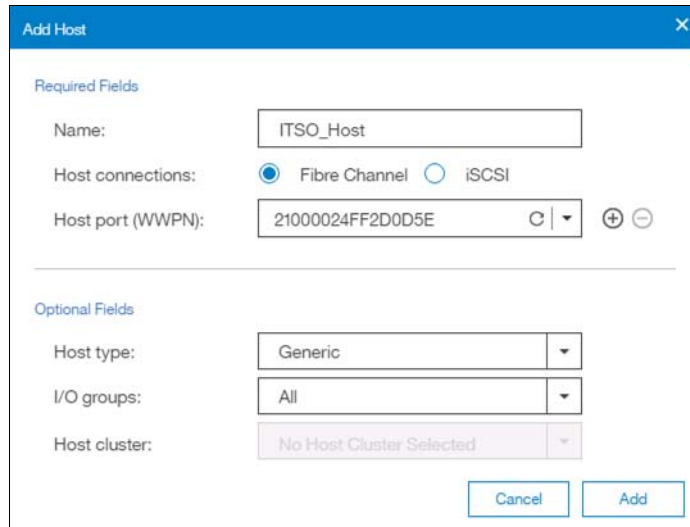


Figure 7-17 Configure Hosts panel

Select your connection type, name the host and assign the ports (in this case, Fibre Channel WWPNS). In the advanced settings, assign the I/O group ownership and host type as shown in Figure 7-18 on page 341. Click **Add** to complete the task. For more information about I/O group assignment, see Chapter 5, “Host configuration” on page 189.



Add Host

Required Fields

Name:

Host connections: ☒ Fibre Channel ☐ iSCSI

Host port (WWPN):

Optional Fields

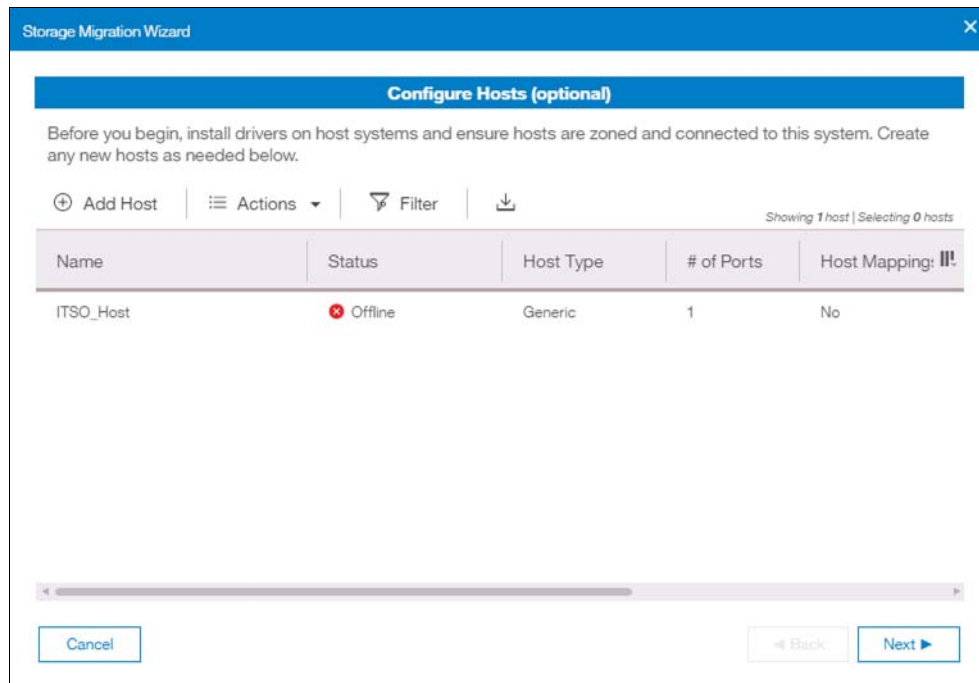
Host type:

I/O groups:

Host cluster:

Figure 7-18 The details to add a host are complete

The host is listed in the original Configure Hosts (optional) panel, as shown in Figure 7-19. Click **Next** to display the Map Volumes to Host (optional) panel.



Storage Migration Wizard

Configure Hosts (optional)

Before you begin, install drivers on host systems and ensure hosts are zoned and connected to this system. Create any new hosts as needed below.

Showing 1 host | Selecting 0 hosts

Name	Status	Host Type	# of Ports	Host Mapping: <input type="button" value="!"/>
ITSO_Host	<input checked="" type="radio"/> Offline	Generic	1	No

Figure 7-19 Configure Hosts panel that lists the host

Mapping volumes to hosts

Note: This step is optional. You can bypass it by selecting **Next** and moving to “Selecting a storage pool” on page 344.

Use this step of the wizard to select volumes that were migrated from the external storage system to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and map them to hosts. Hold Ctrl and click on the volume names to select multiple volumes. Click **Map to Host** to

open the Create Mapping panel. Figure 7-20 shows the Map Volumes to Hosts (optional) panel.

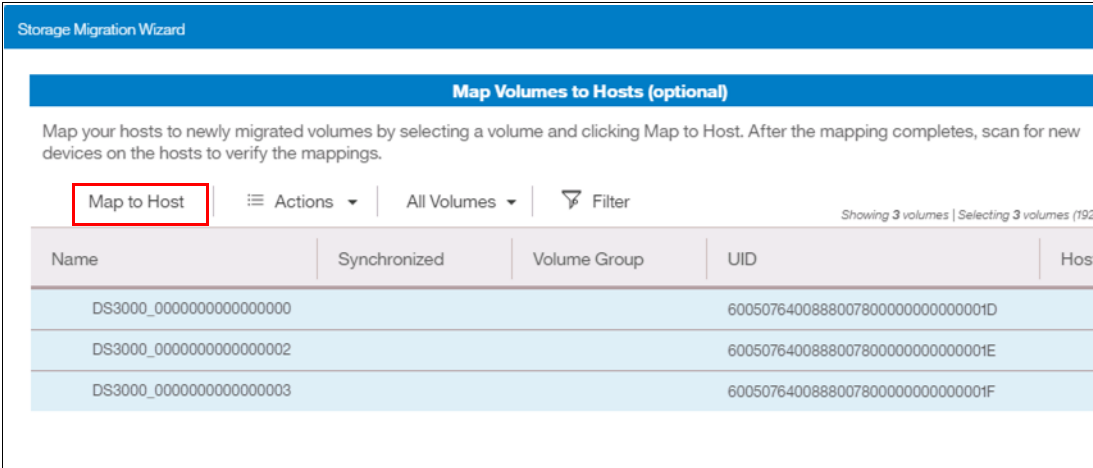


Figure 7-20 Map Volumes to Hosts panel

The image mode volumes are listed, the names are assigned automatically by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems and can be changed to reflect more meaningful names to the user by selecting the volume and clicking **Rename** in the Actions menu.

Names: The names of the image mode volumes must begin with a letter. The name can be a maximum of 63 characters. You can use the following valid characters:

- ▶ Uppercase letters (A - Z)
- ▶ Lowercase letters (a - z)
- ▶ Digits (0 - 9)
- ▶ Underscore (_)
- ▶ Period (.)
- ▶ Hyphen (-)
- ▶ Blank space

The names must not begin or end with a space.

Select from the host list the hosts to which the imported volumes will be mapped as shown in Figure 7-21 on page 343, and click **Next**.

Create Mapping

Create Mappings to:

☒ Hosts
 ☐ Host Clusters

Select hosts to map to 3 volumes

Filter

Showing 1 host | Selecting 1 host

Name	Status	Host Type	Host Mappings
ITSO_Host	Offline	Generic	No

Would you like the system to assign SCSI LUN IDs or manually assign these IDs?

☒ System Assign
 ☐ Self Assign

Cancel

Back

Next

Figure 7-21 Modify host mappings

Note: If you select Host Clusters in the Create Mapping panel you need to ensure that the host cluster has consistent access to I/O groups. If each host in the host cluster does not have access to the same I/O groups, mappings to volumes will fail.

A confirmation screen is displayed with the task summary as shown in Figure 7-22 on page 344. Click **Map Volumes** to finish the task and return to the Map Volumes to Hosts (optional) panel.

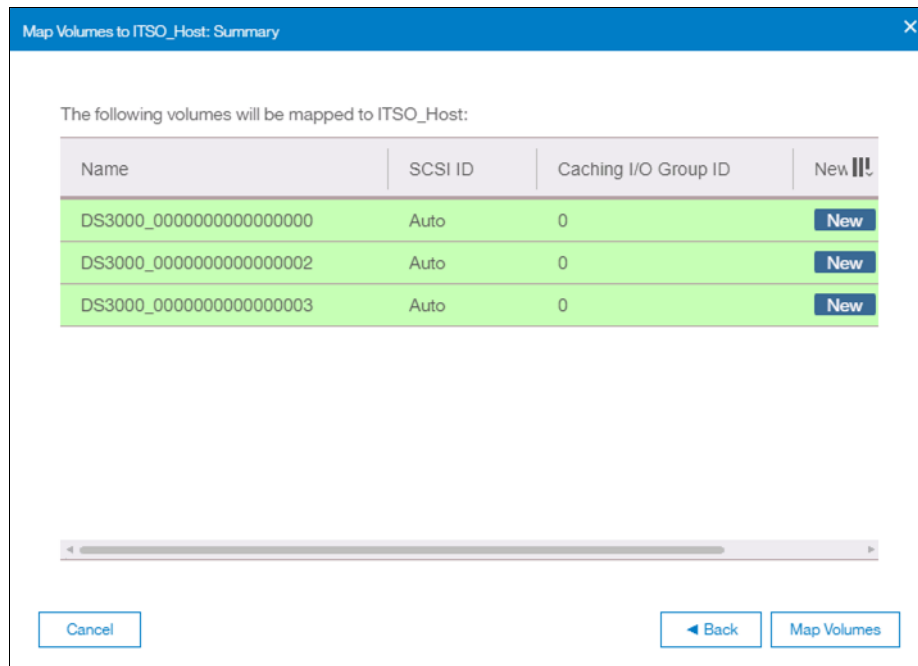


Figure 7-22 Mapping volumes summary

Figure 7-23 shows that the host mappings are in place for the chosen volumes. Click **Next**.

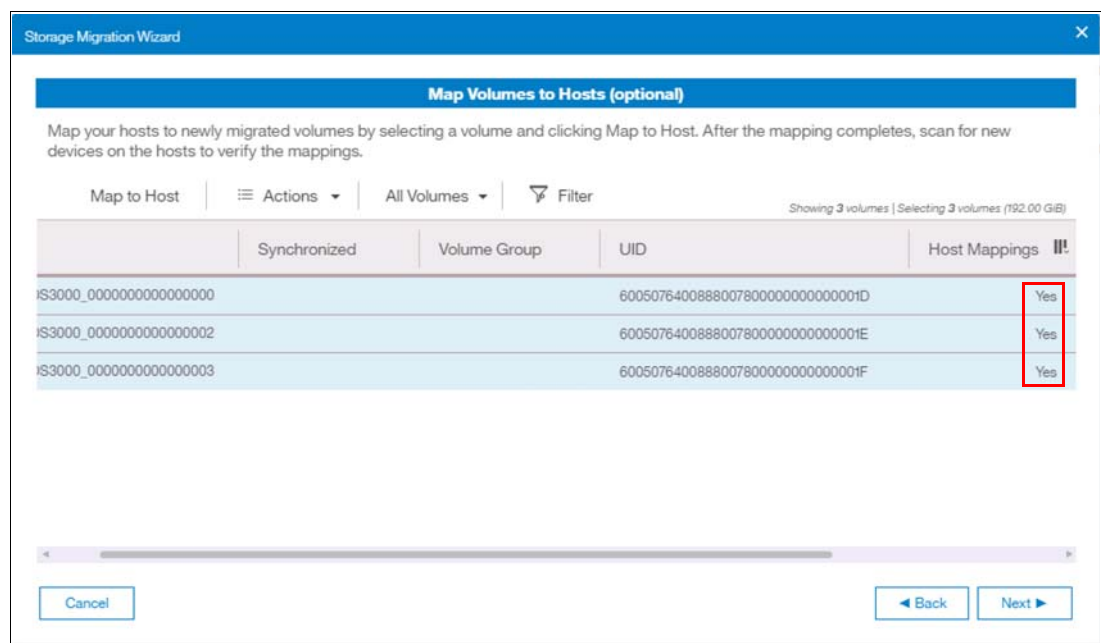


Figure 7-23 Map Volumes to Hosts panel with completed host mappings

Selecting a storage pool

Note: This step is optional. You can bypass it by avoiding a pool selection, clicking **Next** and moving to “Finishing the storage migration wizard” on page 345.

To continue with the storage migration wizard, select a storage pool to migrate the imported volumes to, as shown in Figure 7-24. Click **Next** to proceed to the last panel of the storage migration wizard. The process uses the volume mirroring function that is included within the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

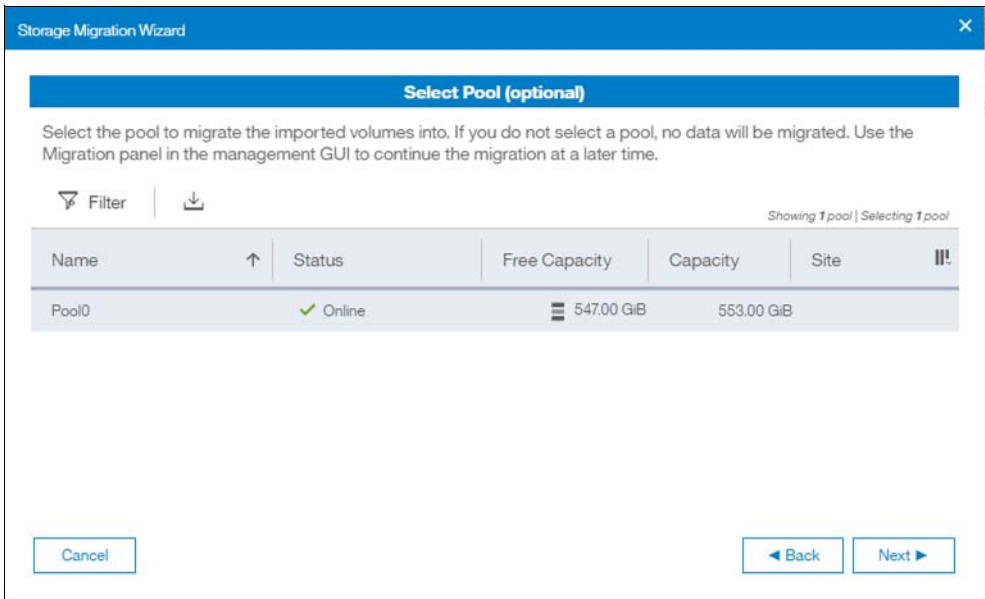


Figure 7-24 Storage pool selection

Finishing the storage migration wizard

Click **Finish** to end the storage migration wizard as shown in Figure 7-25.

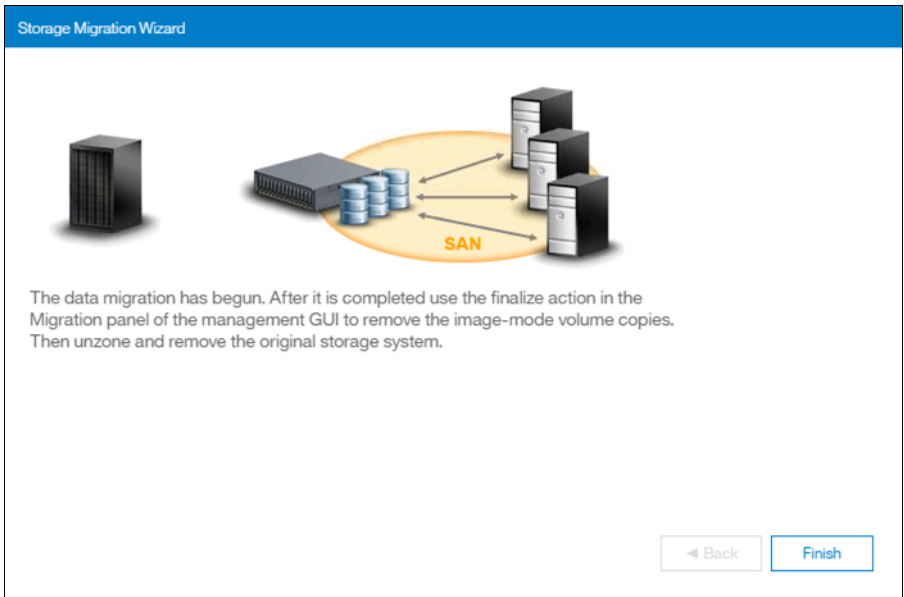


Figure 7-25 Migration wizard complete

The end of the storage migration wizard is not the end of the data migration task. It is still in progress. A percentage indicator is displayed in the Storage Migration panel as shown in Figure 7-26.

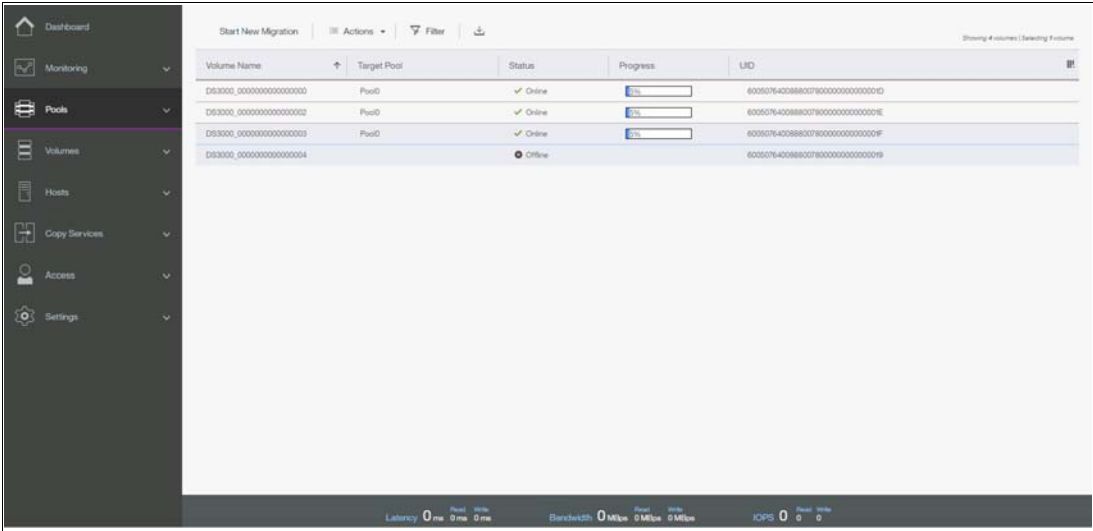


Figure 7-26 Storage migration progress

Finalizing migrated volumes

When the migration completes with all of the progress indicators at 100%, select all of the volume migrations that you want to finalize by holding down Ctrl and clicking the volumes.

Then, select **Actions** → **Finalize** as shown in Figure 7-27. Alternatively, right-click the selected volumes and click **Finalize**.

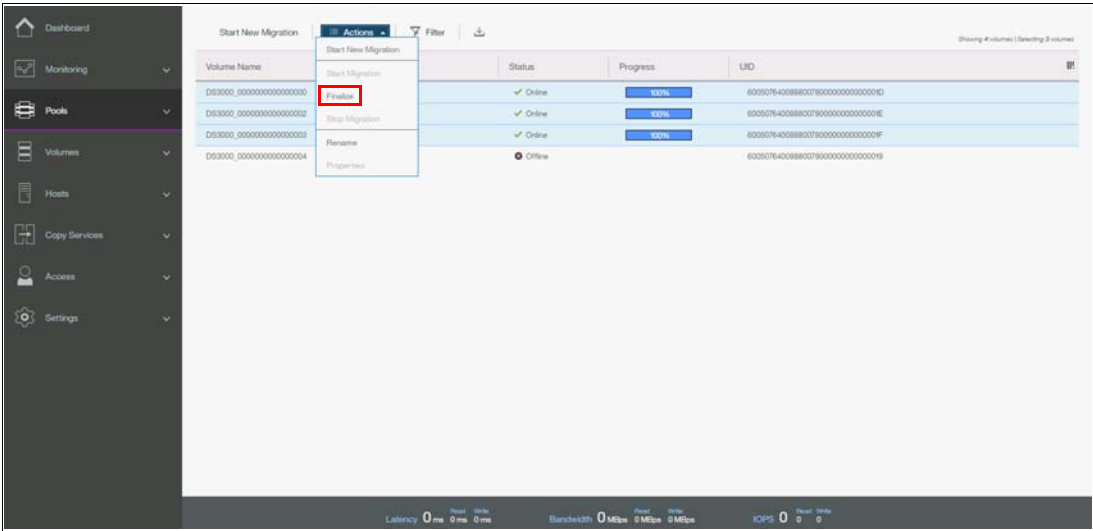


Figure 7-27 Finalize storage migration

You are asked to confirm the number of volume migrations that you want to finalize as shown in Figure 7-28. Verify that the volume names and the number of migrations are correct and click **OK**.

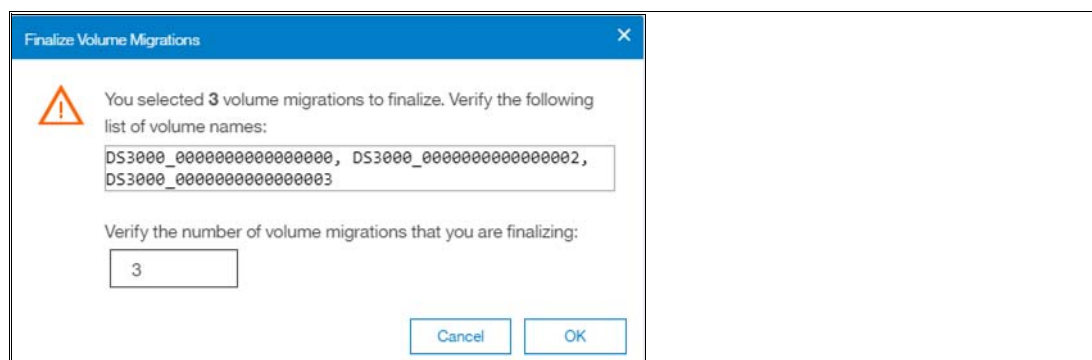


Figure 7-28 Confirm volumes to finalize

When the finalization completes, the data migration to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 is completed. The zoning can be removed and the external storage system can be retired.

Advanced host and volume administration

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 offer many functions for volume and host configuration. The basic host and volume features of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are described in Chapter 5, “Host configuration” on page 189 and Chapter 6, “Volume configuration” on page 269. The chapters also describe how to create hosts and volumes and how to map them to a host.

This chapter focuses on advanced host and volume administration topics. The first part of it describes the following host administration topics:

- ▶ 8.1, “Advanced host administration” on page 350
- ▶ 8.2, “Adding and deleting host ports” on page 367

The second part of the chapter consists of the following volume-related tasks:

- ▶ 8.3, “Advanced volume administration” on page 373
- ▶ 8.4, “Volume properties and volume copy properties” on page 386
- ▶ 8.5, “Advanced volume copy functions” on page 390
- ▶ 8.6, “Volumes by storage pool” on page 399
- ▶ 8.7, “Volumes by host” on page 400

8.1 Advanced host administration

This section describes advanced host administration, including host modification, host mappings, and deleting hosts. Basic host creation and mapping are described in Chapter 5, “Host configuration” on page 189. We assume that hosts are defined and volumes are mapped to them.

The following topics are covered in this section:

- ▶ Modifying hosts, as described in 8.1.1, “Modifying volume mappings” on page 351
- ▶ Ports by host, as described in 8.2, “Adding and deleting host ports” on page 367
- ▶ Host mappings, as described in 8.2, “Adding and deleting host ports” on page 367

The top-level Hosts menu is shown in Figure 8-1.

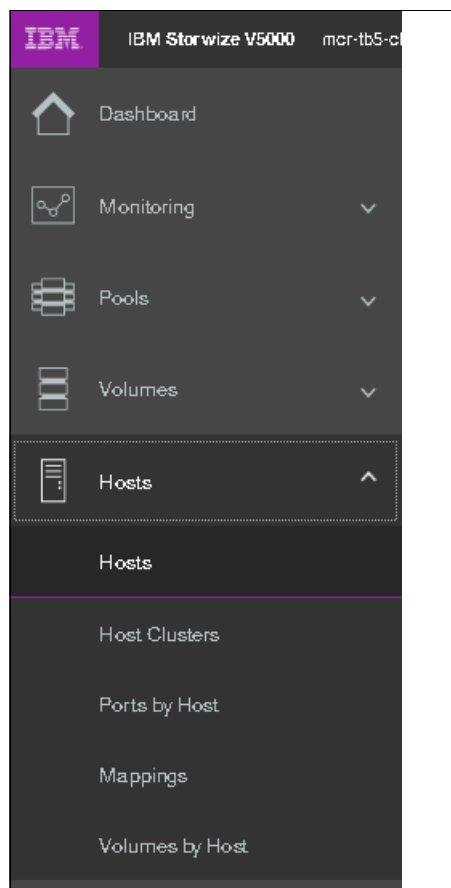


Figure 8-1 Hosts menu

Select **Hosts** to open the Hosts panel, as shown in Figure 8-2.

The image shows the 'Hosts' panel in the management console. At the top, there is a header bar with 'Add Host', an 'Actions' dropdown, a search box, and a download icon. Below this is a table with the following columns: Name, Status, Host Type, # of Ports, Host Mappings, Host Cluster ID, and Host Cluster Name. The table contains two rows of data.

Name	Status	Host Type	# of Ports	Host Mappings	Host Cluster ID	Host Cluster Name
RHEL_HOST_1	Online	Generic	2	Yes	2	RHEL_HOST_CLUSTER
RHEL_HOST_2	Online	Generic	2	No		

Figure 8-2 Hosts panel

Select a host and click **Actions** (as shown in Figure 8-3 on page 351) or right-click the host to show the available actions.

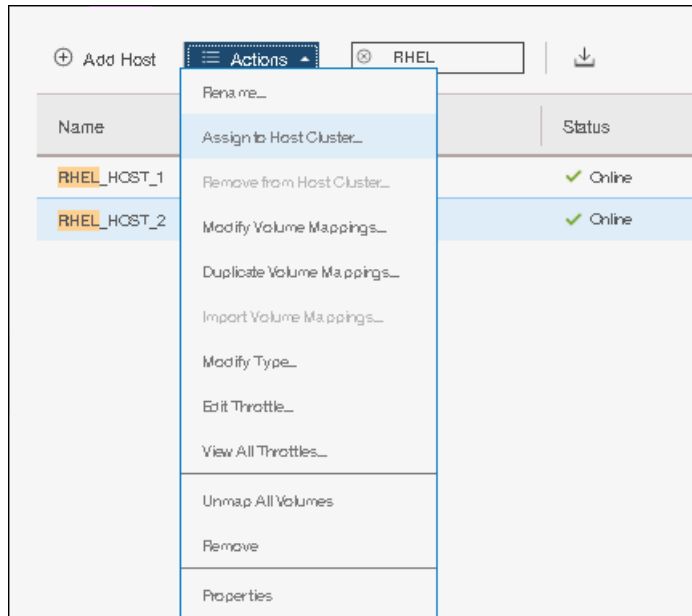


Figure 8-3 Actions menu on the Hosts panel

As shown in Figure 8-3, several actions are associated with host mapping. For more information, see 8.1.1, “Modifying volume mappings” on page 351 and 8.1.2, “Unmapping volumes from a host” on page 354.

8.1.1 Modifying volume mappings

From the Hosts panel, after you select a host and click **Actions**, select **Modify Volume Mappings** to open the Modify Host Mappings panel, as shown in Figure 8-4.

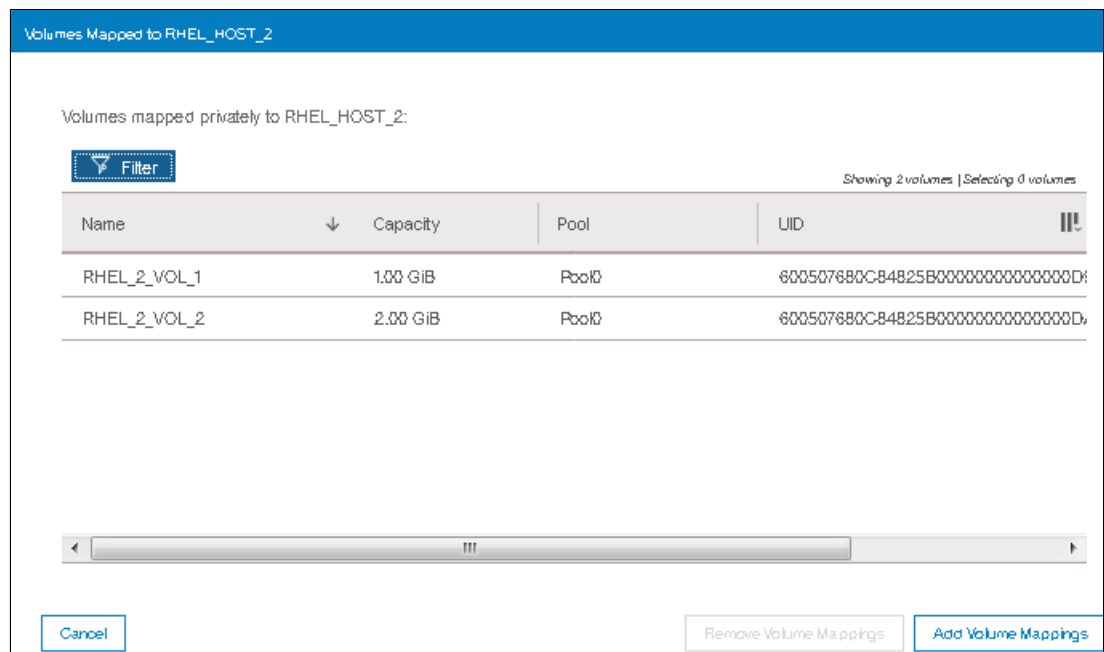


Figure 8-4 Host mappings panel

Click **Add Volume Mappings**. A window will come up listing all the additional volumes which can be mapped to the selected host as shown in Figure 8-5.

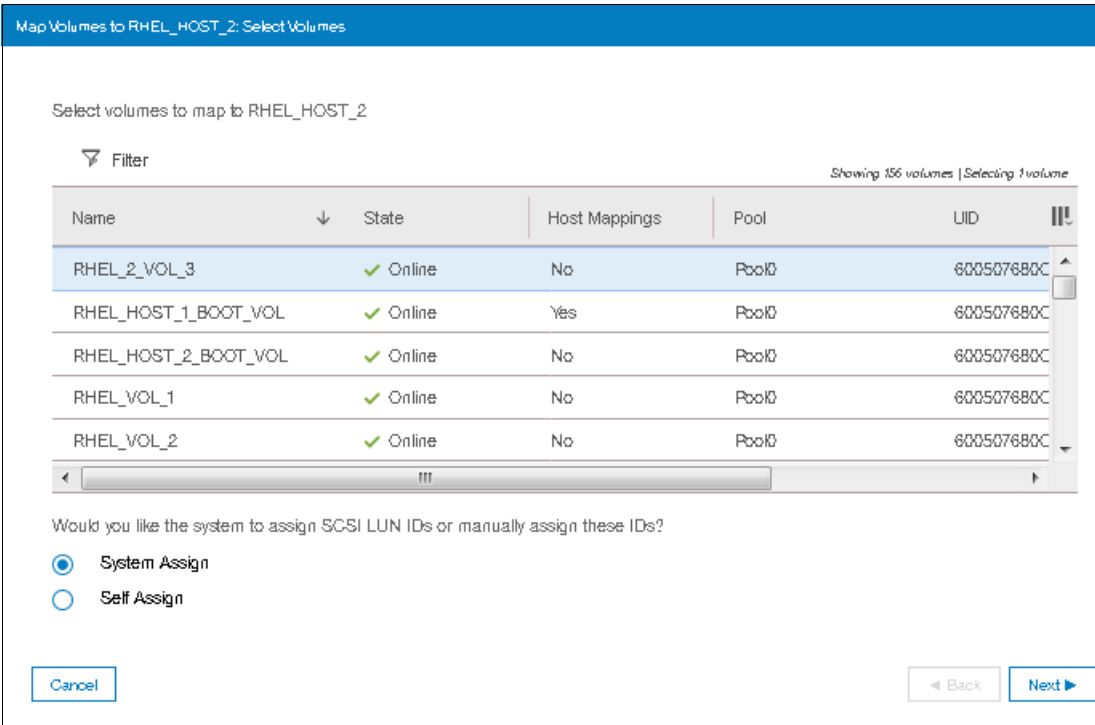


Figure 8-5 Select a host to modify

Select the volume that you need to map to the host. Also indicate whether you would like to let the system assign the SCSI ID or you want to assign the SCSI ID by yourself. In this example, RHEL_2_VOL_3 volume is being mapped with the user supplied SCSI ID as shown in Figure 8-6 on page 353.

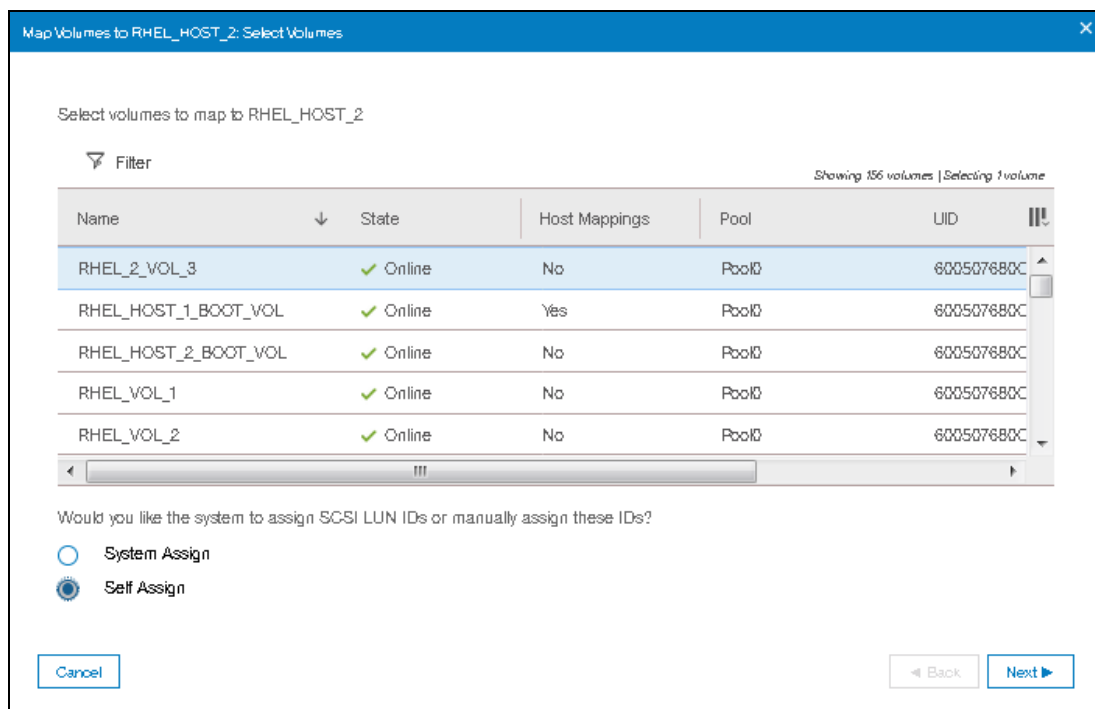


Figure 8-6 Selecting the volume and SCSI ID for mapping to a host

Click **Next**. As shown in Figure 8-7, a window opens where the user can provide the SCSI ID to be used for the mapping. The right hand pane side also shows the current SCSI ID's being used for mapping other volumes to the same host.

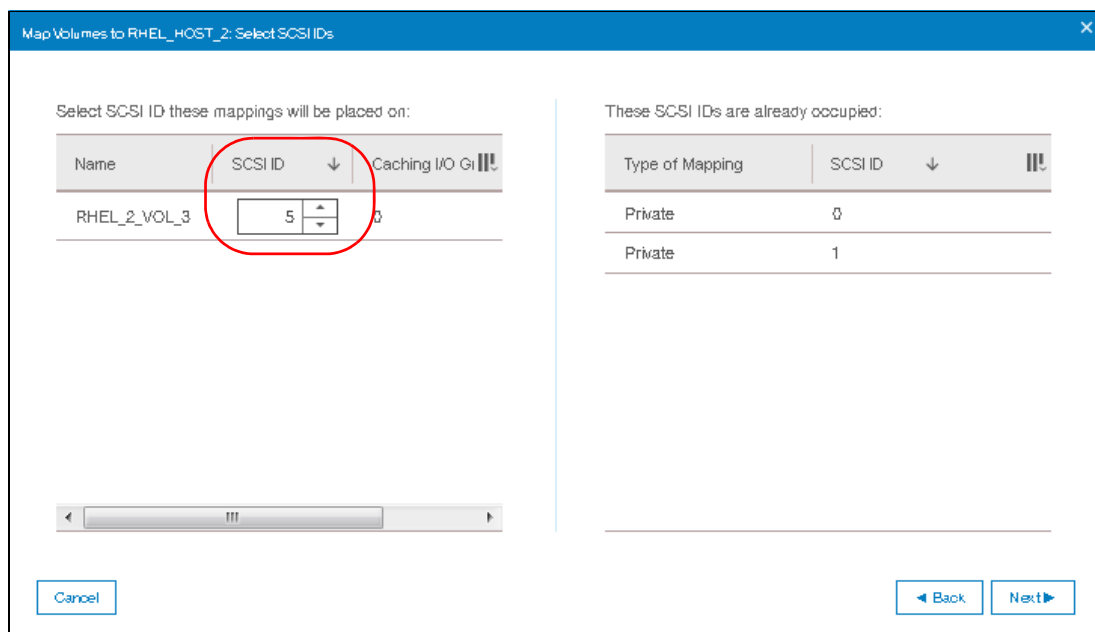


Figure 8-7 Selecting the SCSI ID for the mapping

Important: The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 automatically assign the lowest available SCSI ID if none is specified. However, you can set a SCSI ID for the volume. The SCSI ID cannot be changed while the volume is assigned to the host.

Click **Next**. A window will open indicating new mapping to the host, along with already mapped volumes to that host, as shown in Figure 8-8. SCSI ID's for each of the mappings is also shown as part of this process.

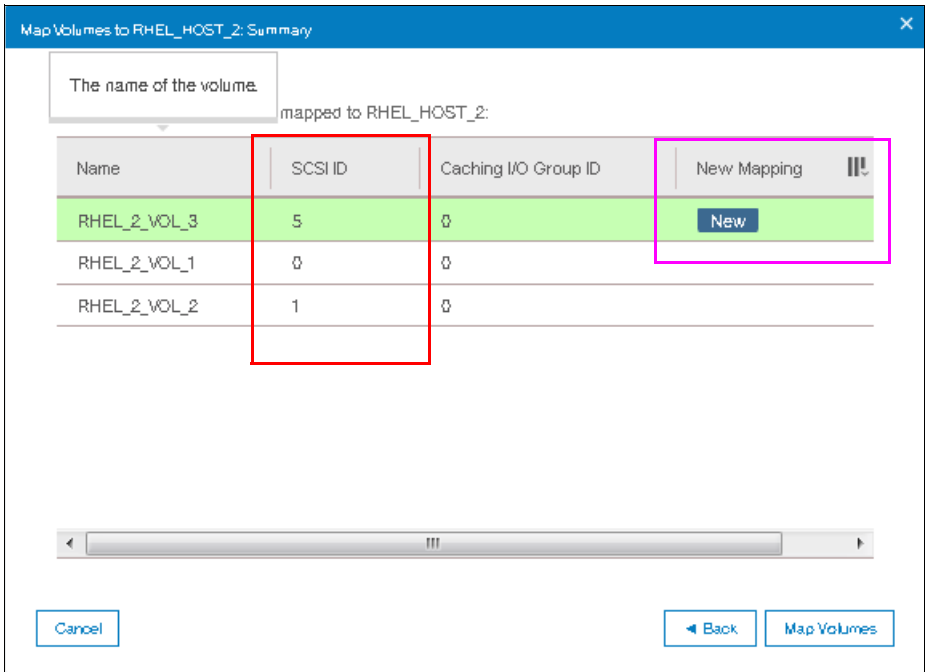


Figure 8-8 Volume map

Click **Map Volumes**. A confirmation window is presented as shown in Figure 8-9.

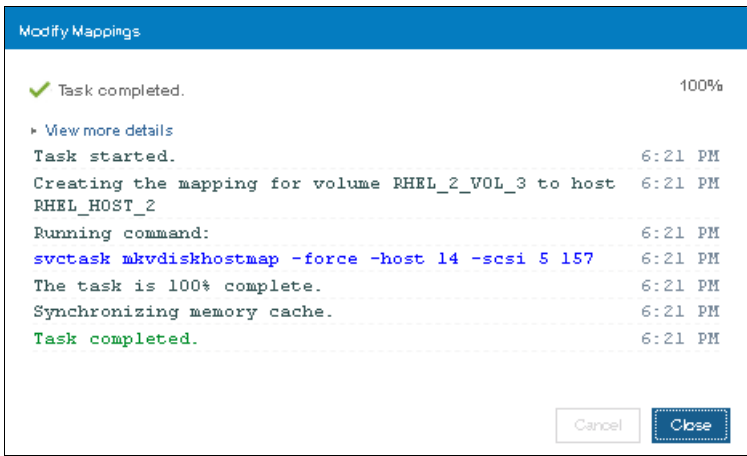


Figure 8-9 Volume mapping confirmation window

8.1.2 Unmapping volumes from a host

To unmap a volume from a host, follow the steps listed here:

1. From the main navigation pane, Click **Hosts** as shown in Figure 8-10 on page 355.




Figure 8-10 Hosts

- From the right hand side pane, select the host for which you want to unmap a volume, and then click Action as shown in Figure 8-11.

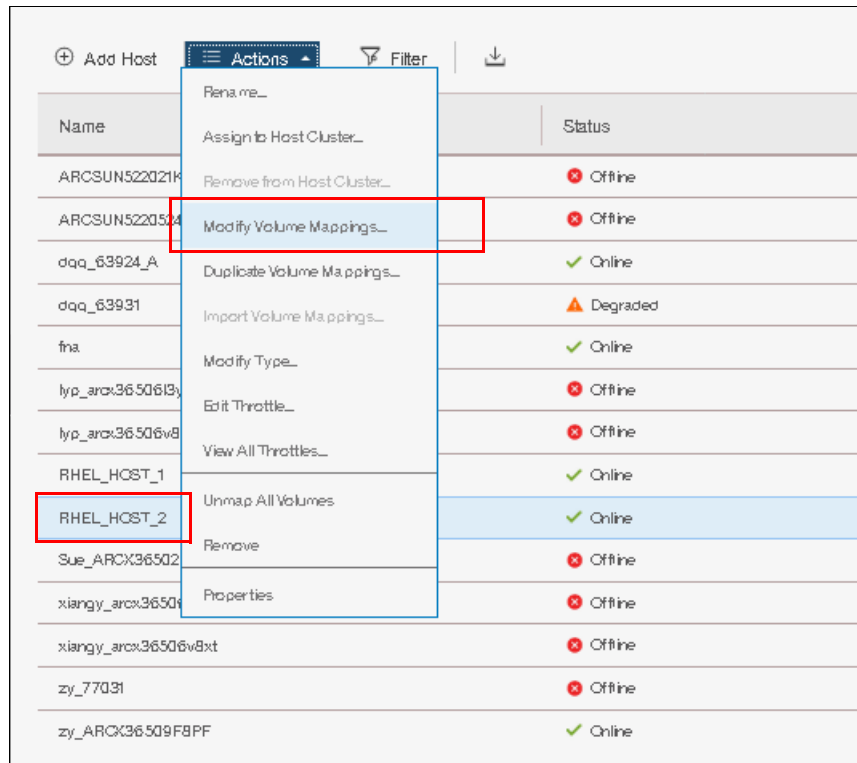


Figure 8-11 Action for the selected host

- Click **Modify Volume Mapping**. A window with a list of volumes currently mapped to the selected host will be shown as in Figure 8-12 on page 356.

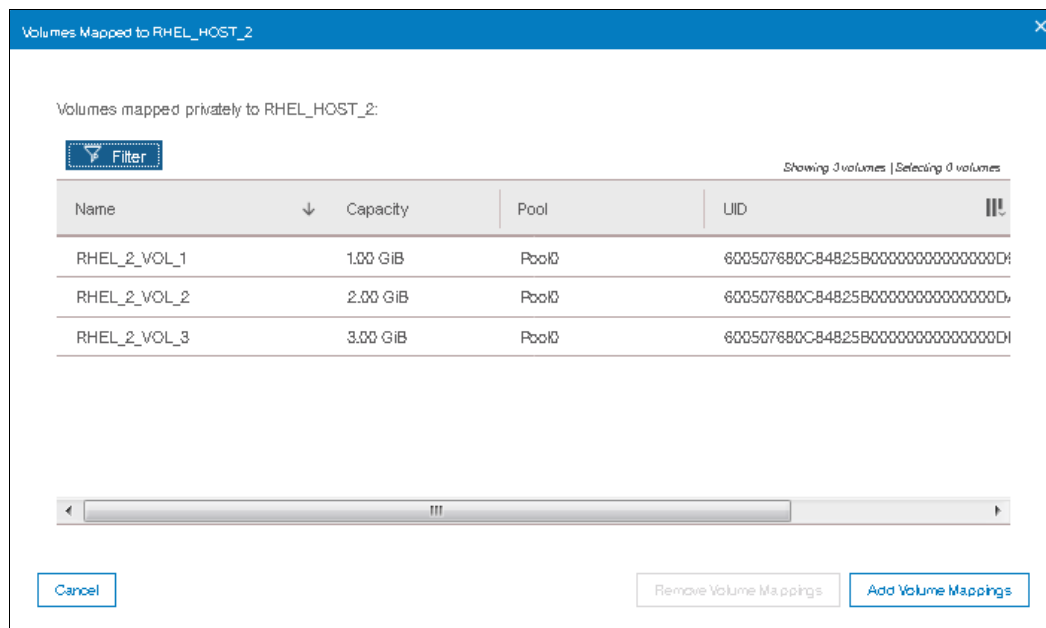


Figure 8-12 List of volumes mapped to the host

4. Select the volume which you want to unmap and then click **Remove Volume Mappings** as shown in Figure 8-13.

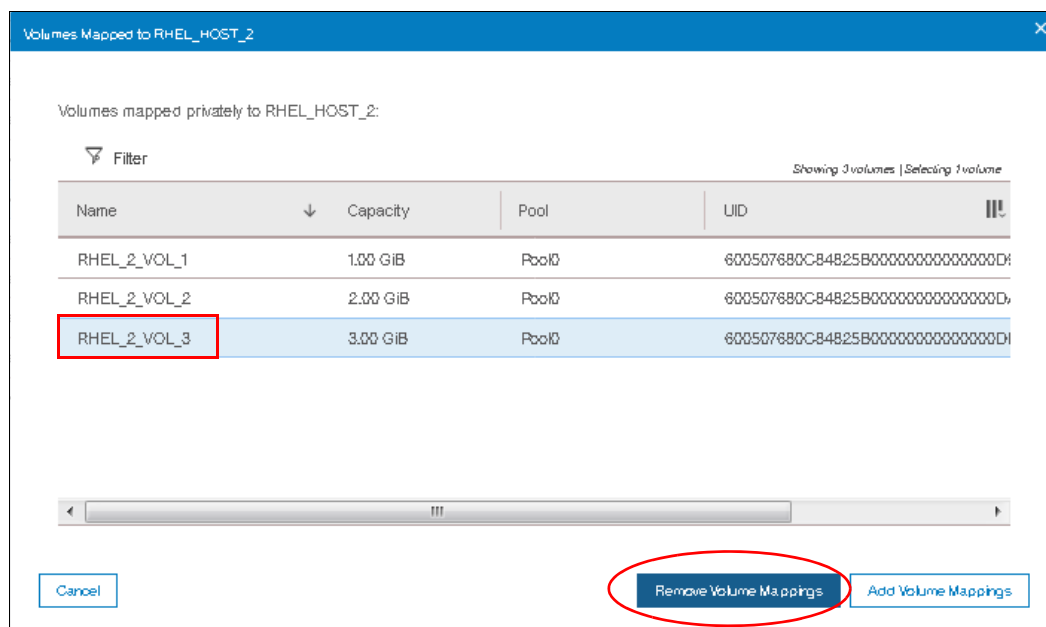


Figure 8-13 Selecting the volume to be unmapped

Note: To unmap multiple volumes click and hold the shift key and then select each consecutive volume in the window. If the multiple volumes you want to unmap are not consecutive, then click the Ctrl key and then select each volume as desired.

5. A window listing the volume that will be unmapped will be as shown in Figure 8-14 on page 357.

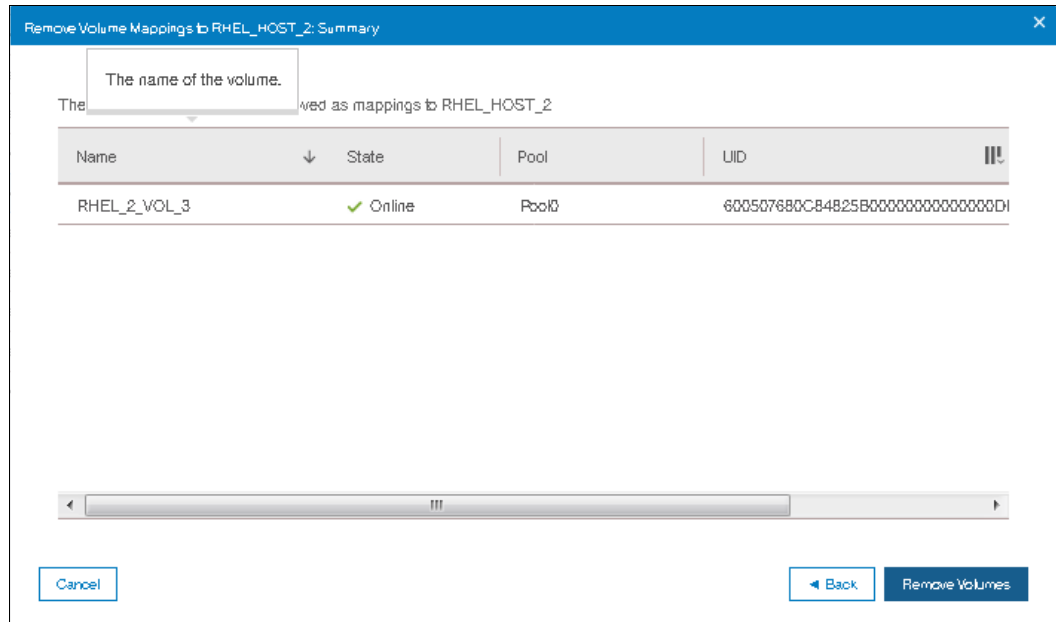


Figure 8-14 Remove volume mapping summary window

- Click **Remove Volumes** and a task completion confirmation window will open as shown in Figure 8-15.

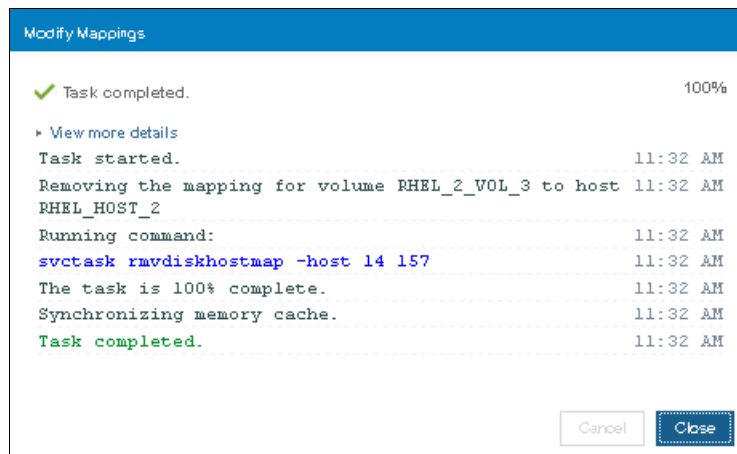


Figure 8-15 Volume unmap task completion

- The selected volume is now unmapped.

8.1.3 Renaming a host

To rename a host, follow the steps as listed:

- From the main navigation panel, select **Hosts** as shown in Figure 8-16 on page 358.




Figure 8-16 Hosts

2. Select the host that needs to be renamed and click **Action** as shown in Figure 8-17 on page 359.

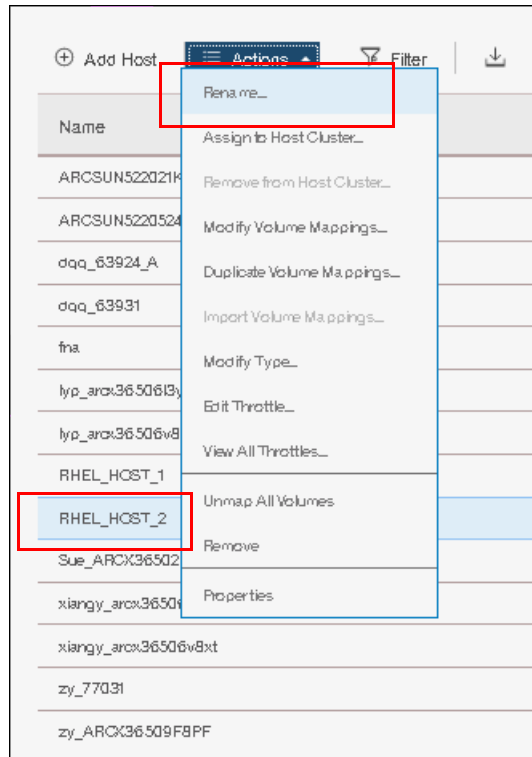


Figure 8-17 Selecting Rename action for host

3. A window will open where you can enter the new name as shown in Figure 8-18.




Figure 8-18 Rename a host window

4. Type in the new host name and click **Rename** as shown in Figure 8-19




Figure 8-19 Renaming a host

5. The selected host will be renamed and a task completion confirmation window will be shown as in Figure 8-20 on page 360.

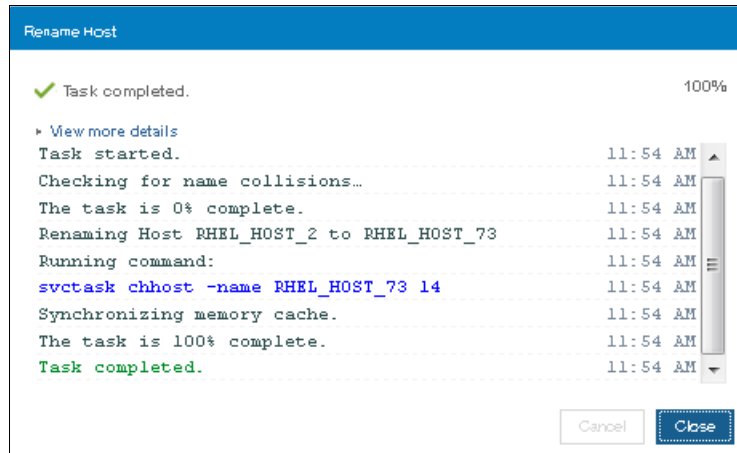


Figure 8-20 Rename a Host task completion

6. Click **Close**.

8.1.4 Removing a host

o rename a host, follow the steps as listed:

1. From the main navigation panel, select **Hosts** as shown in Figure 8-16 on page 358.

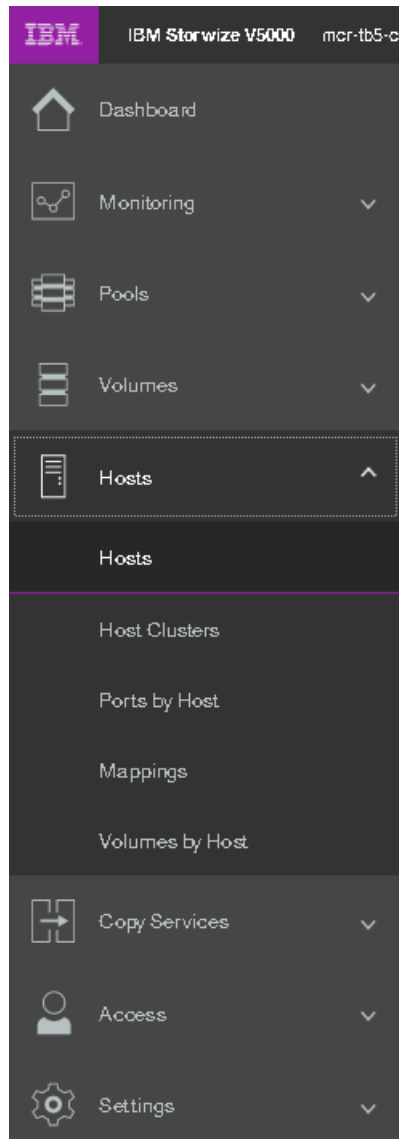


Figure 8-21 Hosts

2. Select the host that needs to be renamed and click **Action** as shown in Figure 8-22 on page 362.

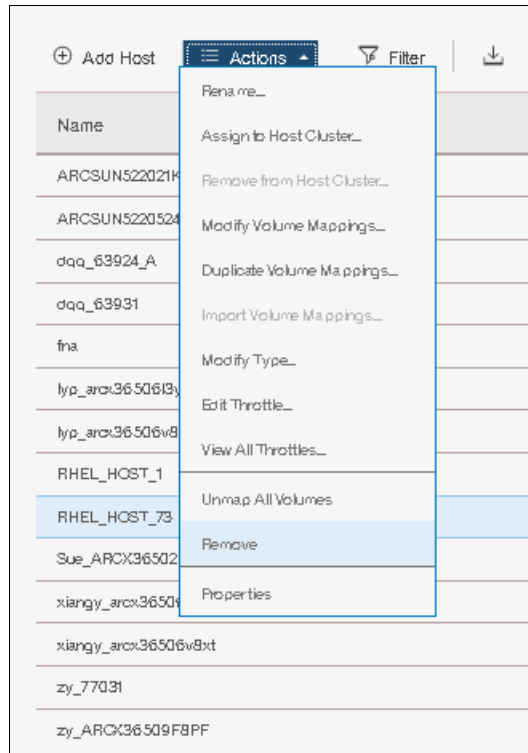


Figure 8-22 Selecting Remove operation

3. Verify the number of hosts you are removing, along with confirmation to remove the host even if the host has mapped volumes, as shown in Figure 8-23.




Figure 8-23 Removing a host

Note: If a host has volumes mapped, then to remove the host, you will have to check the box for **Remove the hosts even if volumes are mapped to them. These volumes will no longer be accessible to the hosts** to force the action.

4. A task completion window will open as shown in Figure 8-24 on page 363.

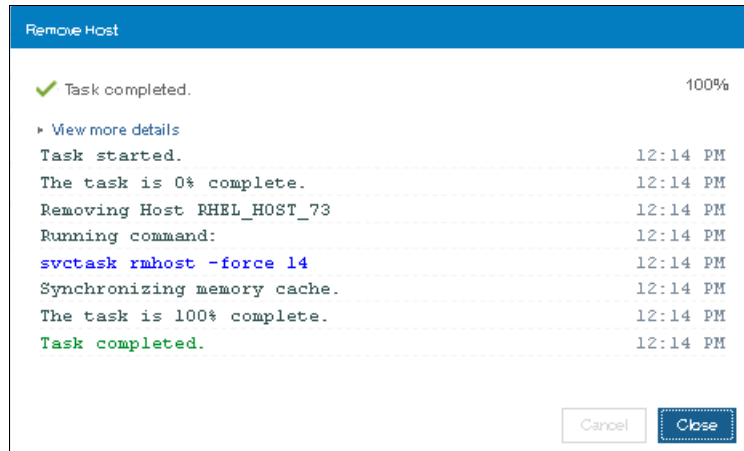


Figure 8-24 Host removal task completion

5. Click **Close**.

8.1.5 Host properties

This section describes the host properties, which provide the following information:

- Overview
- Mapped volumes
- Port definitions

Overview

To open the Host Details panel, select the host. From the Actions menu, click **Properties**. You also can highlight the host and right-click to access the Actions menu, as shown in Figure 8-25 on page 364.

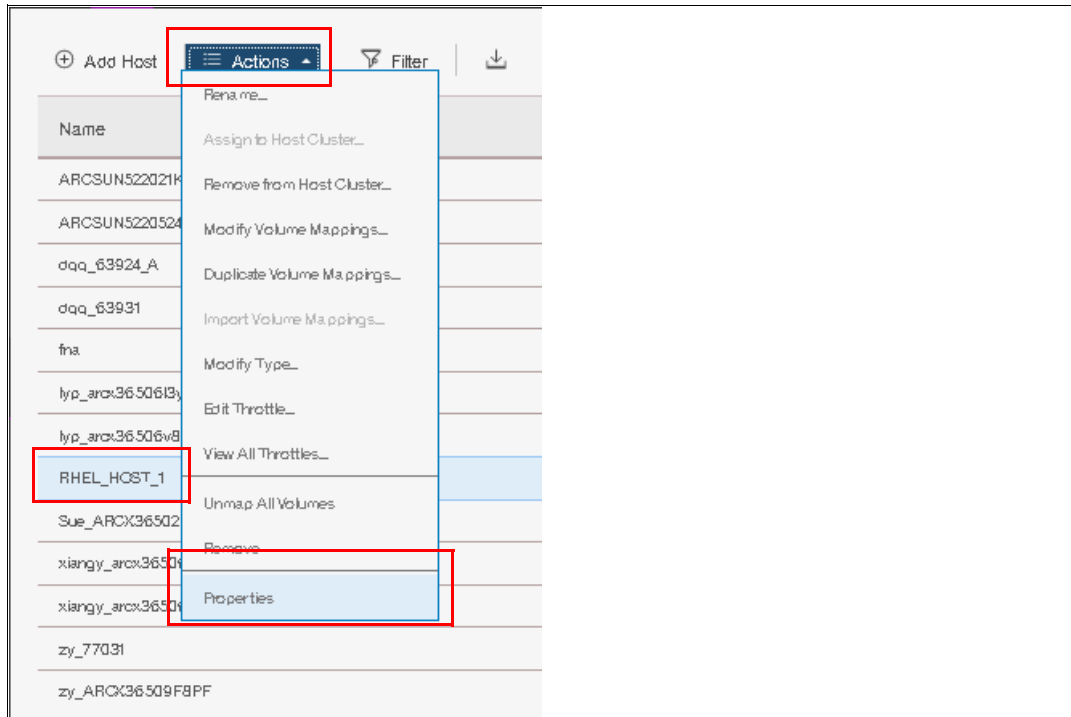


Figure 8-25 Opening host properties

Figure 8-26 shows the Overview tab of the Host Details panel. Select the **Show Details** check box in the lower-left corner of the window to see more information about the host.

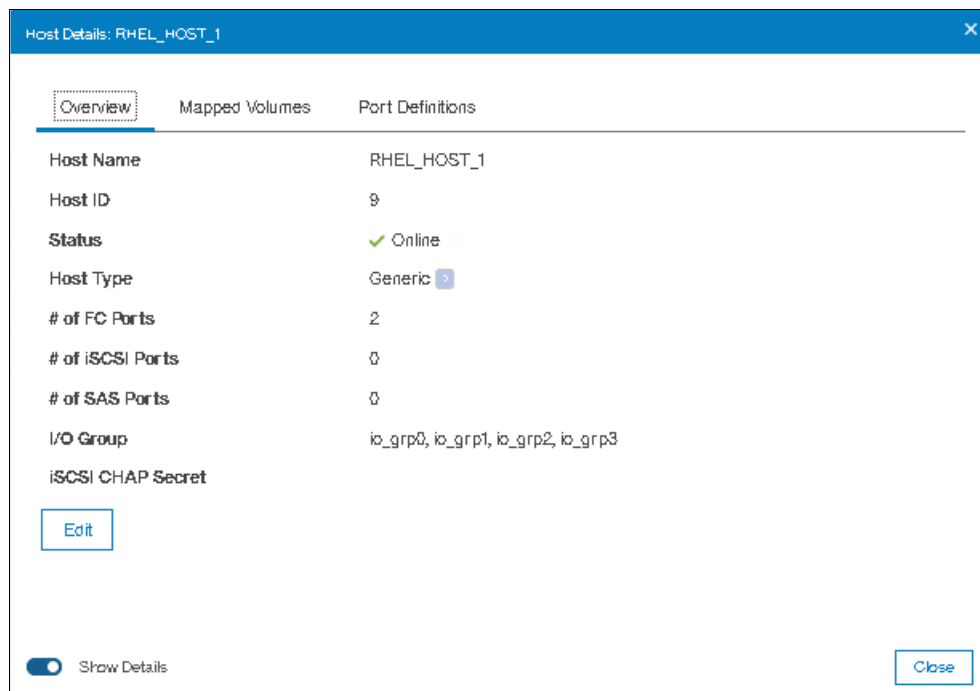


Figure 8-26 Host details

The Overview tab provides the following information:

- Host name: Host object name.

- ▶ **Host ID:** Host object identification number.
- ▶ **Status:** The current host object status. This value can be Online, Offline, or Degraded.
- ▶ **Host type:** The type of host can be Generic, Generic (hidden secondary volumes), HP/UX, OpenVMS, Target Port Group Support (TPGS), and VMware Virtual Volume (VVOL).
- ▶ **Number of Fibre Channel (FC) ports:** The number of host Fibre Channel ports.
- ▶ **Number of Internet SCSI (iSCSI) ports:** The number of host iSCSI names or host iSCSI qualified names (IQN) IDs.
- ▶ **Number of serial-attached SCSI (SAS) ports:** The number of host SAS ports.
- ▶ **I/O group:** The I/O group from which the host can access a volume (or volumes).
- ▶ **iSCSI Challenge Handshake Authentication Protocol (CHAP) secret:** The CHAP information if it exists or if it is configured.

To change the host properties, click **Edit**. Several fields can be edited, as shown in Figure 8-27.

Figure 8-27 Host properties: Editing the host information

For the host type, choose one of these values: Generic, Generic (hidden secondary volumes), HP/UX, OpenVMS, TPGS, or VVOL.

After you change any host information, click **Save** to apply your changes.

Mapped Volumes

Figure 8-28 on page 366 shows the Mapped Volumes tab, which provides an overview of the volumes that are mapped to the host. This tab provides the following information:

- ▶ SCSI ID
- ▶ Volume name
- ▶ Unique identifier (UID)
- ▶ Caching I/O group ID

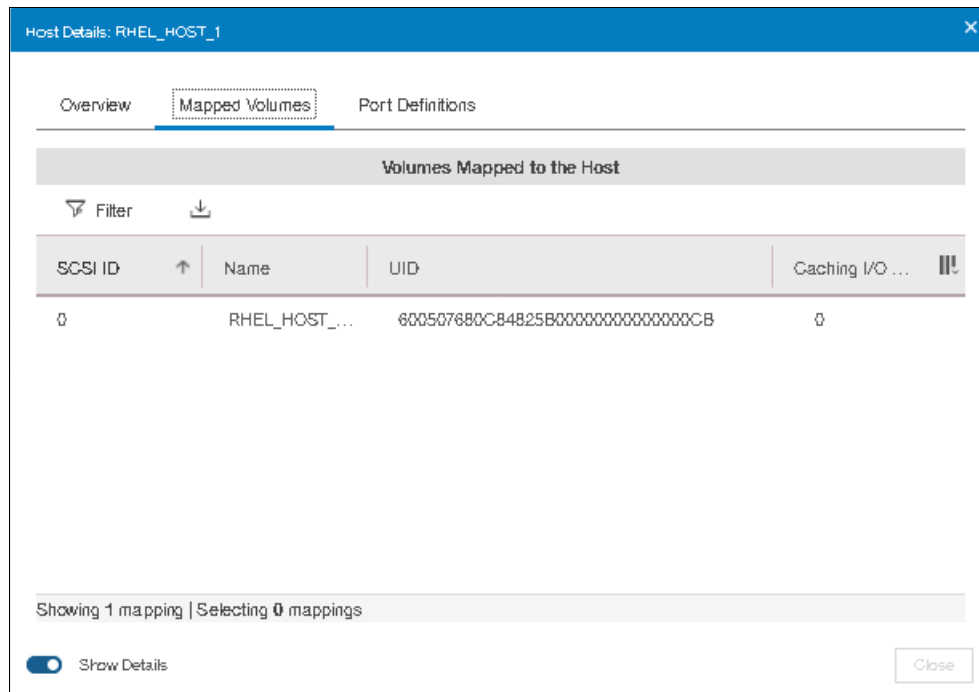


Figure 8-28 Host Details: Mapped Volumes tab

Port Definitions

Figure 8-29 on page 367 shows the Port Definitions tab, which shows the configured host ports and their status. This tab provides the following information:

- ▶ Name: The worldwide port names (WWPNs) (for SAS and FC hosts) or iSCSI Qualified Name (IQN) for iSCSI hosts
- ▶ Type: Port type
- ▶ Status: Current port status
- ▶ Number of nodes that are logged in: Lists the number of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 node canisters that each port (initiator port) is logged in to

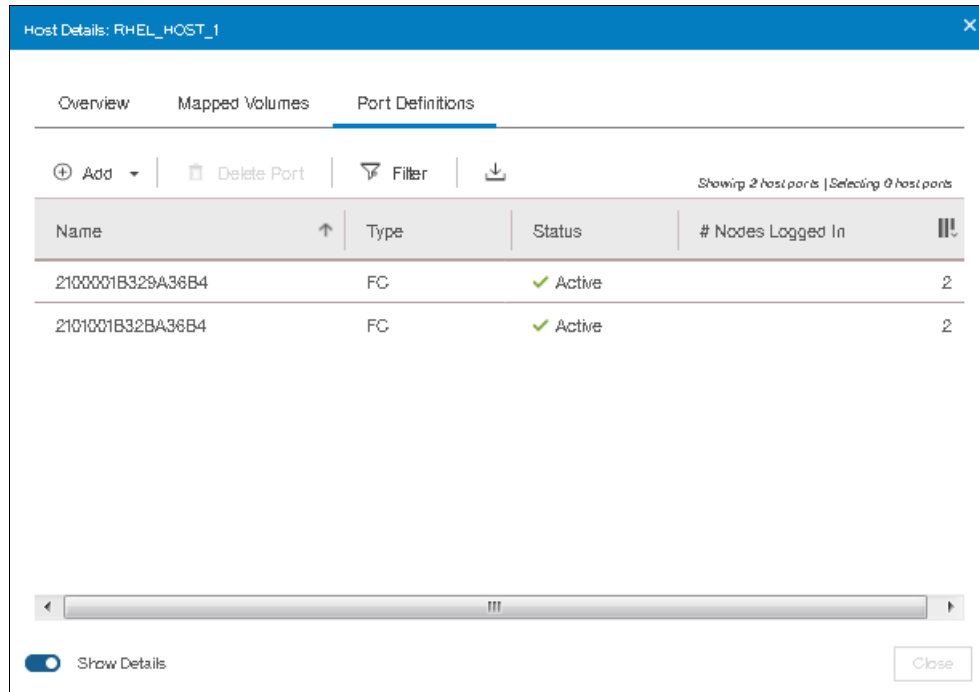


Figure 8-29 Host port definitions

Note: You can also add additional ports for the host using this window.

Click **Close** to close the Host Details panel.

8.2 Adding and deleting host ports

In this section we describe how to add and delete port(s) to and from a host definition. The examples shown here are for Fibre Channel ports, but the steps shown are equally applicable for SAS and iSCSI ports as well.

8.2.1 Adding host port

To add or delete host ports, go to the **Port Definitions** tab as shown in “Port Definitions” on page 366, then follow the steps listed here.

1. Click **Add** as shown in Figure 8-30 on page 368.

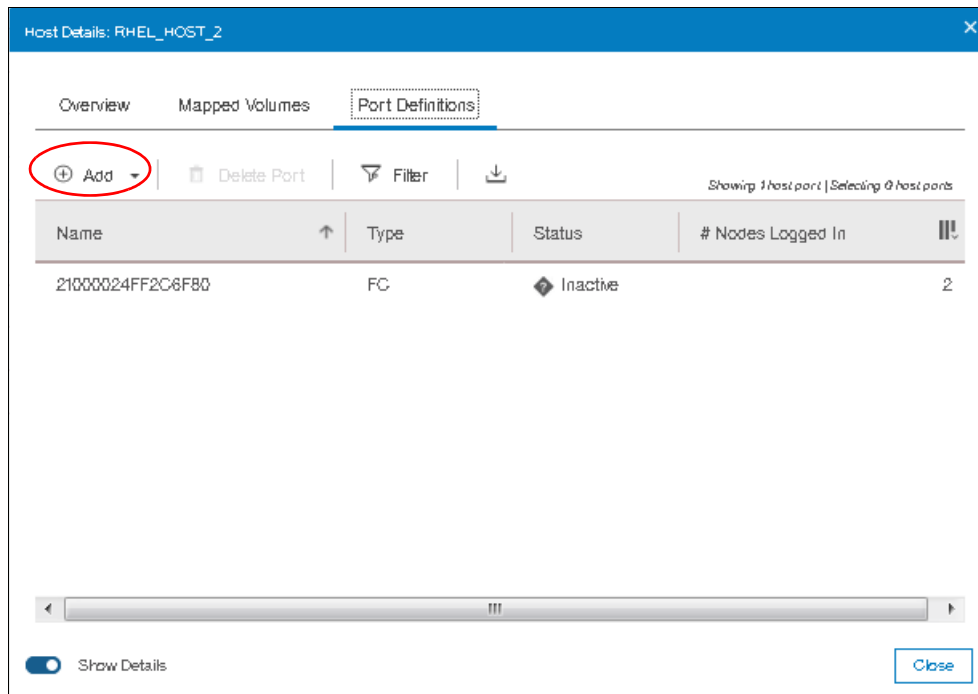


Figure 8-30 Add host port

2. Select the type of port that you want to add. In this example, we chose Fibre Channel port as shown in Figure 8-31.




Figure 8-31 Type of port to be added

3. A window opens with a drop-down for you to choose the desired WWPN to be added as shown in Figure 8-32 on page 369.

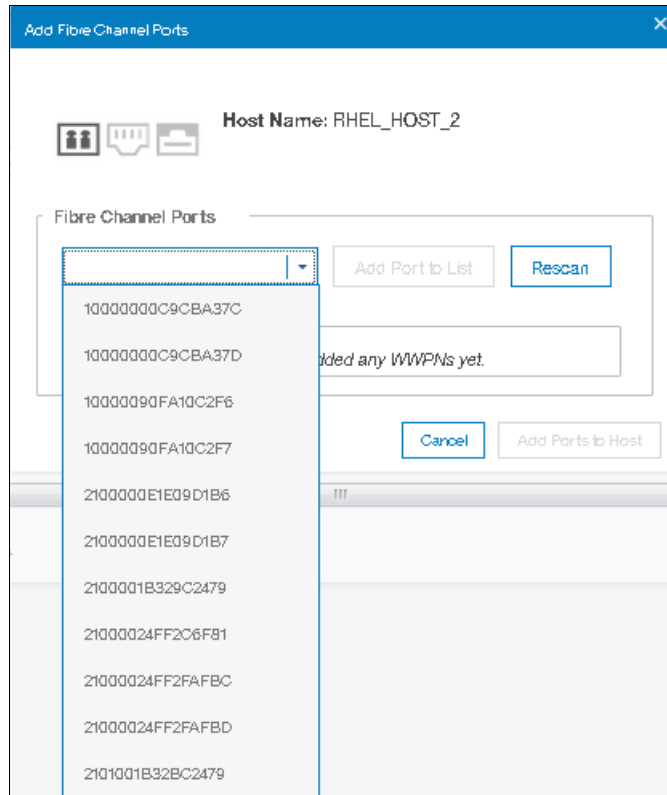


Figure 8-32 Drop-down of WWPNs

Note: If the WWPN does not show in the drop-down list, click Rescan and try again. If the port does not show up even after the rescanning, then check the zoning.

4. Select the desired WWPN and click **Add Port to List** as shown in Figure 8-33.

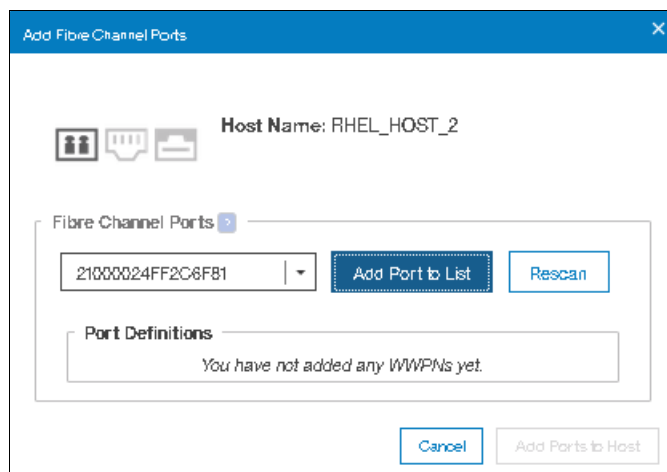


Figure 8-33 Add Port to List

5. The selected port will be shown under **Port Definitions** as shown in Figure 8-34 on page 370.




Figure 8-34 Port Definitions

Note: If the selected port is not the desired one, then you can click on the red cross to delete it from the selection.

6. Click **Add Ports to List**. A task completion window will be shown as in Figure 8-35.

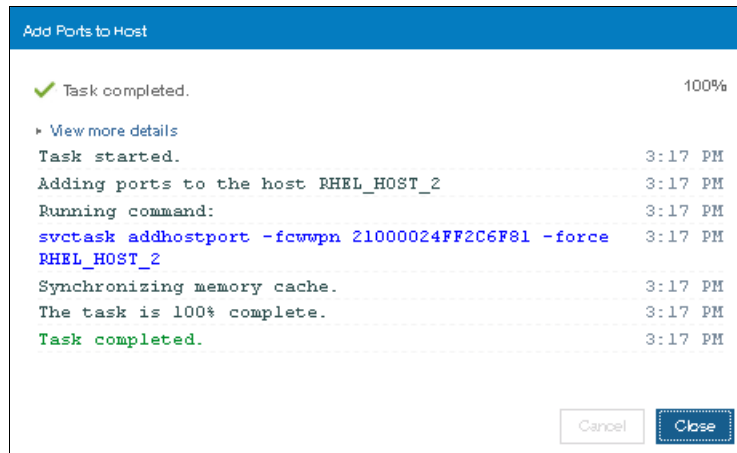


Figure 8-35 Add Ports to Host task completed

7. The Host Details window will now show the ports defined for the host, including the recently added one as shown in Figure 8-36 on page 371.

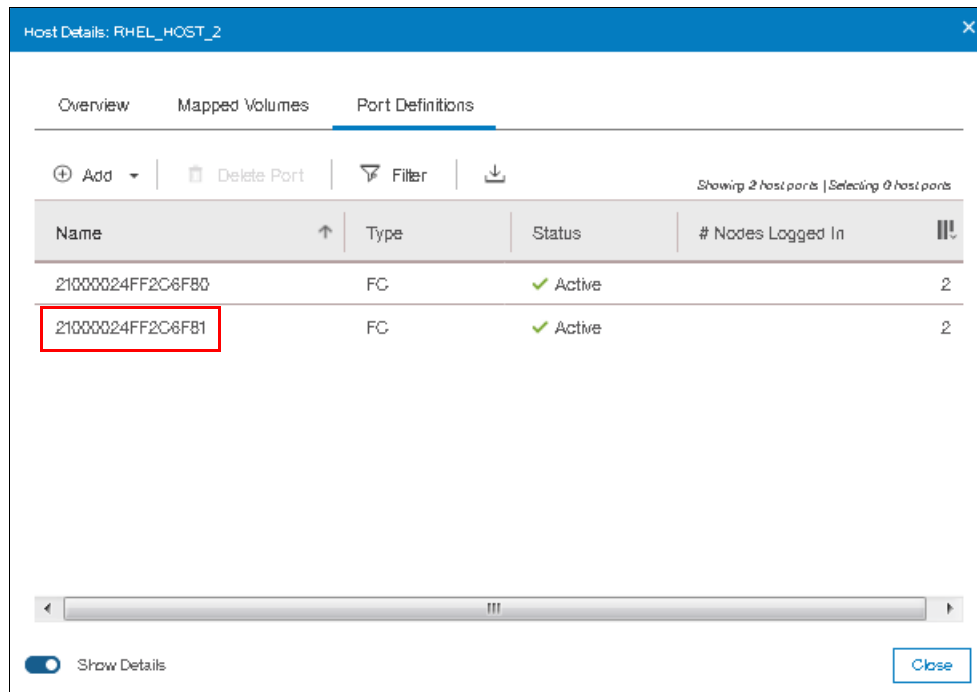


Figure 8-36 Port Definitions after adding a port

8.2.2 Deleting a host port

To add or delete host ports, go to the **Port Definitions** tab as shown in “Port Definitions” on page 366, then follow the steps listed here.

1. Select the port that you want to delete as shown in Figure 8-37.




Figure 8-37 Selecting the port to delete

2. Select **Delete Port** as shown in Figure 8-38.




Figure 8-38 Delete Port operation

3. Verify the number of ports and the port to be deleted as shown in Figure 8-39.




Figure 8-39 Delete port

4. Click **Delete**. A window indicating port deletion task completed will be shown as in Figure 8-40 on page 373.

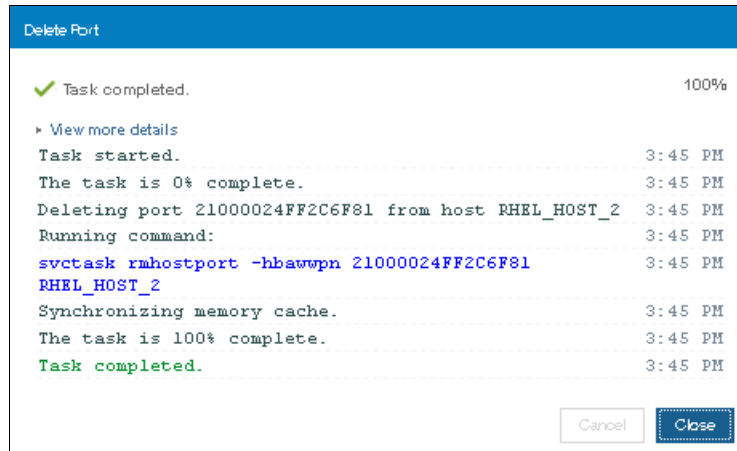


Figure 8-40 Port deletion task completed

5. A window will be shown with the current ports for the selected host as shown in Figure 8-41.

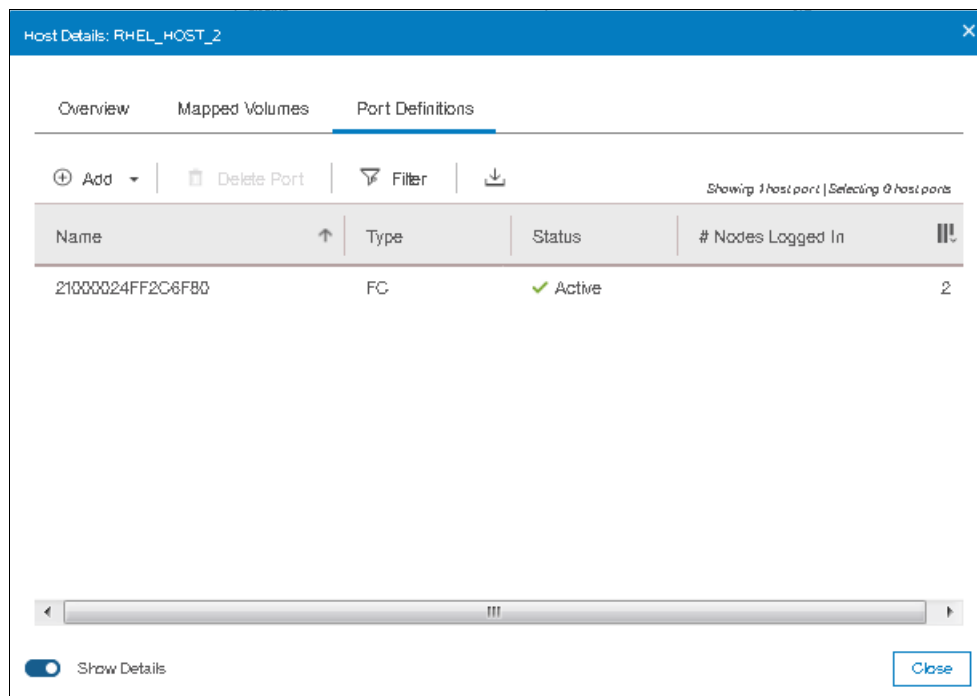


Figure 8-41 Host ports after deletion of a port.

8.3 Advanced volume administration

This section describes volume administration tasks, such as volume modification and the creation of volume copies. We assume that you completed Chapter 6, “Volume configuration” on page 269 and that you are familiar with volume creation and generic, thin-provisioned, mirrored, thin-mirrored and compressed volumes.

Figure 8-42 shows the following options, which are available in the Volumes menu for advanced feature administration:

- ▶ Volumes
- ▶ Volumes by Pool
- ▶ Volumes by Host

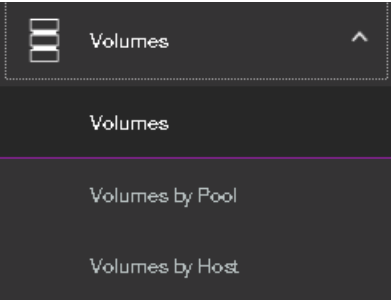


Figure 8-42 Volumes menu

8.3.1 Advanced volume functions

Click **Volumes** (as shown in Figure 8-42) and the Volumes panel opens, as shown in Figure 8-43.

Create Volumes Actions All Volumes Filter								
Showing 19 volumes (selecting 7 volume (70.00 GiB))								
Name	State	Synchronized	Pool	UID	Host Mappings	Capacity	Global Minor Change Vol...	IP
> Compresses_Volume	Online		Multi_Tier_Pool	6005076300A600000000000000000062	Yes	1000 GiB	No	
GMCV_SRC_VOL	Online		Enterprise_Pool	6005076300A600000000000000000063	No	1000 GiB	No	
GMCV_VHD_SRC_VOL	Online		Enterprise_Pool	6005076300A600000000000000000064	No	1500 GiB	Yes	
ITSQ_BASIC_VOL_FC	Online		Enterprise_Pool	6005076300A600000000000000000030	No	100 GiB	No	
ITSQ_RC_SRC_VOL_1	Online		Enterprise_Pool	6005076300A600000000000000000048	No	500 GiB	No	
ITSQ_RC_SRC_VOL_2	Online		Enterprise_Pool	6005076300A600000000000000000049	No	200 GiB	No	
ITSQ_RC_SRC_VOL_3	Online		Enterprise_Pool	6005076300A600000000000000000048	No	300 GiB	No	
ITSQ_Volume	Online		Multi_Tier_Pool	6005076300A600000000000000000060	No	1500 GiB	No	
MY_SRC_VOL	Online		Enterprise_Pool	6005076300A600000000000000000065	No	500 GiB	No	
RC_SRC_VOL_4	Online		Enterprise_Pool	6005076300A600000000000000000005E	No	200 GiB	No	
SRC_VOL_1	Online		Enterprise_Pool	6005076300A600000000000000000005C	No	300 GiB	No	
Thin_Provision_Volume	Online		Multi_Tier_Pool	6005076300A600000000000000000006	No	1000 GiB	No	
vol60	Online		Enterprise_Pool	6005076300A600000000000000000059	No	500 GiB	No	
vol61	Online		Enterprise_Pool	6005076300A600000000000000000005A	No	200 GiB	No	
vol62	Online		Enterprise_Pool	6005076300A60000000000000000005B	No	300 GiB	No	
vol63	Online		Enterprise_Pool	6005076300A60000000000000000005D	No	300 GiB	No	
vol64	Online		Enterprise_Pool	6005076300A60000000000000000005F	No	200 GiB	Yes	
vol65	Online		Enterprise_Pool	6005076300A600000000000000000066	No	500 GiB	Yes	

Figure 8-43 Volumes panel

This panel lists all configured volumes on the system and provides the following information:

- ▶ Name: Shows the name of the volume. If a twisty sign (>) appears before the name, two copies of this volume exist. Click the twisty sign (>) to expand the view and list the volume copies, as shown in Figure 8-43.
- ▶ State: Provides the status information about the volume, which can be online, offline, or degraded.
- ▶ Synchronized: For mirrored volumes, whether the copies are synchronized or not.
- ▶ Pool: Shows in which storage pool the volume is stored. The primary copy, which is marked with an asterisk (*), is shown unless you expand the volume copies.
- ▶ UID: The volume unique identifier.
- ▶ Host mappings: Shows whether a volume has host mapping: Yes when host mapping exists and No when no hosting mappings exist.

- **Capacity:** The disk capacity that is presented to the host. If a blue volume is listed before the capacity, this volume is a thin-provisioned volume. Therefore, the listed capacity is the virtual capacity, which might be larger than the real capacity on the system.
- **Global Mirror Change Volume:** Indicates whether a volume is a change volume for a Global Mirror relationship or not.

Tip: Right-click anywhere in the blue title bar to customize the volume attributes that are displayed. You might want to add useful information, such as the caching I/O group and the real capacity.

To create a volume, click **Create Volumes** and complete the steps as described in Chapter 6, “Volume configuration” on page 269.

Right-clicking or selecting a volume and opening the Actions menu shows the available actions for a volume, as shown in Figure 8-44 on page 376.

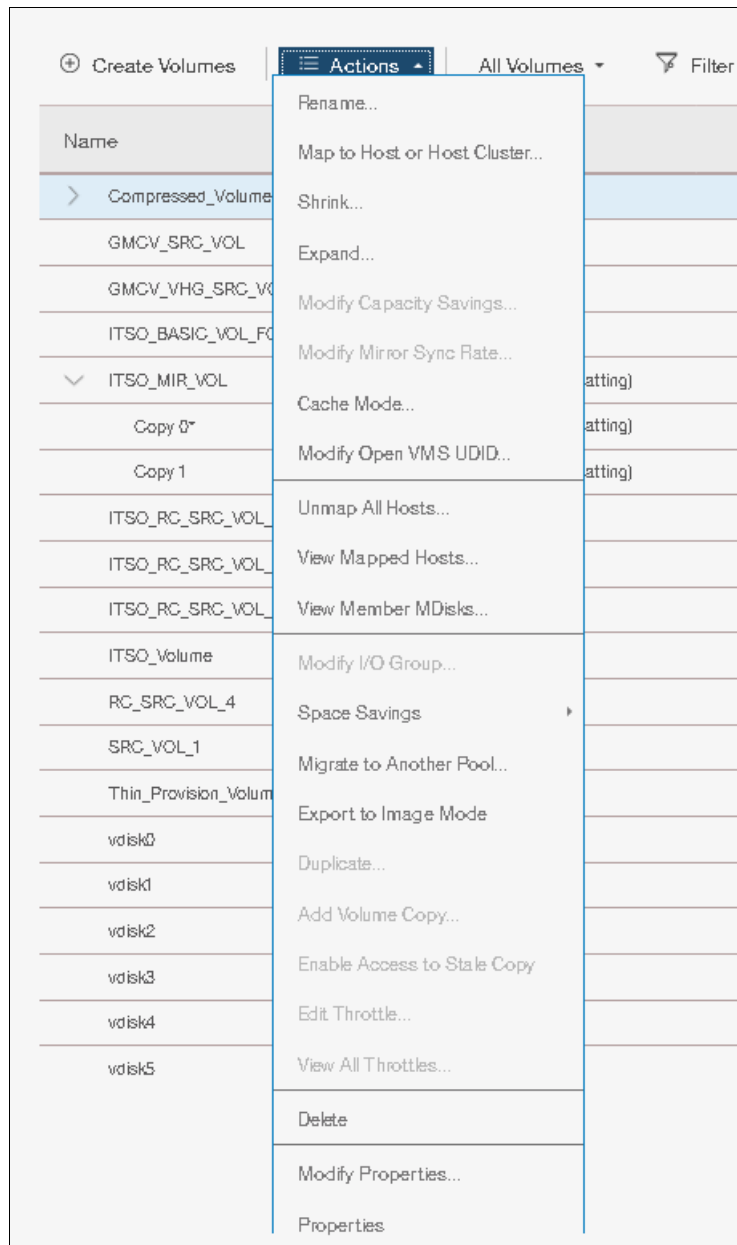


Figure 8-44 Actions menu for a volume

Depending on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 model capabilities and volume types, the following volume actions are available:

- ▶ Rename (8.3.4, “Renaming a volume” on page 378)
- ▶ Map to Host or Host Cluster (8.3.2, “Other actions are available for copies of volumes. For more information, see 8.5, “Advanced volume copy functions” on page 390. Unmapping volumes from all hosts” on page 377)
- ▶ Shrink (8.3.5, “Shrinking a volume” on page 378)
- ▶ Expand (8.3.6, “Expanding a volume” on page 380)
- ▶ Modify capacity savings (Choose between none, Thin Provisioning, and Compression.)
- ▶ Modify mirror synchronization rate (Set the synchronization rate value. For more information, see 8.4, “Volume properties and volume copy properties” on page 386.)

- ▶ Cache mode (Choose between Enabled, Read Only, and Disabled.)
- ▶ Modify open VMS unit device identifier (UDID)
- ▶ Unmap all hosts (8.3.2, “Other actions are available for copies of volumes. For more information, see 8.5, “Advanced volume copy functions” on page 390. Unmapping volumes from all hosts” on page 377)
- ▶ View mapped hosts (8.3.3, “Viewing which host is mapped to a volume” on page 378)
- ▶ Modify I/O group (only applicable to multiple I/O group systems)
- ▶ Space Savings (only for compressed volumes)
- ▶ Migrate to another pool (8.3.7, “Migrating a volume to another storage pool” on page 380)
- ▶ Export to image mode (8.3.8, “Exporting to an image mode volume” on page 381)
- ▶ Duplicate (8.3.10, “Duplicating a volume” on page 383)
- ▶ Add volume copy (8.3.11, “Adding a volume copy” on page 385)
- ▶ Enable access to stale copy (Available for IBM HyperSwap volumes if the copy is not up-to-date and inaccessible but contains consistent data from an earlier time)
- ▶ Edit Throttle
- ▶ View All Throttles
- ▶ Delete (8.3.9, “Deleting a volume” on page 383)
- ▶ Volume Copy Actions (see 8.5, “Advanced volume copy functions” on page 390)
- ▶ Modify Properties
- ▶ Properties (8.4, “Volume properties and volume copy properties” on page 386)

8.3.2 Other actions are available for copies of volumes. For more information, see 8.5, “Advanced volume copy functions” on page 390. **Unmapping volumes from all hosts**

To remove all host mappings from a volume, select **Unmap All Hosts** from the Actions menu. This action removes all host mappings, which means that no hosts can access this volume. Confirm the number of mappings to remove, and click **Unmap**, as shown in Figure 8-45.

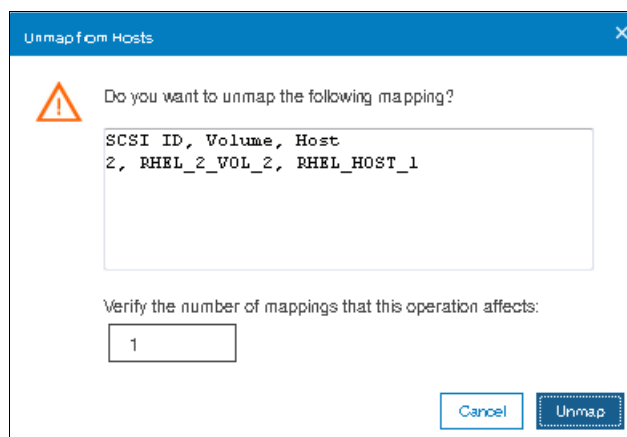


Figure 8-45 Unmapping from host or hosts

After the task completes, click **Close** to return to the Volumes panel.

Important: Ensure that the required procedures are run on the host OS before you run the unmapping procedure.

8.3.3 Viewing which host is mapped to a volume

To determine which host mappings are configured, highlight a volume and select **View Mapped Host** from the Actions menu. The Host Maps tab of the Volume Details panel opens, as shown in Figure 8-46. In this example, host RHEL_HOST_1 is mapped to the RHEL_HOST_1_BOOT_VOL volume.

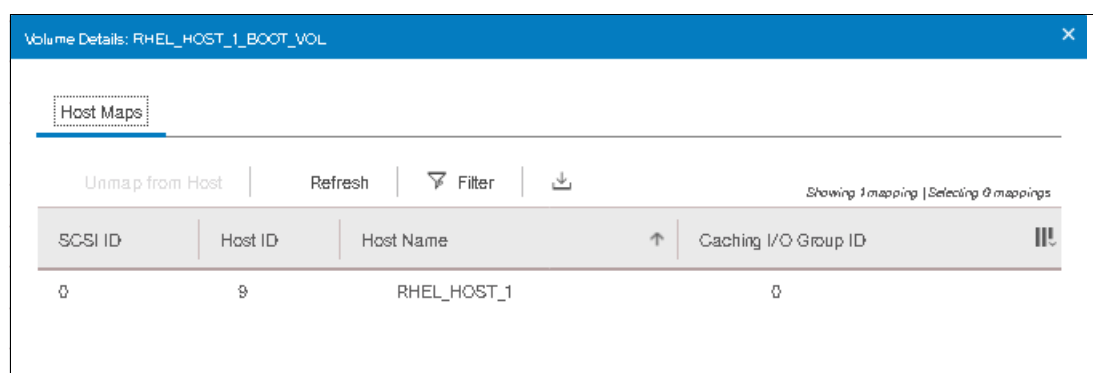


Figure 8-46 Volume to host mapping

To remove a mapping, highlight the host and click **Unmap from Host**. If several hosts are mapped to this volume (for example, in a cluster), only the selected host is removed.

8.3.4 Renaming a volume

To rename a volume, select **Rename** from the Actions menu. The Rename Volume window opens. Enter the new name, as shown in Figure 8-47.

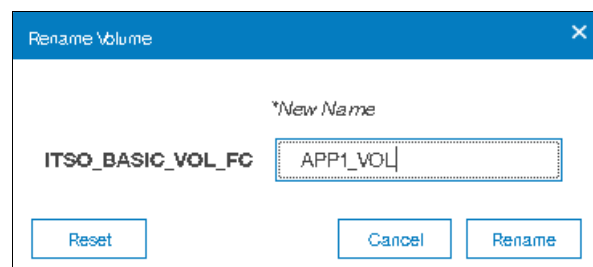


Figure 8-47 Renaming a volume

Click **Reset** to reset the name field to the original name of the volume. Click **Rename** to apply the changes. Click **Close** to close the panel.

8.3.5 Shrinking a volume

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can shrink volumes. Use this feature only if the host OS supports it. This capability reduces the capacity that is allocated to the particular volume by the amount that is specified. To shrink a volume, select **Shrink** from

the Actions menu. Enter the new volume size or the value by which to shrink the volume, as shown in Figure 8-48.

Important: Before you shrink a volume, ensure that the host OS supports this capability. If the OS does not support shrinking a volume, log disk errors and data corruption can occur.

Shrink Volume - ITSO_Volume

You selected to shrink the capacity of volume ITSO_Volume. The system arbitrarily reduces the capacity of the volume by removing one or more extents that are allocated to the volume. You cannot control which extents are removed and cannot ensure that unused capacity is removed.

If the volume contains data that is being used, do not attempt to shrink a volume under any circumstances without first backing up your data.

Current size: GiB

Shrink by: ▼

Final size: GiB

Figure 8-48 Shrink Volume panel

Click **Shrink** to start the process. Click **Close** when the task completes to return to the Volumes panel.

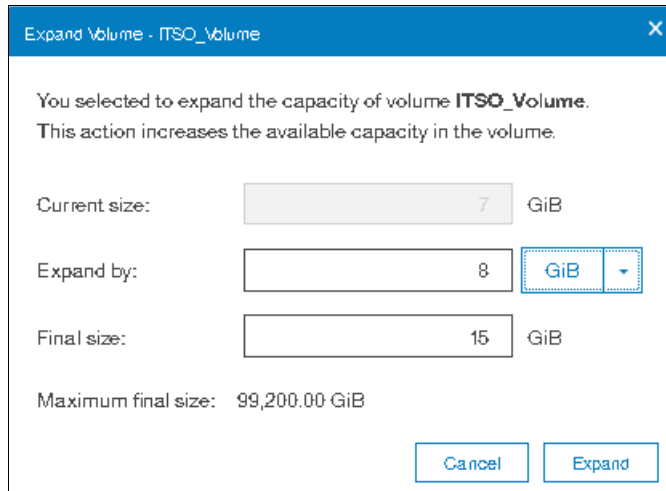
Run the required procedures on the host OS after the shrinking process.

Important: For volumes that contain more than one copy, you might receive a CMMVC6354E error. Check the Running tasks window and wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by increasing the Mirror Sync Rate value in the Actions menu. When the copy is synchronized, resubmit the shrink process.

Similar errors might occur if other tasks, for example, volume expand or format operations, are running on the volume. The solution is to wait until these operations finish, then restart the shrink process.

8.3.6 Expanding a volume

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 can expand volumes. Use this feature only if the host OS supports it. This capability increases the capacity that is allocated to the particular volume by the amount that is specified. To expand a volume, select **Expand** from the Actions menu. Enter the new volume size or enter the amount by which the volume needs to expand. Click **Expand**, as shown in Figure 8-49.



The dialog box titled "Expand Volume - ITSO_Volume" contains the following information:

- Message: "You selected to expand the capacity of volume **ITSO_Volume**. This action increases the available capacity in the volume."
- Current size: 7 GiB
- Expand by: 8 GiB (with a dropdown arrow)
- Final size: 15 GiB
- Maximum final size: 99,200.00 GiB
- Buttons: "Cancel" and "Expand"

Figure 8-49 Expand Volume panel

If the task completion dialog stays open, review the results of the operation and click **Close** to return to the Volumes panel.

Run the required procedures on the host OS to use the full available space.

Note: You can expand the capacity of volumes in Metro Mirror and Global Mirror relationships that are in **consistent_synchronized** state if those volumes are using thin-provisioned or compressed copies.

You cannot expand the following types of volumes:

- ▶ Volumes in HyperSwap relationships or in Global Mirror relationships that are operating in cycling mode
- ▶ Volumes in relationships where a change volume is configured
- ▶ Volumes that have a fully allocated copy

For further information on volume expansion refer to:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_expandvdisktask_22fift.html

8.3.7 Migrating a volume to another storage pool

To migrate a volume to another storage pool refer to Chapter 6, "Volume configuration" on page 269.

8.3.8 Exporting to an image mode volume

Image mode provides a direct block-for-block translation from a managed disk (MDisk) to a volume with no virtualization. An image mode MDisk is associated with one volume only. This feature can be used to export a volume to a non-virtualized disk and to remove the volume from storage virtualization.

Note: Among the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 families, this feature is available only on the Lenovo Storage V5030 storage system.

Select the volume that you want. From the Actions menu, choose **Export to Image Mode**, as shown in Figure 8-50.

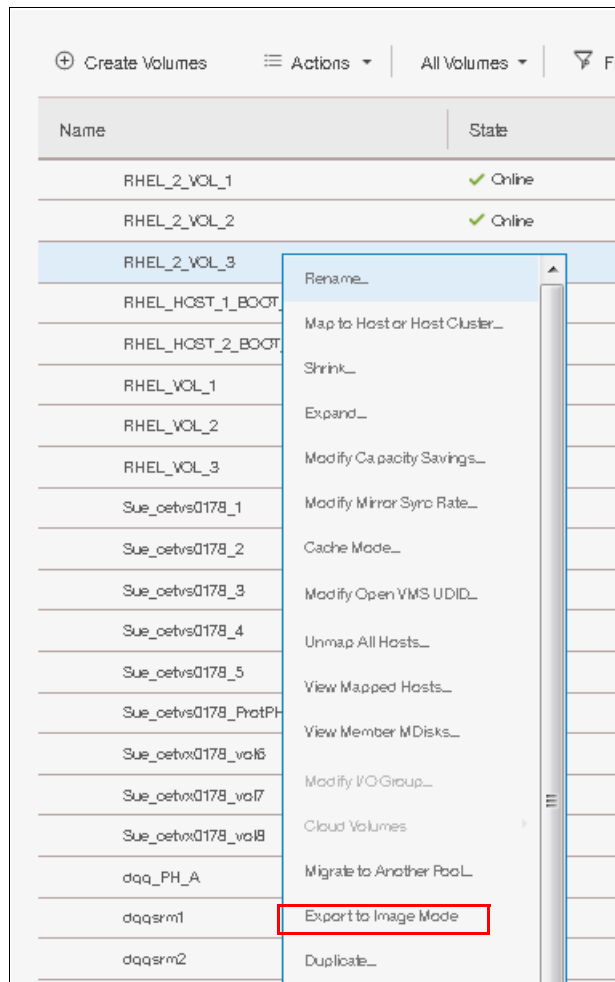


Figure 8-50 Exporting a volume to an image mode

The Export to Image Mode wizard opens and displays the available MDisk. Select the MDisk to which to export the volume, and click **Next**, as shown in Figure 8-51 on page 382.

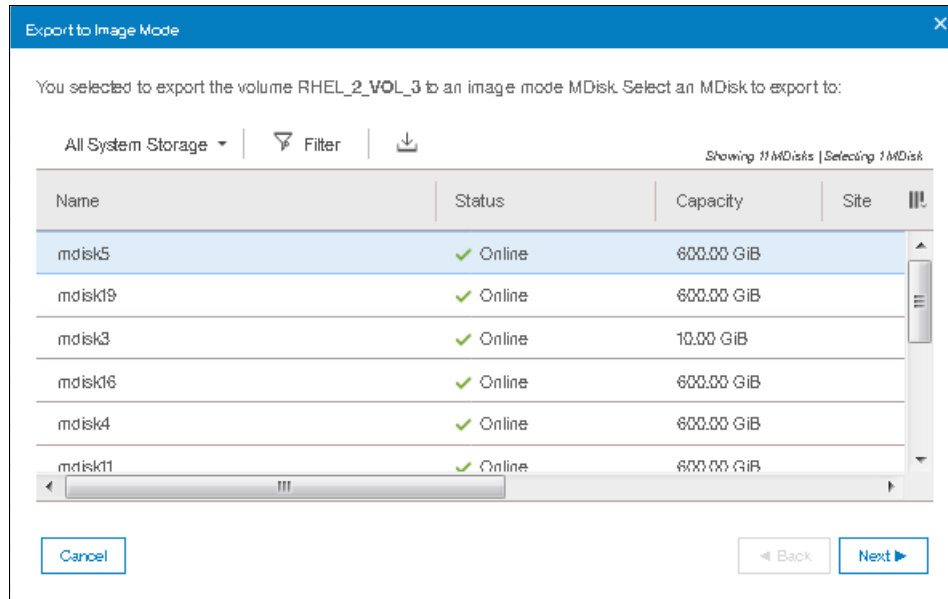


Figure 8-51 Selecting the MDisk to which to export the volume

Select a storage pool into which the image-mode volume is placed after the migration completes, as shown in Figure 8-52.

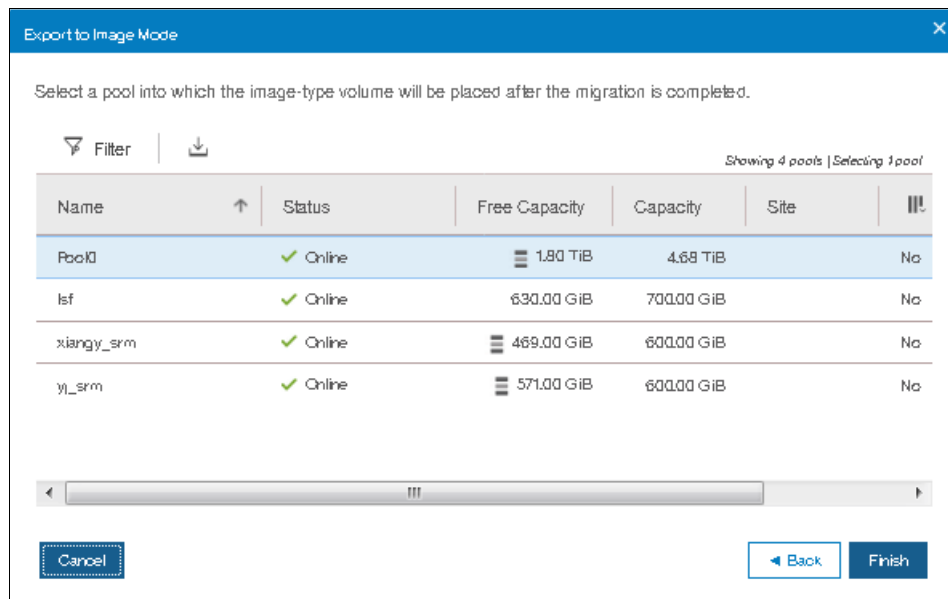


Figure 8-52 Select the storage pool

Click **Finish** to start the migration. After the task is complete, check the results and click **Close** to return to the Volumes panel.

Important: Use image mode to import or export existing data into or out of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. Migrate data from image mode MDisks to other storage pools to benefit from storage virtualization.

For more information about importing volumes from external storage, see Chapter 7, “Storage migration” on page 323 and Chapter 4, “Storage pools” on page 139.

8.3.9 Deleting a volume

To delete a volume, select **Delete** from the Actions menu. Confirm the number of volumes and select the check box if you want to force the deletion. Figure 8-53 shows the Delete Volume panel.

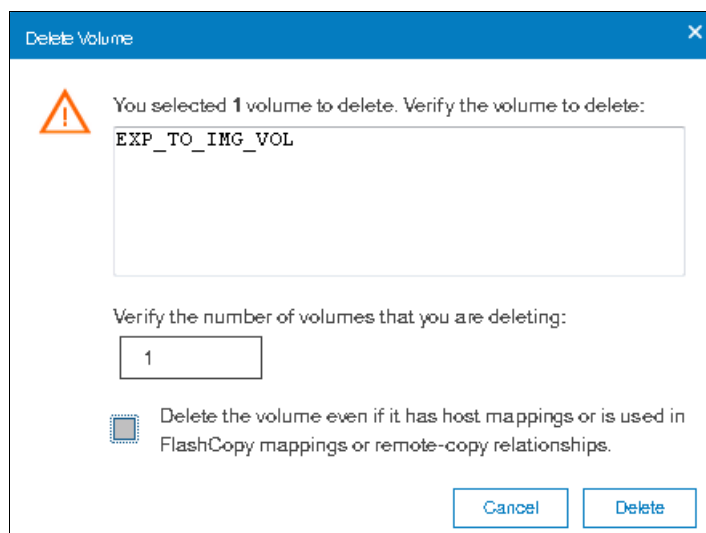


Figure 8-53 Delete Volume panel

Click **Delete** to remove the selected volume or volumes from the system. After the task completes, click **Close** to return to the Volumes panel.

Important: You must force the deletion if the volume has host mappings or if the volume is used in FlashCopy mappings. To be cautious, always ensure that the volume has no association before you delete it.

8.3.10 Duplicating a volume

You can create a new volume by using the same presets and parameters as an existing volume. These parameters are shown:

- ▶ Volume preset (generic, thin-provision, and compressed)
- ▶ Volume size
- ▶ Storage pool
- ▶ Access and caching I/O group
- ▶ Caching mode
- ▶ Easy Tier status
- ▶ Virtualization type

Important: *Duplicating a volume does not duplicate the volume data.* The duplicating task creates a volume with the same preset and volume parameters as the source volume. *Duplicating mirrored and image-mode volumes is not supported.*

To duplicate a volume, select **Duplicate** from the Actions menu (Figure 8-54).

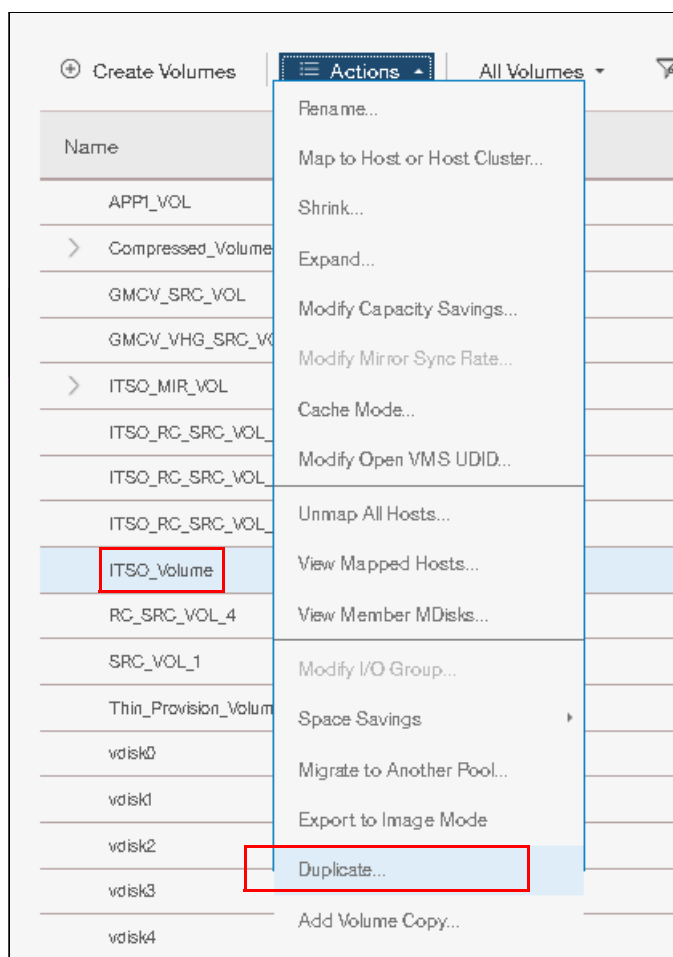


Figure 8-54 Duplicate volume option

The Duplicate Volume window, which is shown in Figure 8-55, can be used to change the name of the new volume. By default, a sequence integer is appended to the name of the volume to be duplicated.




Figure 8-55 Duplicate Volume

Click **Duplicate** to start the process. If the task completion dialog stays on the window, review the process results and click **Close**.

8.3.11 Adding a volume copy

If a volume consists of only one copy, you can add a second mirrored copy of the volume. This second copy can be generic or thin-provisioned.

You can also use this method to migrate data across storage pools with different extent sizes.

To add a second copy, select the volume and click **Actions** → **Add Volume Copy**, as shown in Figure 8-56.

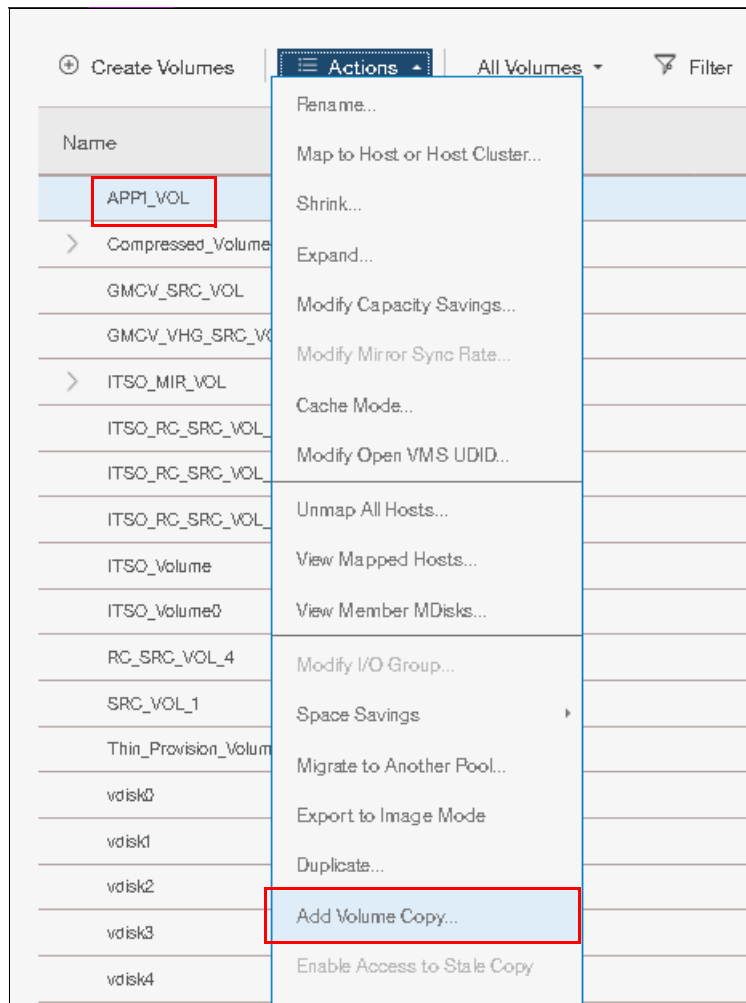


Figure 8-56 Add Volume Copy option

Select the storage pool in which to create the copy. Select capacity savings, between None, Thin-provisioned or Compressed. Click **Add**, as shown in Figure 8-57 on page 386.

Add Volume Copy

Create preset volumes with copies in multiple pools but at a single site.

Pool:

Copy 1: Enterprise_Pool Total 552.00 GiB

Copy 2: Multi_Tier_Pool Total 4.24 TiB

Capacity savings: None

Summary

- 1 volume
- 2 mirrored copies
- 1 copy in pool Enterprise_Pool
- 1 copy in pool Multi_Tier_Pool

Buttons: Cancel, Add

Figure 8-57 Add Volume Copy: Select a storage pool

The copy is created after you click **Add** and data starts to synchronize as a background task. If the task completion dialog stays on the window, review the results and click **Close**.

Now, the volume that is named APP1_VOL has two volume copies, which are stored in two separate storage pools (Figure 8-58).

<div> <div>⊕ Create Volumes</div> <div>≡ Actions ▾</div> <div>All Volumes ▾</div> <div>🔍 Filter</div> </div>				
Name	↑	State	Synchronized	Pool
APP1_VOL		✓ Online		Enterprise_Pool
Copy 0*		✓ Online	Yes	Enterprise_Pool
Copy 1		✓ Online	No	Multi_Tier_Pool
> Compressed_Volume		✓ Online		Multi_Tier_Pool
GMCV_SRC_VOL		✓ Online		Enterprise_Pool

Figure 8-58 Volume copies

8.4 Volume properties and volume copy properties

This section provides an overview of all available information that relates to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 volumes.

To open the advanced view of a volume, select **Properties** from the Actions menu. Click **View more details** to show the full list of volume properties, as shown in Figure 8-59 on page 387.

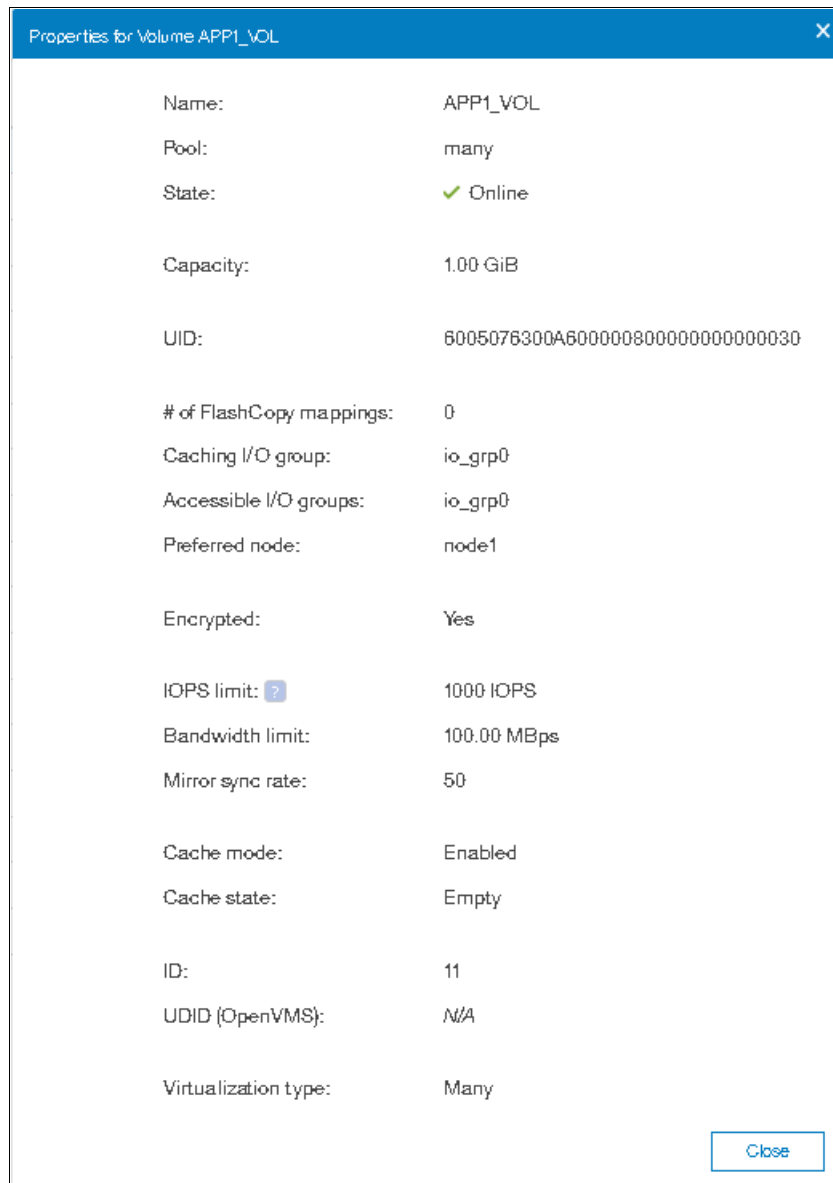


Figure 8-59 Volume details overview

The following details are available:

- ▶ **Name:** Shows the name of the volume.
- ▶ **Pool:** Gives you information about the storage pool in which the volume exists. The value “many” indicates that the volume has multiple copies, which can be in different pools.
- ▶ **State:** Gives status information about the volume, which can be online, offline, or degraded.
- ▶ **Capacity:** Shows the capacity of the volume. If the volume is thin-provisioned, this number is the virtual capacity.
- ▶ **UID:** The volume unique identifier.
- ▶ **Number of FlashCopy mappings:** The number of existing FlashCopy relationships. For more information, see Chapter 10, “Copy services” on page 451.
- ▶ **Caching I/O group:** Specifies the volume caching I/O group.

- ▶ Accessible I/O groups: Shows the I/O group that the host can use to access the volume.
- ▶ Preferred node: Specifies the ID of the preferred node for the volume.
- ▶ Encrypted: Shows whether the volume is encrypted.
- ▶ I/O throttling: You can set a limit on the number of I/O operations that are accepted for a volume. The limit is set in terms of I/O operations per second (IOPS) or bandwidth. For more details on I/O throttling, 6.10, “I/O throttling” on page 317.
- ▶ Mirror sync rate: After creation, or if a volume copy is offline, the mirror sync rate weights the synchronization process. Volumes with a high sync rate (100%) complete the synchronization faster than volumes with a lower priority. By default, the rate is set to 50% for all volumes.
- ▶ Cache mode: Shows whether the cache is enabled or disabled for this volume.
- ▶ Cache state: Indicates whether open I/O requests are in the cache that is not destaged to the disks.
- ▶ ID: Shows the identifier (ID) of the volume.
- ▶ UDID (OpenVMS): The unit device identifiers (UDIDs) are used by OpenVMS hosts to access the volume.
- ▶ Virtualization type: Specifies the virtualization type of the volume. The value can be striped, seq, image, or many. The value “many” indicates that the volume has multiple copies, which can have different virtualization types.

To open the advanced view of a volume copy, select the copy of the volume that you want and click **Actions** → **Properties**, as shown in Figure 8-60.

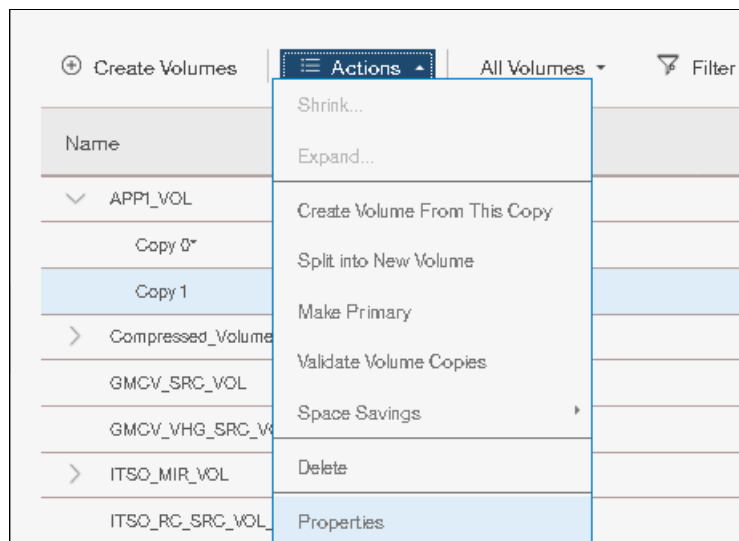


Figure 8-60 Volume copy properties

The Properties panel opens. Click **View more details** to show the full list of the volume copy properties, which is shown in Figure 8-61 on page 389.

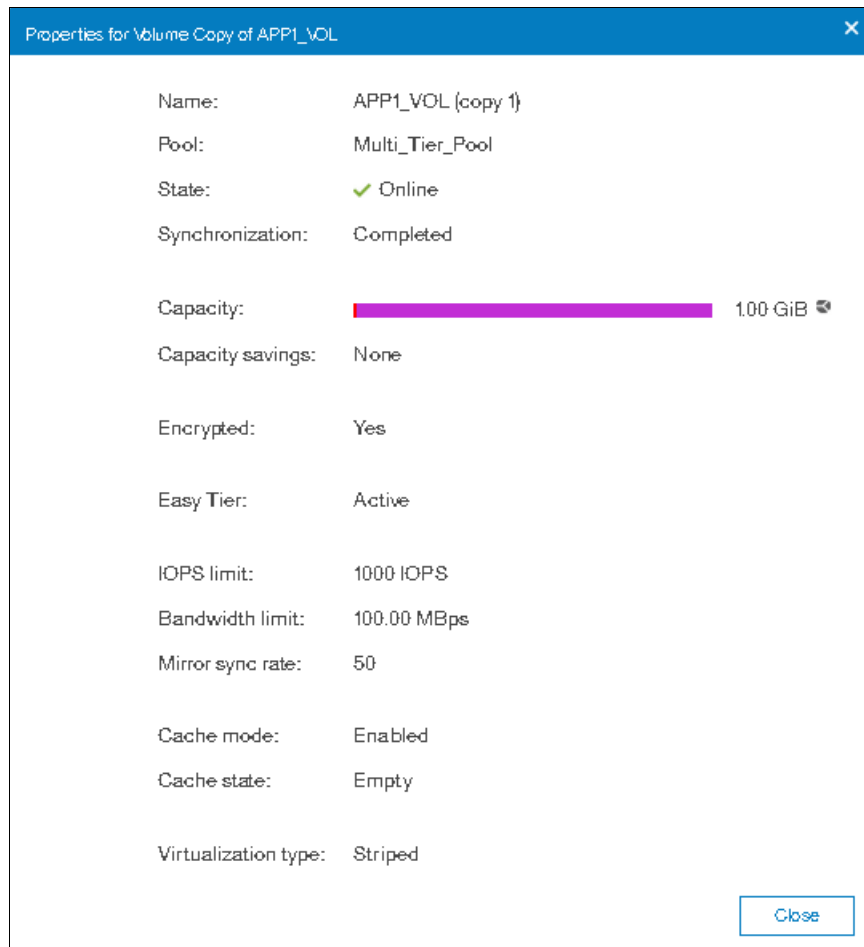


Figure 8-61 Volume copy properties

The following additional volume copy details are available:

- **Capacity:** Shows the allocated (used) and the virtual (real) capacity. You can hover your mouse pointer over the pie chart to display information for additional tiers.
- **Capacity savings:** Shows the method of capacity savings for the volume copy. It can be None for a generic copy, Compressed, or Thin-provisioned.
- **Easy Tier:** Indicates the Easy Tier state.

Mirror synchronization rate: Setting the Mirror Sync Rate to 0% disables synchronization.

To change the volume copy properties for the thin-provisioned or compressed volume copy type, select a volume copy from the Volumes panel, and click **Actions** → **Modify Properties**. Use the Modify Properties dialog, which is shown in Figure 8-62, to customize the thin-provisioning values: Warning Threshold and Enable Autoexpand.

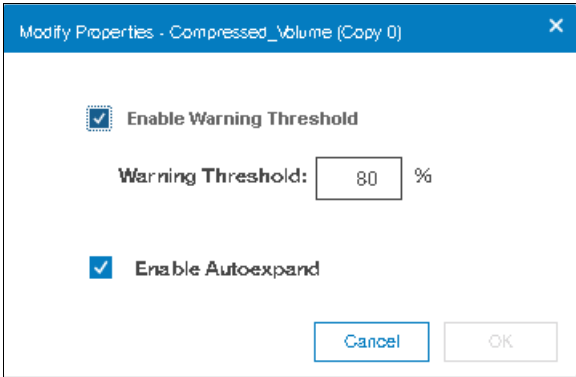


Figure 8-62 Modify volume copy properties

Modify the values if needed and click **OK**. After the task completes, check the results of the operation and click **Close** to return to the Volumes panel.

8.5 Advanced volume copy functions

In 8.3.1, “Advanced volume functions” on page 374, we described all of the available actions at a volume level and how to create a second volume copy. In this section, we focus on volumes that consist of two volume copies and how to apply the concept of two copies for business continuity and data migration.

Select the desired volume having multiple copies. Click on the twisty (>) to show all the copies. Then, select a volume copy. Open the Actions menu to display the following volume copy actions (Figure 8-63).

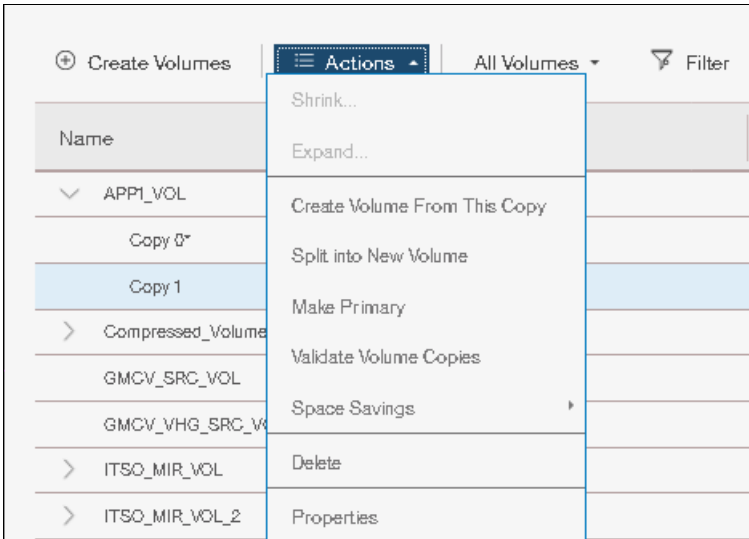


Figure 8-63 Volume Copy Actions menu

The following functions are available on the Actions menu:

- ▶ Create a volume from this copy
- ▶ Split into a new volume (8.5.2, “Splitting into a new volume” on page 393)
- ▶ Make primary (8.5.1, “Volume copy: Make Primary” on page 391)
- ▶ Validate volume copies (8.5.3, “Validate Volume Copies option” on page 394)
- ▶ Delete (8.5.4, “Delete volume copy option” on page 397)
- ▶ Modify properties (8.4, “Volume properties and volume copy properties” on page 386)
- ▶ Properties (8.4, “Volume properties and volume copy properties” on page 386)

8.5.1 Volume copy: Make Primary

When you look at the volume copies, you can see that one of the copies shows an asterisk (*) next to its name, as shown in Figure 8-64.

Name	State	Synchronized	Pool
APP1_VOL	✓ Online		Enterprise_Pool
Copy 0*	✓ Online	Yes	Enterprise_Pool
Copy 1	✓ Online	Yes	Multi_Tier_Pool

Figure 8-64 Volume copy names

Each volume has a primary and a secondary copy, and the asterisk indicates the primary copy. The two copies are always synchronized, which means that all writes are destaged to both copies, but all reads are always performed from the primary copy. The maximum configurable number of copies per volume is two. The roles of the copies can be changed.

To accomplish this task, select the secondary copy. Then, click **Actions** → **Make Primary**. Usually, it is a preferred practice to place the volume copies on storage pools with similar performance because the write performance is constrained if one copy is placed on a lower-performance pool.

Figure 8-65 shows the secondary copy Actions menu.

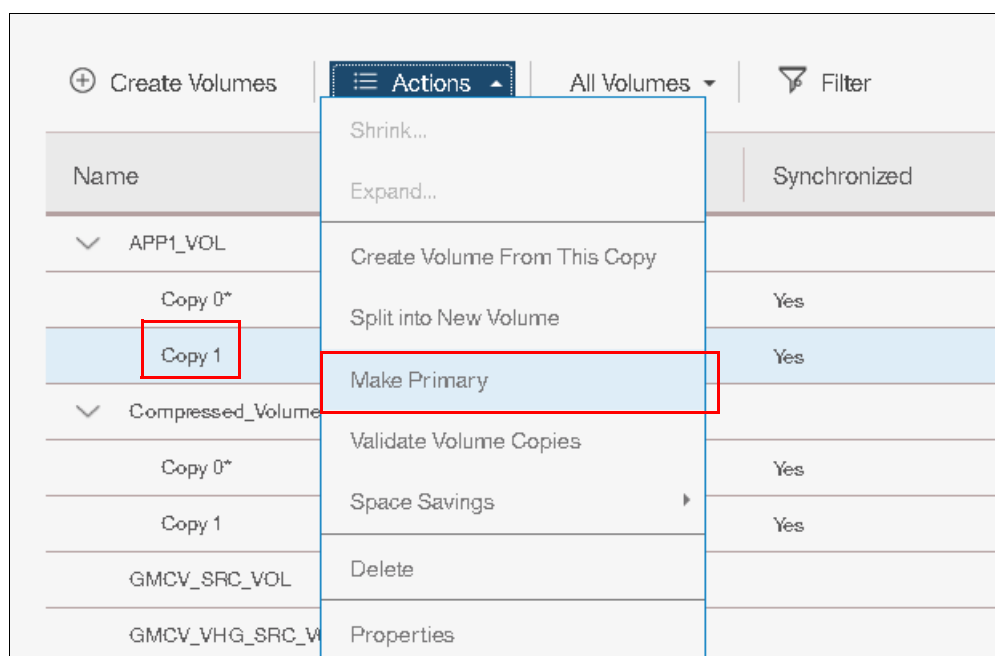


Figure 8-65 Make primary option

If you require high read performance, you can place the primary copy in a solid-state drive (SSD) pool or an externally virtualized Flash System and then place the secondary copy in a normal disk storage pool. This action maximizes the read performance of the volume and guarantees that a synchronized second copy is in your less expensive disk pool. You can migrate online copies between storage pools. For more information about how to select the copy that you want to migrate, see 8.3.7, “Migrating a volume to another storage pool” on page 380.

Click **Make Primary** and the role of the Copy 1 is changed to Primary, as shown in Figure 8-66.

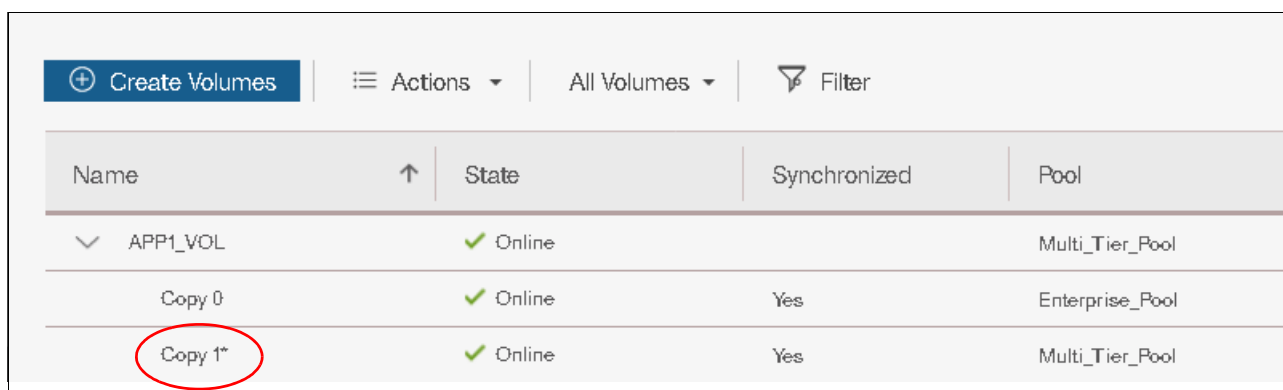


Figure 8-66 Change the primary volume copy

If the task completion dialog stays on the window, check the process output and click **Close**.

The volume copy feature is also a powerful option for migrating volumes, as described in 8.5.5, “Migrating volumes by using the volume copy features” on page 398.

8.5.2 Splitting into a new volume

If the two-volume copies are synchronized, you can split one of the copies to a new volume and map this volume to another host. From a storage point of view, this procedure can be performed online, which means that you can split one copy from the volume and create a copy from the remaining volume without affecting the host. However, if you want to use the split copy for testing or backup, you must ensure that the data inside the volume is consistent. Therefore, the data must be flushed to storage to make the copies consistent.

For more information about flushing the data, see your operating system documentation. The easiest way to flush the data is to shut down the hosts or application before a copy is split.

In our example, volume APP1_VOL has two copies: Copy 0 is primary and Copy 1 is secondary. To split a copy, click **Split into New Volume** (Figure 8-67) on any copy and the remaining secondary copy automatically becomes the primary for the source volume.

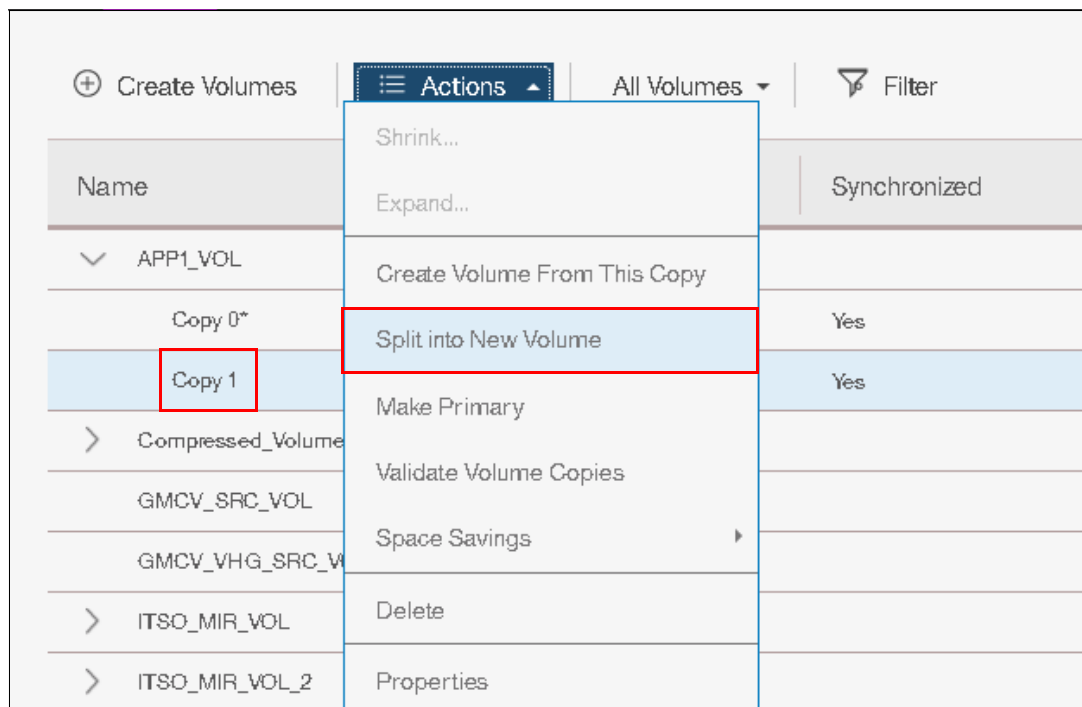


Figure 8-67 Split into New Volume option

Figure 8-68 shows the Split Volume Copy panel to specify a name for the new volume.

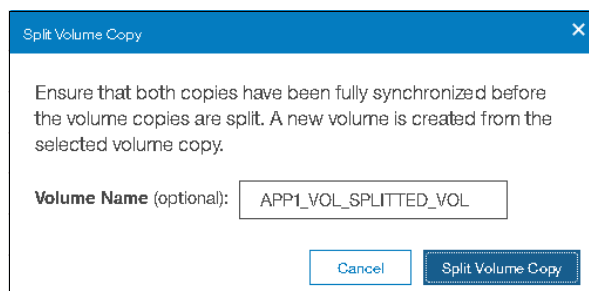
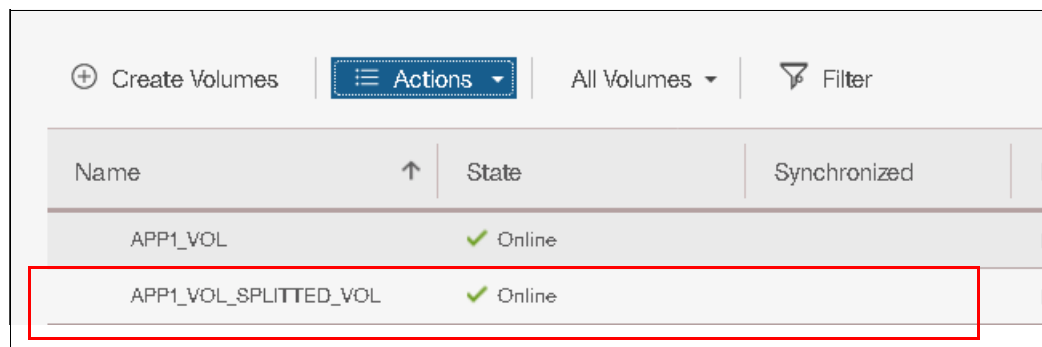


Figure 8-68 Split Volume Copy panel

If the task completion dialog stays on the window after the task completes, review the results and click **Close** to return to the Volumes panel.

As shown in Figure 8-69, the copy appears as a new volume that is named APP1_VOL_SPLITTED_VOL (as specified during the split process). The new volume can be mapped to a host.



⊕ Create Volumes ⋮ Actions ▾ All Volumes ▾ 🔍 Filter			
Name	↑	State	Synchronized
APP1_VOL		✓ Online	
APP1_VOL_SPLITTED_VOL		✓ Online	

Figure 8-69 Volumes: New volume from the split copy operation

Important: If you receive error message code CMMVC6357E while you are splitting a volume copy, click the **Running Tasks** icon to view the synchronization status. Then, wait for the copy to synchronize and repeat the splitting process.

8.5.3 Validate Volume Copies option

You can check whether the volume copies are identical and process the differences between them if they are not.

To validate the copies of a mirrored volume, complete the following steps:

1. From the Actions menu, select **Validate Volume Copies**, as shown in Figure 8-70 on page 395.

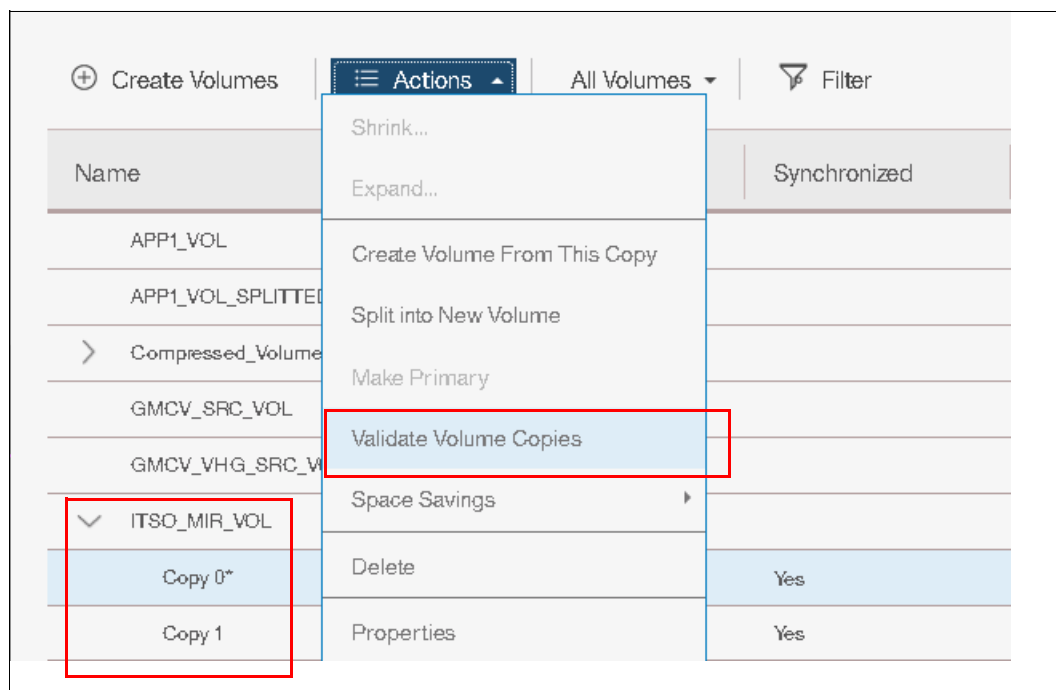


Figure 8-70 Actions menu: Validate Volume Copies

2. The Validate Volume Copies dialog opens, as shown in Figure 8-71.

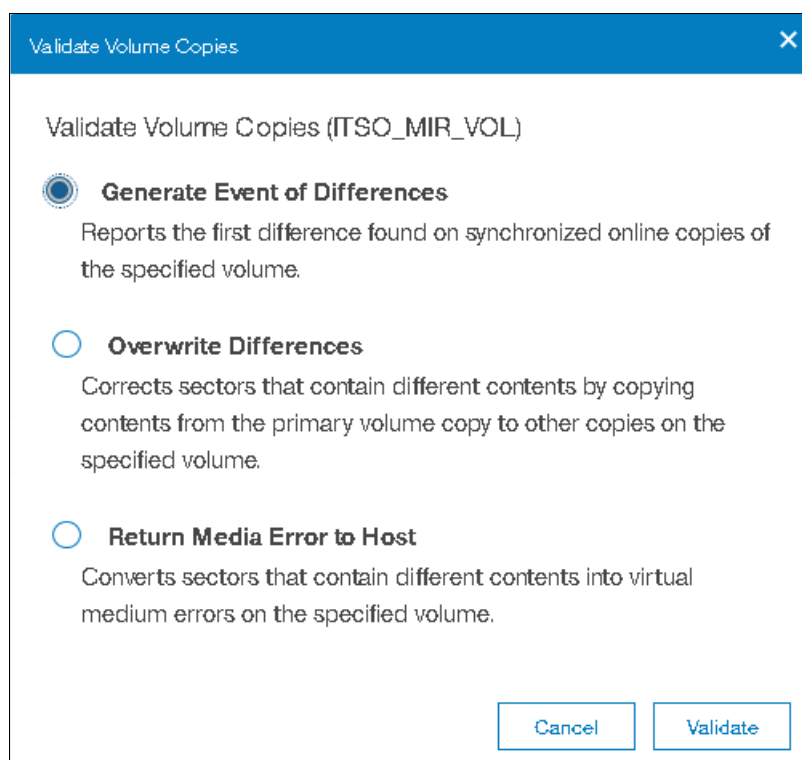


Figure 8-71 Validate Volume Copies panel

The following options are available:

- Generate Event of Differences

Use this option if you want to verify that the mirrored volume copies are identical. If any difference is identified, the command stops and logs an error that includes the logical block address (LBA) and the length of the first difference. You can use this option, starting at a different LBA each time, to count the number of differences on a volume.

- Overwrite Differences

Use this option to overwrite contents from the primary volume copy to the other volume copy. The command corrects any differing sectors by copying the sectors from the primary copy to the copies, which are compared. Upon completion, the command process logs an event, which indicates the number of differences that were corrected. Use this option if you are sure that the primary volume copy data is correct or that your host applications can handle incorrect data.

- Return Media Error to Host

Use this option to convert sectors on all volume copies, which contain different contents, into virtual medium errors. Upon completion, the command logs an event, which indicates the number of differences that were found, the number that were converted into medium errors, and the number that were not converted. Use this option if you are unsure what the correct data is and you do not want an incorrect version of the data to be used.

3. Select which action to perform and click **Validate** to start the task.
4. The volume is now checked. If the task dialog stays open, review the task results and click **Close**.
5. Figure 8-72 shows the output when you select the volume copy **Generate Event of Differences** option.

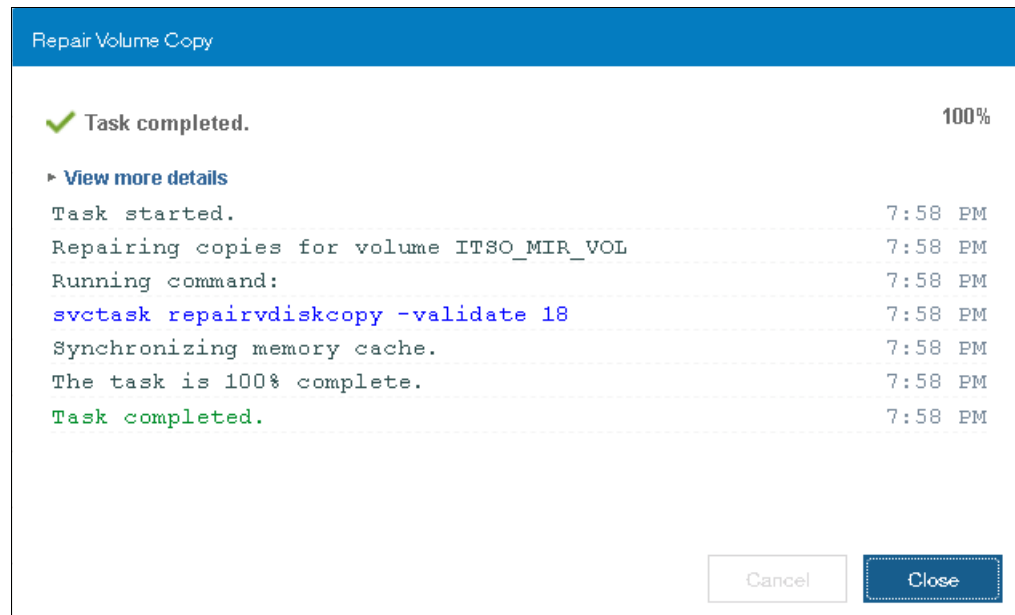


Figure 8-72 Volume copy validation output

The validation process runs as a background process and might take time, depending on the volume size. You can check the status in the Running Tasks window, as shown in Figure 8-73 on page 397.

Select a running task to see its progress.	Progress: Volume Repairs	
	<div>2 Volume Repairs</div>	
	Name	Progress
	ITSO_MIR_VOL, copy 0	8%
	ITSO_MIR_VOL, copy 1	8%

Figure 8-73 Validate Volume Copies: Running Tasks

8.5.4 Delete volume copy option

Click **Delete** (as shown in Figure 8-74) to delete a volume copy.

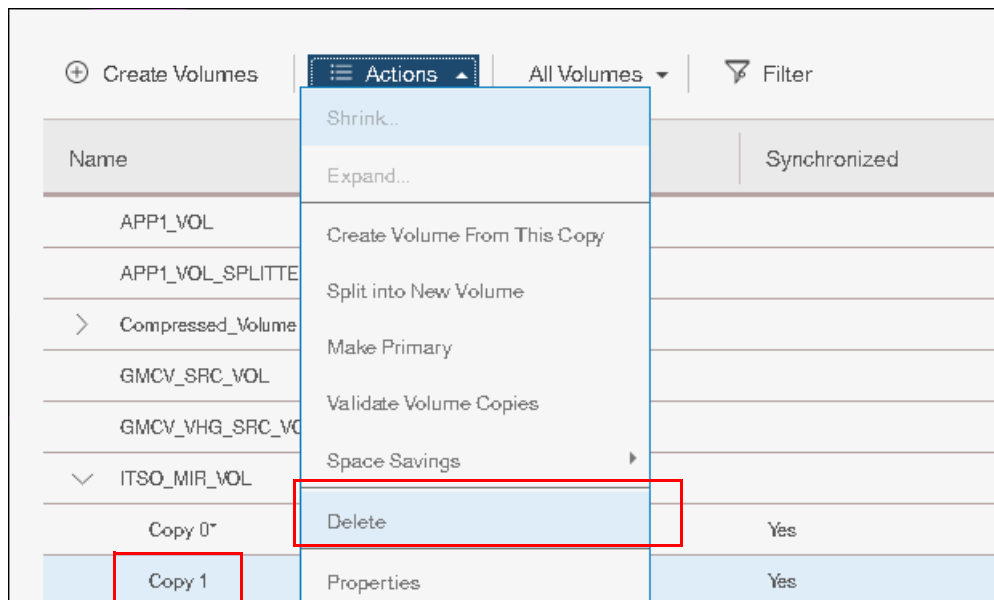


Figure 8-74 Actions menu: Delete a volume copy

Confirm the deletion process by clicking **Yes**. Figure 8-75 shows the copy deletion warning panel.

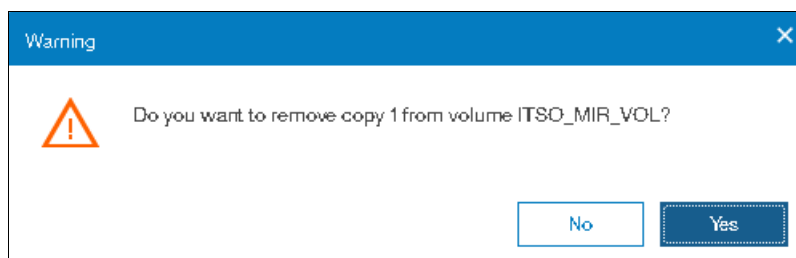


Figure 8-75 Delete a copy confirmation

If the task completion dialog is still open after the copy is deleted, review the results of the operation and click **Close** to return to the Volumes panel.

The copy is deleted, but the volume stays online by using the remaining copy (Figure 8-76).

ITSO_MIR_VOL	✓ Online	Enterprise_Pool
--------------	----------	-----------------

Figure 8-76 Volume remains online after the copy deletion

8.5.5 Migrating volumes by using the volume copy features

In the previous sections, we showed how to create, synchronize, split, and delete volume copies. A combination of these tasks can be used to migrate volumes to other storage pools.

The easiest way to migrate volume copies is to use the migration feature that is described in 8.3.7, “Migrating a volume to another storage pool” on page 380. By using this feature, one extent after another is migrated to the new storage pool. However, the use of volume copies provides another way to migrate volumes if the storage pool extent sizes differ.

To migrate a volume, complete the following steps:

1. Create a second copy of your volume in the target storage pool. For more information, see 8.3.11, “Adding a volume copy” on page 385.
2. Wait until the copies are synchronized.
3. Change the role of the copies and make the new copy the primary copy. For more information, see 8.5, “Advanced volume copy functions” on page 390.
4. Split or delete the old copy from the volume. For more information, see 8.5.2, “Splitting into a new volume” on page 393 or 8.5.4, “Delete volume copy option” on page 397.

This migration process requires more user interaction with the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI, but it offers benefits. For example, we look at migrating a volume from a tier 1 storage pool to a lower-performance tier 2 storage pool.

In step 1, you create the copy on the tier 2 pool, while all reads are still performed in the tier 1 pool to the primary copy. After the synchronization, all writes are destaged to both pools, but the reads are still only from the primary copy.

Because the copies are fully synchronized, you can switch their roles online (step 3), and analyze the performance of the new pool. After you test your lower performance pool, you can split or delete the old copy in tier 1 or switch back to tier 1 in seconds if the tier 2 storage pool did not meet your requirements.

8.6 Volumes by storage pool

To see the layout of volumes within pools, click **Volumes by Pool**, as shown in Figure 8-77.




Figure 8-77 Volumes by Pool menu

The Volumes by Pool panel opens, as shown in Figure 8-78.

Pool Filter

- 14 Volume copies
70.00 GiB Allocated / 700.00 GiB
- 113 Volume copies
2.55 TiB Allocated / 4.68 TiB
- 20 Volume copies
131.00 GiB Allocated / 600.00 GiB
- 11 Volume copies
230.00 GiB Allocated / 600.00 GiB

lzf
Online
2 MDisks, 14 Volume copies
Easy Tier Balanced

Volume Allocation

Name	State	Synchronized	Volume Group	UUID	Host Mappings	Capacity	IP
lzf0	✓ Online			600507669C948258000000000000007A	Yes	5.00 GiB	
lzf1	✓ Online			600507669C948258000000000000007B	Yes	5.00 GiB	
lzf2	✓ Online			600507669C948258000000000000007C	Yes	5.00 GiB	
lzf3	✓ Online			600507669C948258000000000000007D	Yes	5.00 GiB	
lzf4	✓ Online			600507669C948258000000000000007E	Yes	5.00 GiB	
lzf5	✓ Online			600507669C948258000000000000007F	Yes	5.00 GiB	
lzf6	✓ Online			600507669C9482580000000000000080	Yes	5.00 GiB	
lzf7	✓ Online			600507669C9482580000000000000081	Yes	5.00 GiB	
lzf8	✓ Online			600507669C9482580000000000000082	Yes	5.00 GiB	
lzf9	✓ Online			600507669C9482580000000000000083	Yes	5.00 GiB	
lzf_vols0	✓ Online			600507669C948258000000000000008D	Yes	5.00 GiB	
lzf_vols1	✓ Online			600507669C948258000000000000008E	Yes	5.00 GiB	
lzf_vols2	✓ Online			600507669C948258000000000000008F	Yes	5.00 GiB	
lzf_vols3	✓ Online			600507669C9482580000000000000090	Yes	5.00 GiB	

Figure 8-78 Volumes by Pool panel

The left pane is called the *pool filter*. The storage pools are displayed in the pool filter. For more information about storage pools, see Chapter 4, “Storage pools” on page 139.

In the upper right, you see information about the pool that you selected in the pool filter. The following information is also shown:

- ▶ **Pool icon:** Because storage pools can have different characteristics, you can change the storage pool icon by clicking it. For more information, see 4.2, “Working with storage pools” on page 150.
- ▶ **Pool name:** The name that was entered when the storage pool was created. Click it to change the name, if needed.
- ▶ **Pool details:** Shows you the information about the storage pools, such as the status, number of managed disks, and Easy Tier status.
- ▶ **Volume allocation:** Shows you the amount of capacity that is allocated to volumes from this storage pool.

The lower-right section lists all volumes with at least one copy in the selected storage pool. The following information is provided:

- ▶ **Name:** Shows the name of the volume.
- ▶ **State:** Shows the status of the volume.
- ▶ **UID:** Shows the volume unique identifier (UID).
- ▶ **Host mappings:** Shows whether host mappings exist.
- ▶ **Capacity:** Shows the capacity that is presented to hosts.

Also, you can create volumes from this panel. Click **Create Volumes** to open the Volume Creation panel. The steps are described in Chapter 6, “Volume configuration” on page 269.

Selecting a volume and opening the Actions menu or right-clicking the volume shows the same options as described in 8.3, “Advanced volume administration” on page 373.

8.7 Volumes by host

To see an overview of the volumes that a host can access, click **Volumes by Host**, as shown in Figure 8-79.




Figure 8-79 Volumes by Host option

The Volumes by Host panel opens, as shown in Figure 8-80.

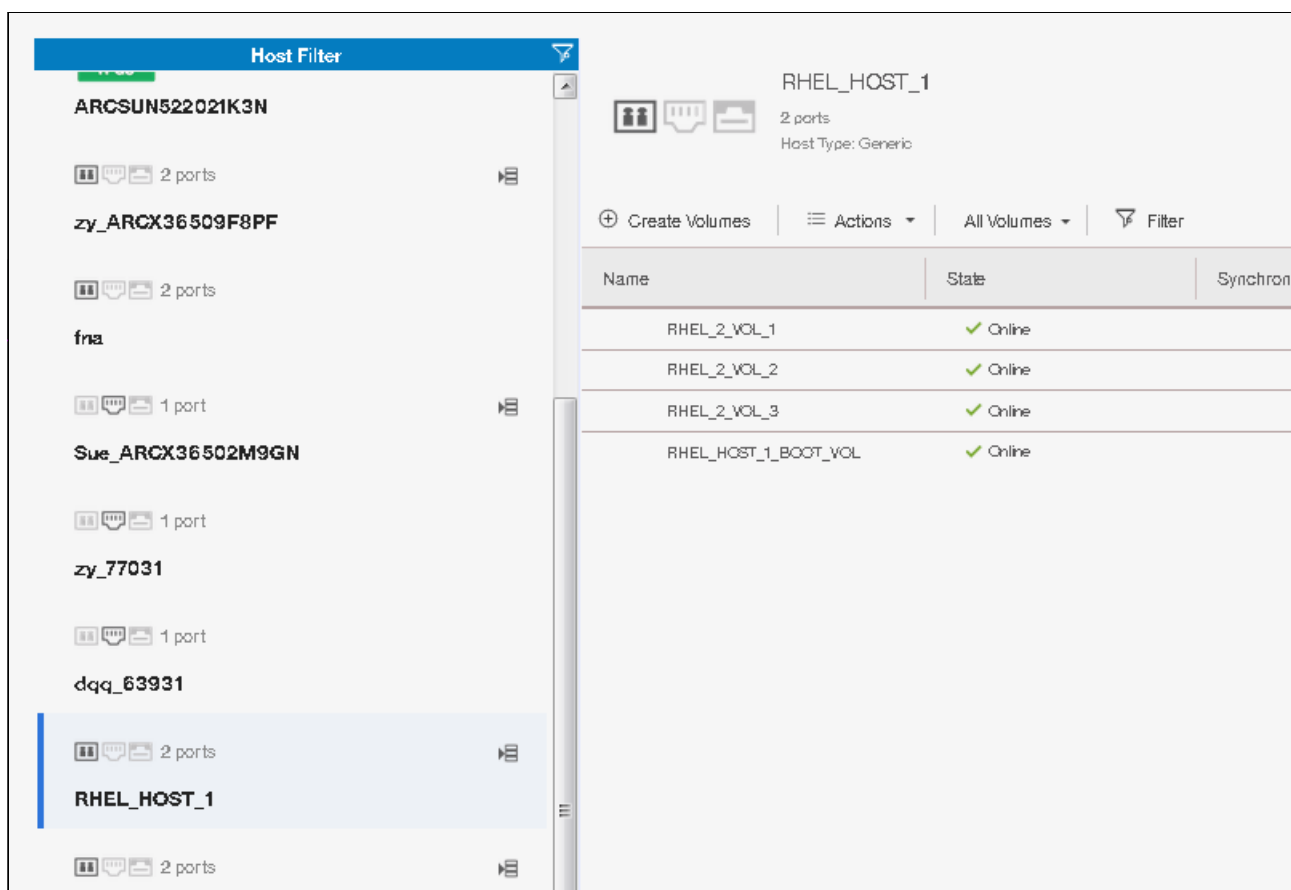


Figure 8-80 Volumes by Host panel

The *host filter* is in the left pane of the view. Selecting a host shows its properties in the right pane, such as the host name, number of ports, host type, and the I/O group to which it has access.

The right pane, next to the host name, shows icons for Fibre Channel, iSCSI and SAS connectivity. Depending on the type of host connectivity, the respective icon will be highlighted and the other icons will be grayed out.

The volumes that are mapped to this host are listed in the table in lower-right part of the panel.

You can create a volume from this panel. Click **Create Volumes** to open the same wizard as described in Chapter 6, “Volume configuration” on page 269.

Selecting a volume and opening the Actions menu or right-clicking the volume shows the same options as described in 8.3, “Advanced volume administration” on page 373.

Advanced features for storage efficiency

This chapter introduces the basic concepts of dynamic data relocation and storage optimization features. The IBM Spectrum Virtualize software running inside Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 offer IBM Easy Tier, thin provisioning and IBM Real-time Compression functions for storage efficiency. It provides only a basic technical overview and benefits of each feature.

Specifically, this chapter provides information about the following topics:

- ▶ 9.1, “Introduction” on page 404
- ▶ 9.2, “Easy Tier” on page 404
- ▶ 9.3, “Thin provisioning” on page 432
- ▶ 9.4, “Real-time Compression Software” on page 438

9.1 Introduction

In modern and complex application environments, the increasing and often unpredictable demands for storage capacity and performance, lead to planning and optimization issues related to storage resources.

Consider the following typical storage management issues:

- Usually when a storage system is implemented, only a portion of the configurable physical capacity is deployed. When the storage system runs out of its initial installed capacity and more capacity becomes necessary, a hardware upgrade is implemented to add physical resources to the storage system. This new physical capacity can hardly be configured to keep an even spread of the overall storage resources.

Typically, the new capacity is allocated to fulfill only new storage requests. The existing storage allocations do not benefit from the new physical resources. Similarly, the new storage requests do not benefit from the existing resources. Only new resources are used.

- In a complex production environment, it is not always possible to optimize storage allocation for performance. The unpredictable rate of storage growth and the fluctuations in throughput requirements, which are input/output (I/O) operations per second (IOPS), often lead to inadequate performance.

Furthermore, the tendency to use even larger volumes to simplify storage management works against the granularity of storage allocation, and a cost-efficient storage tiering solution becomes difficult to achieve. With the introduction of high-performing technologies, such as Flash drives or all-flash arrays, this challenge becomes even more important.

- The move to larger and larger physical disk drive capacities means that previous access densities that were achieved with low-capacity drives can no longer be sustained.
- Any business has applications that are more critical than others and there is a need for specific application optimization. Therefore, the ability to relocate specific application data to a faster storage media is required.
- Although more servers are purchased with internal SSD drives attached for better application response time, the data distribution across these internal SSD drives and external storage arrays must be carefully planned. An integrated and automated approach is crucial to achieve performance improvement without compromising data consistency, especially in a disaster recovery (DR) situation.

All of these issues deal with data placement, relocation capabilities or data volume reduction. Most of these challenges can be managed by having spare resources available, by moving data and by using data mobility tools or operating systems features (such as host level mirroring) to optimize storage configurations.

However, all of these corrective actions are expensive in terms of hardware resources, labor, and service availability. Relocating data among the physical storage resources that dynamically or effectively reduces the amount of data, transparently to the attached host systems, is becoming increasingly important.

9.2 Easy Tier

In today's storage industry, flash drives are emerging as an attractive alternative to hard disk drives (HDDs). Because of their low response times, high throughput, and I/O per second (IOPS) energy-efficiency, flash drives can help a storage infrastructure to achieve significant

savings in operational costs. However, the current acquisition cost per GB for flash is higher than for Enterprise serial-attached Small Computer System Interface (SCSI) (SAS) and Nearline (NL) SAS.

Enterprise SAS drives replaced the old SCSI drives. They are common in the storage market. They are offered in various capacities, spindle speeds and form factors. Nearline SAS is the low-cost, large-capacity storage drive class, which is commonly offered at 7200 rpm spindle speed.

It is critical to choose the correct mix of drives and the correct data placement to achieve optimal performance at the lowest cost. Maximum value can be achieved by placing “hot” data with high I/O density and low response time requirements on Flash. Enterprise class disks are targeted for “warm” and Nearline for “cold” data that is accessed sequentially and at lower rates.

In this section is described the Easy Tier disk performance optimization function of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. It also describes how to activate the Easy Tier process for both evaluation purposes and for automatic extent migration.

9.2.1 Easy Tier overview

Easy Tier is an optional licensed function of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 that bring enterprise storage enhancements to the entry and midrange segment. It enables automated subvolume data placement throughout different storage tiers to intelligently align the system with its current workload requirements and to optimize storage usage. This function includes the ability to automatically and non disruptively relocate data (at the extent level) from one tier to another tier in either direction to achieve the best available storage performance workload for an specific environment.

Easy Tier reduces the I/O latency for hot spots, but it does not replace storage cache. Easy Tier and storage cache solve a similar access latency workload problem, but these methods weigh differently in the algorithmic construction based on “locality of reference,” recency and frequency. Because Easy Tier monitors I/O performance from the extent end (after cache), it is able to pick up the performance issues that cache cannot solve and complement the overall storage system performance.

In general, the storage environment I/O is monitored on volumes and the entire volume is always placed inside one appropriate storage tier. Determining the amount of I/O on single extents is too complex for monitoring I/O statistics, to move them manually to an appropriate storage tier and to react to workload changes.

Easy Tier is a performance optimization function that overcomes this issue because it automatically migrates (or moves) extents that belong to a volume between different storage tiers, as shown in Figure 9-1 on page 406. Because this migration works at the extent level, it is often referred to as *sublogical unit number (LUN) migration*.

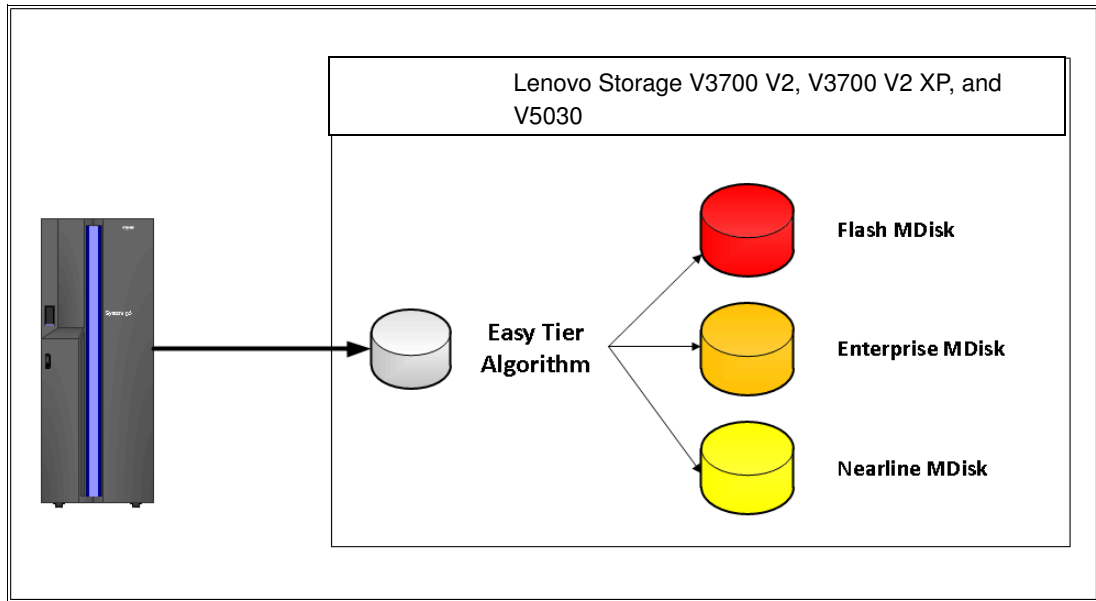


Figure 9-1 Easy Tier concept

You can enable Easy Tier for storage on a volume basis. It monitors the I/O activity and latency of the extents on all volumes that are enabled for Easy Tier over a 24-hour period. Based on the performance log, it creates an extent migration plan and dynamically moves high activity or hot extents to a higher disk tier within the same storage pool. It also moves extents in which the activity rate dropped off (or cooled) from higher disk tier managed disks (MDisks) back to a lower tier MDisk.

To enable the migration between MDisks with different tier levels, the target storage pool must consist of MDisks with different characteristics. These pools are named as *multi-tiered storage pools*. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Easy Tier is optimized to boost the performance of storage pools that contain Flash, Enterprise and Nearline drives.

To identify the potential benefits of Easy Tier in your environment before you install higher MDisk tiers (such as Flash), you can enable the Easy Tier monitoring on volumes in single-tiered storage pools. Although the Easy Tier extent migration is not possible within a single-tiered pool, the Easy Tier statistical measurement function is possible. Enabling Easy Tier on a single-tiered storage pool starts the monitoring process and logs the activity of the volume extents.

IBM Storage Tier Advisor Tool (STAT) is a no-cost tool that helps you analyze this data. If you do not have a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, use the Disk Magic tool to get a better idea about the required number of different drive types that are appropriate for your workload.

Easy Tier is available for all the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 internal volumes and volumes on external virtualized storage subsystems (V5030).

9.2.2 Tiered storage pools

With the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, we must differentiate between the following types of storage pools:

- ▶ Single-tiered storage pools
- ▶ Multi-tiered storage pools

Figure 9-2 shows single-tiered storage pools which include one type of disk tier attribute. Each disk, ideally, has the same size and performance characteristics. Multi-tiered storage pools are populated with two or more different disk tier attributes, high-performance flash drives, enterprise SAS drives, and Nearline drives.

A volume migration occurs when the complete volume is migrated from one storage pool to another storage pool. An Easy Tier data migration moves only extents inside the storage pool to different performance attributes.

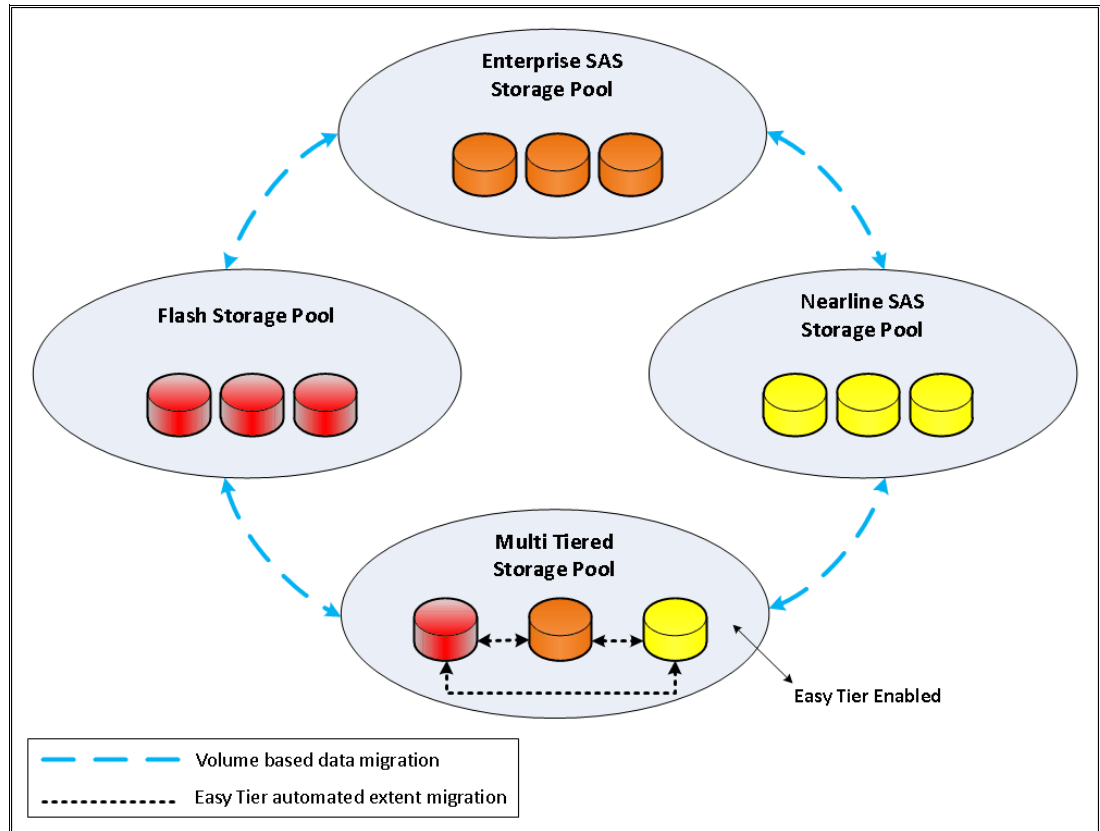


Figure 9-2 Tiered storage pools

By default, Easy Tier is enabled on any pool that contains two or more classes of disk drives. The Easy Tier function manages the extent migration:

- **Promote**
Moves the candidate hot extent to a higher performance tier.
- **Warm demote:**
 - Prevents performance overload of a tier by demoting a warm extent to a lower tier.
 - Triggered when bandwidth or I/O per second (IOPS) exceeds a predefined threshold.
- **Cold demote**
Coldest extent moves to a lower tier.
- **Expanded or cold demote**
Demote appropriate sequential workload to the lowest tier to better use nearline bandwidth.
- **Swap**

This operation exchanges a cold extent in a higher tier with a hot extent in a lower tier or vice versa.

Note: Extent migrations occur only between adjacent tiers within the same pool.

Figure 9-3 shows the Easy Tier extent migration.




Figure 9-3 Easy Tier actions

9.2.3 Easy Tier process

Easy Tier is based on an algorithm with a threshold to evaluate if an extent is cold, warm, or hot. Easy Tier consists of four main processes. These processes ensure that the extent allocation in multi-tiered storage pools is optimized for the best performance, based on your workload in the last 24 hours. The processes are listed:

- ▶ I/O Monitoring
- ▶ Data Placement Advisor
- ▶ Data Migration Planner
- ▶ Data Migrator

Figure 9-4 shows the flow between these processes.




Figure 9-4 Easy Tier process flow

The four main processes and the flow between them are described in the following sections.

9.2.4 I/O Monitoring

The I/O Monitoring (IOM) process operates continuously and monitors host volumes for I/O activity. It collects performance statistics for each extent at 5-minute intervals and derives averages for a rolling 24-hour period of I/O activity.

Easy Tier permits large block I/Os and considers only I/Os up to 64 KB as migration candidates.

IOM is an efficient process and adds negligible processing impact to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 node canisters.

9.2.5 Data Placement Advisor

The Data Placement Advisor (DPA) uses workload statistics to make a cost benefit decision about the extents that need to be candidates for migration to a higher-performance tier.

This process also identifies extents that must be migrated back to a lower tier.

9.2.6 Data Migration Planner

By using the previously identified extents, the Data Migration Planner (DMP) process builds the extent migration plan for the storage pool.

9.2.7 Data Migrator

The Data Migrator (DM) process involves scheduling and the actual movement, or migration, of the volume's extents up to, or down from, the high disk tier.

The extent migration rate is described.

Easy Tier cycles moves extents at a rate of approximately 12 GB every 5 minutes as a maximum speed:

- ▶ This speed is applied to all functions except warm promotes and warm demotes.

- ▶ If an Easy Tier cycle can generate cold demotes among the other operations, the speed is reduced to 11GB every 5 minutes.
- ▶ If an Easy Tier cycle has only cold demotes to be performed, the full 12GB limit can be used.

This rate equates to around 3 TB a day that is migrated between disk tiers. Figure 9-5 shows the Easy Tier Data Migrator flow.




Figure 9-5 Easy Tier Data Migrator

9.2.8 Easy Tier accelerated mode

Under normal production conditions, Easy Tier works properly to process daily workloads.

Easy Tier considers migration scenarios and scenarios where large amounts of data need to be rebalanced in its internal algorithms.

Easy Tier accelerated mode was introduced in controller firmware version 7.5. Easy Tier accelerated mode allows the system to cope with migration situations where the user needs to speed up the Easy Tier function temporarily.

Normal Easy Tier migration speed is 12 GB every 5 minutes for all functions, except cold demote, which is 1 GB every 10 minutes.

Accelerated mode allows an Easy Tier migration speed of 48 GB every 5 minutes with no limit on cold demotes and no support for warm demotes.

You enable Easy Tier accelerated mode from the command line by using **chsystem -easytieracceleration on/off**.

Note: Accelerated mode is not intended for day-to-day Easy Tier traffic. Turn on accelerated mode when necessary. Because Easy Tier accelerated mode can increase the workload on the system temporarily, use Easy Tier accelerated mode during periods of lower system activity.

9.2.9 Easy Tier operating modes

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 offer the following operating modes for Easy Tier:

- Easy Tier: Off

Easy Tier can be turned off. No statistics are recorded and no extents are moved.

- Easy Tier: On

When the Easy Tier function is turned on, Easy Tier measures the I/O activity for all extents. With a multi-tiered pool, the extents are migrated dynamically by the Easy Tier processes to achieve the best performance. The movement is transparent to the host server and applications.

A statistic summary file is created. This file can be off-loaded and analyzed with the IBM Storage Tier Advisory Tool, as described in 9.2.16, “IBM Storage Tier Advisor Tool” on page 431. Easy Tier can be turned on for any single-tiered or multi-tiered pool, but its functionality will differ on each of them.

- Easy Tier: Measured mode

When Easy Tier is in measured mode, Easy Tier measures the I/O activity for all extents but it does not move any extents within the storage pool. A statistics summary file is created. This file can be off-loaded from the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. This file can be analyzed with the IBM Storage Tier Advisory Tool. This analysis shows the workload benefits of adding or removing different drive classes from a pool before any hardware is acquired. No license is required to set Easy Tier to measured mode.

- Easy Tier: Auto

Auto mode is the operating mode default. If Easy Tier is set to auto for a single-tiered storage pool, Easy Tier is set to off for all volumes inside the storage pool and no extents are moved. If Easy Tier is set to auto for a multi-tiered storage pool, the Easy Tier status becomes active and Easy Tier is set to on for all volumes inside the storage pool and the extents are migrated dynamically by the Easy Tier process. However, the extents are not migrated if the Easy Tier function is not licensed.

9.2.10 Easy Tier status

Depending on the Easy Tier mode attributes, the storage pool Easy Tier status can be one of the following values:

- Active: This status indicates that Easy Tier is actively managing the extents of the storage pool.
- Balanced: This status applies to homogeneous storage pools and indicates that Easy Tier is actively managing the extents to provide enhanced performance by rebalancing the extents among the MDisk within the tier. This rebalancing characteristic is called *Storage Pool Balancing*, which is described in 9.2.11, “Storage Pool Balancing” on page 412.

- **Measured:** This status means that Easy Tier is constantly measuring the I/O activity for all extents to generate an I/O statistics report, but no extents are being moved within that pool.
- **Inactive:** When the Easy Tier status is inactive, no extents are monitored and no statistics are recorded.

9.2.11 Storage Pool Balancing

Storage Pool Balancing is associated with Easy Tier. It operates independently and does not require an specific license. Storage Pool Balancing works with Easy Tier when multiple MDisks exist in a single pool.

Note: At the time of the creation of a new pool the default Easy Tier status is shown on the pool properties as *Ba*lanced, but the pool will not benefit from the Storage Pool Balancing feature without multiple MDisks within it.

It assesses the extents in a storage tier and balances them automatically across all MDisks within that tier. Storage Pool Balancing moves the extents to achieve a balanced workload distribution and avoid hotspots. Storage Pool Balancing is an algorithm that is based on MDisk IOPS usage, which means that it is not capacity-based but performance-based. It works on a 6-hour performance window.

When a new MDisk is added to an existing storage pool, Storage Pool Balancing can automatically balance the extents across all MDisks in the pool, if required.

Figure 9-6 represents an example of Storage Pool Balancing.




Figure 9-6 Storage Pool Balancing

9.2.12 Easy Tier rules

The following operating rules apply when IBM System Storage Easy Tier is used on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030:

- Automatic data placement and extent I/O activity monitors are supported on each copy of a mirrored volume. Easy Tier works on each copy independently of each other.

Volume mirroring: *Volume mirroring* can have different workload characteristics for each copy of the data because reads are normally directed to the primary copy and writes occur to both copies. Therefore, the number of extents that Easy Tier migrates probably differs for each copy.

- ▶ Easy Tier works with all striped volumes, including these types of volumes:
 - Generic volumes
 - Thin-provisioned volumes
 - Mirrored volumes
 - Thin-mirrored volumes
 - Global and Metro Mirror sources and targets
- ▶ Easy Tier automatic data placement is not supported for image mode or sequential volumes. I/O monitoring for these volumes is supported, but you cannot migrate extents on these volumes unless you convert image or sequential volume copies to striped volumes.
- ▶ The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 create volumes or volume expansions by using extents from MDisk from the Enterprise and Nearline tier. Extents from MDisk in the Flash tier are used if Enterprise space and Nearline space are not available.
- ▶ When a volume is migrated out of a storage pool that is managed with Easy Tier, Automatic Data Placement Mode is no longer active on that volume. Automatic Data Placement is also turned off while a volume is migrated, even if it is between pools that both have Easy Tier Automatic Data Placement enabled. Automatic Data Placement for the volume is re-enabled when the migration is complete.
- ▶ Flash drive performance depends on block size. (Small blocks perform better than large blocks.) Easy Tier measures I/O blocks that are smaller than 64 KB, but it migrates the entire extent to the appropriate disk tier.
- ▶ As extents are migrated, the use of smaller extents makes Easy Tier more efficient.
- ▶ The first migration starts about 1 hour after Automatic Data Placement Mode is enabled. It takes up to 24 hours to achieve optimal performance.
- ▶ In the current Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Easy Tier implementation, it takes about two days before hotspots are considered moved from tier to tier, which prevents hotspots from being moved from a fast tier if the workload changes over a weekend.
- ▶ If you run an unusual workload over a longer period, Automatic Data Placement can be turned off and turned on online to avoid data move.

Depending on which storage pool and which Easy Tier configuration is set, a volume copy can have the Easy Tier states that are shown in Table 9-1 on page 414.

Table 9-1 Easy Tier states

Storage pool	Single-tiered or multi-tiered storage pool	Volume copy Easy Tier setting	Easy Tier status on volume copy
Off	Single-tiered	Off	Inactive
Off	Single-tiered	On	Inactive
Off	Multi-tiered	Off	Inactive
Off	Multi-tiered	On	Inactive
Auto ^a	Single-tiered	Off	Measured ^b
Auto ^a	Single-tiered	On	Balanced (see footnote e)
Auto ^a	Multi-tiered	Off	Measured ^b
Auto ^a	Multi-tiered	On	Active ^{c d}
On	Single-tiered	Off	Measured ^b
On	Single-tiered	On	Balanced (see footnote e)
On	Multi-tiered	Off	Measured ^b
On	Multi-tiered	On	Active ^c
Measure	Single-tiered	Off	Measured ^b
Measure	Single-tiered	On	Measured ^b
Measure	Multi-tiered	Off	Measured ^b
Measure	Multi-tiered	On	Measured ^b

a. The default Easy Tier setting for a storage pool is Auto, and the default Easy Tier setting for a volume copy is On. This scenario means that Easy Tier functions are disabled for storage pools with a single tier and only Storage Pool Balancing is active.

b. When the volume copy status is measured, the Easy Tier function collects usage statistics for the volume, but automatic data placement is not active.

c. If the volume copy is in image or sequential mode or is being migrated, the volume copy Easy Tier status is measured instead of active.

d. When the volume copy status is active, the Easy Tier function operates in automatic data placement mode for that volume.

e. When the volume Easy Tier status is balanced, Easy Tier is actively managing the extents by rebalancing them among the MDisk within the tier.

9.2.13 Creating multi-tiered pools: Enabling Easy Tier

In this section, we describe how to create multi-tiered storage pools by using the GUI.

When a storage pool changes from single-tiered to multi-tiered, Easy Tier is enabled by default for the pool and on all volume copies inside this pool. The current release of Easy Tier supports up to three tiers of storage (Flash, Enterprise, and Nearline).

In this example, we create a pool that contains Enterprise and Nearline MDisk.

To create a multi-tiered pool, complete the following steps:

1. Navigate to **Pools** → **Pools** as shown in Figure 9-7 on page 415. Click **Create** to open the Create Pool panel.

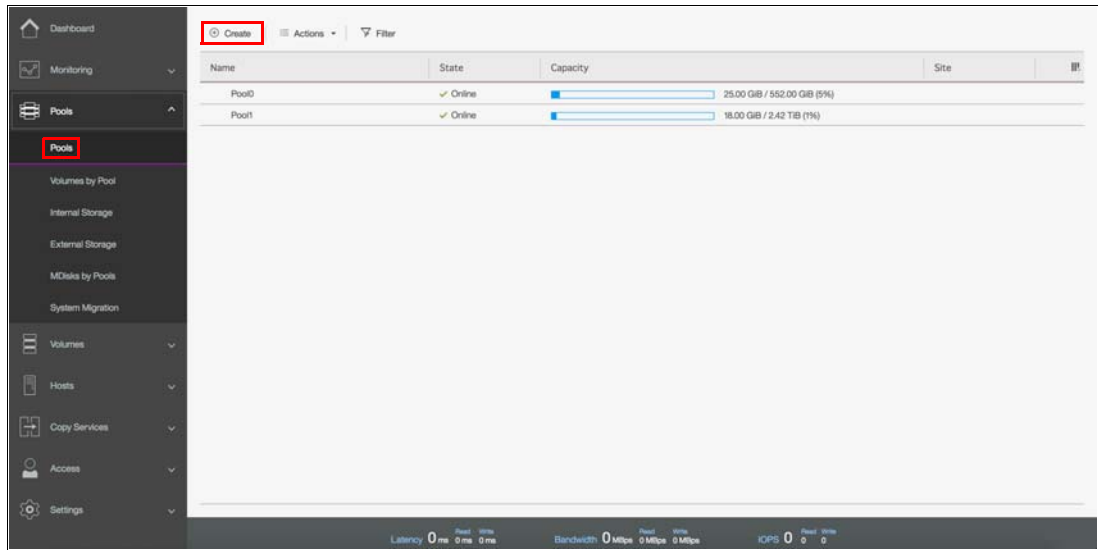


Figure 9-7 Selecting create in the Pools panel

2. Provide a name for the new pool and click **Create**. Encryption option will be available if the system has an encryption license enabled, click **Enable** if you want to enable pool encryption. If you navigate to **Settings** → **GUI Preferences** and click **General**, the **Advanced pool settings** can be selected, which allows you to define the extent size during a pool creation as shown in Figure 9-8.

Create Pool

Name:

Easy_Tier_Pool

Extent size:

1.00 GiB

▼

Maximum addressable capacity: 4.00 PiB

Encryption:

☒ Enable

Cancel

Create

Figure 9-8 Creating an Easy Tier pool

3. After creating the pool, it will be displayed in the Pools list as shown in Figure 9-9.

Name	State	Capacity	Site
Pool0	✓ Online	25.00 GB / 552.00 GB (5%)	
Pool1	✓ Online	18.00 GB / 2.42 TB (1%)	
Easy_Tier_Pool	✓ Online	0 bytes	

Latency: 0 ms Read 0 ms Write 0 ms
Bandwidth: 0 MBps Read 0 MBps Write 0 MBps
IOPS: 0 Read 0 Write 0

Figure 9-9 Pools list

4. To show the pool properties, select the pool and select **Properties** from the Actions menu. Alternatively, right-click the pool and select **Properties**. Clicking **View more details** in the bottom-left of the panel will display additional pool information. No storage is assigned to the pool at the time of its creation and the Easy Tier default status is set to **Balanced** as shown in Figure 9-10 on page 417.

Properties for Pool Easy_Tier_Pool

Name:

Easy_Tier_Pool

State:

☒ No Storage

Capacity:

Volumes:

0

MDisks:

0

Thin provisioning savings:

0 bytes

Compression savings:

0 bytes

Total savings:

0 bytes

Encryption:

Encrypted

Easy tier:

Balanced

Extent size:

1.00 GiB

Close

Figure 9-10 Pool properties panel

5. To add storage to a pool you can either select the pool and click **Add Storage** from the Actions menu or right-clicking the pool, as shown in Figure 9-11.

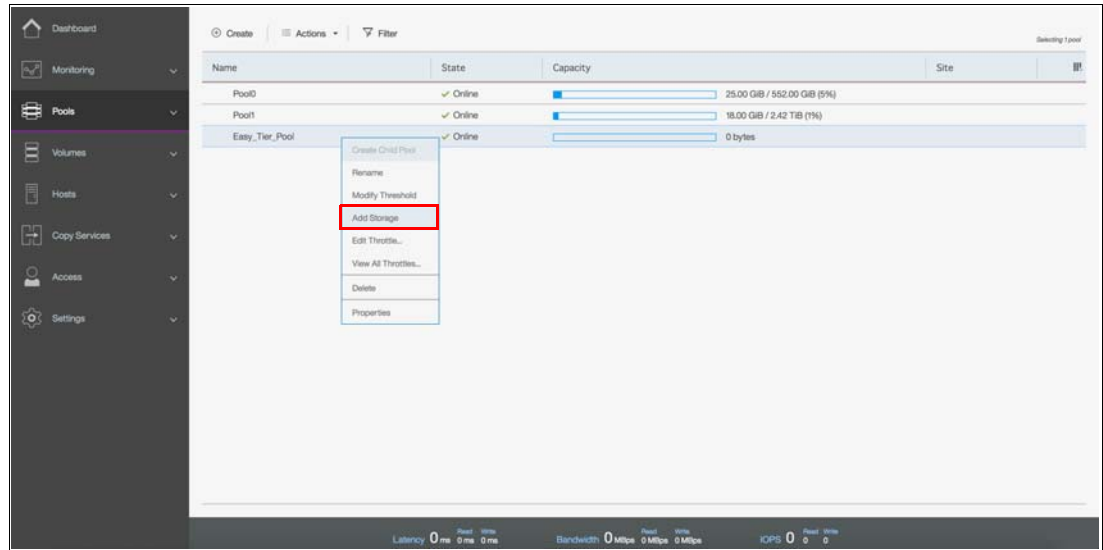


Figure 9-11 Add Storage to a pool

6. The Assign Storage to Pool panel offers two options to configure the storage into the pool: Quick Internal or Advanced Internal Custom. Figure 9-12 shows the Quick Internal panel.

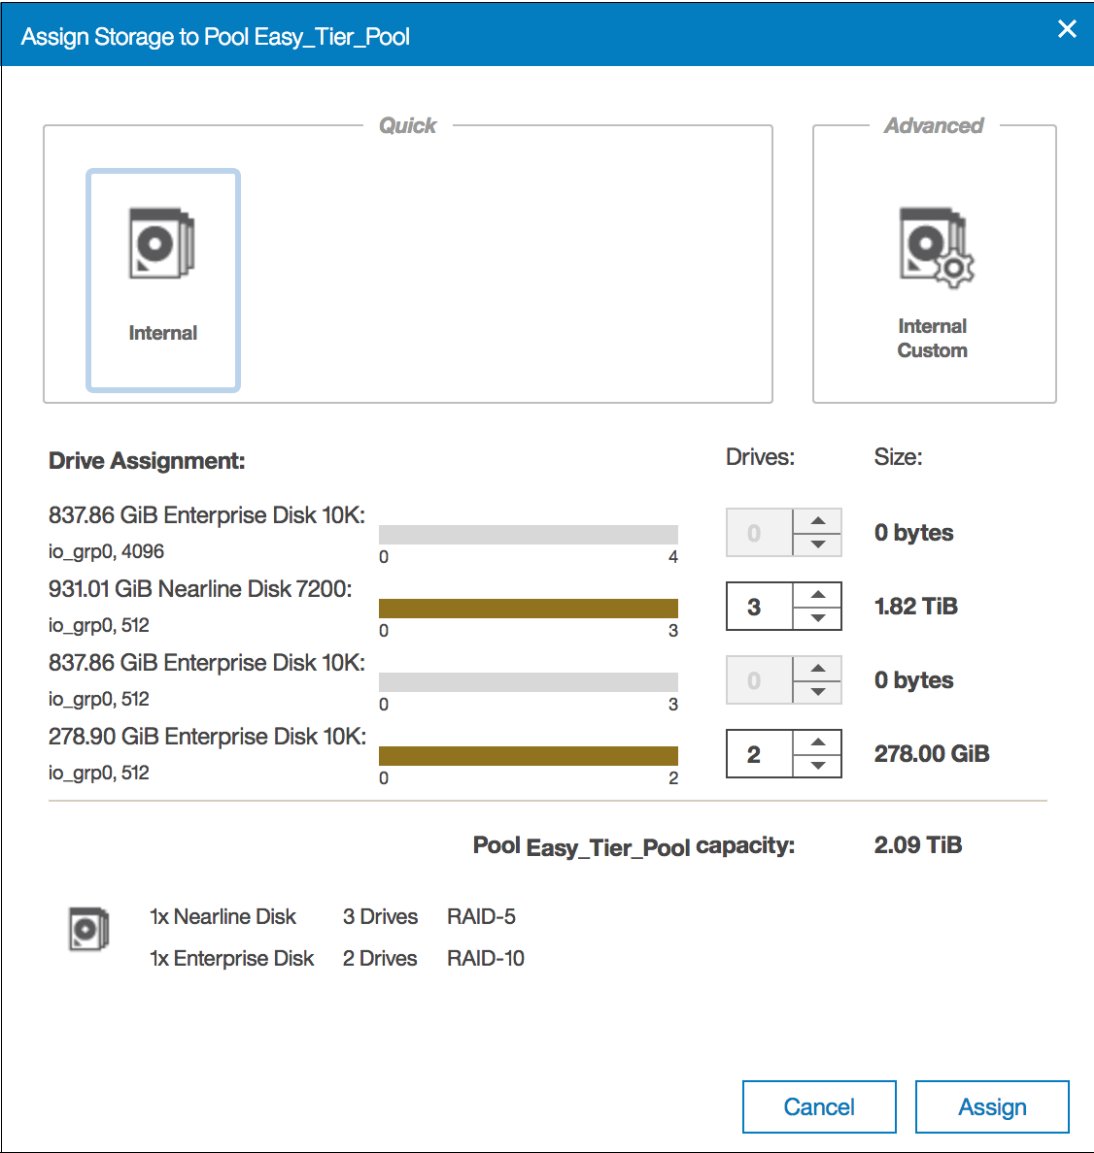


Figure 9-12 Add Storage - Quick Internal panel

7. Figure 9-13 shows the Advanced Internal Custom panel.

The screenshot shows a window titled "Assign Storage to Pool Easy_Tier_Pool" with a close button (X) in the top right corner. The window is divided into two main sections: "Quick" and "Advanced".

The "Quick" section on the left contains a disk icon and the label "Internal".

The "Advanced" section on the right is highlighted with a blue border and contains a disk icon with a gear, labeled "Internal Custom".

Below these sections is a blue header bar labeled "Drive Assignment".

Under "Drive Assignment", the "Drive Class" is set to "837.86 GiB Enterprise Disk 10K" with a dropdown arrow.

The "Drives" section shows a progress bar from 0 to 4, with a value of 4 selected in a box with up/down arrows.

The "RAID" section shows "RAID-5" selected in a dropdown.

The "Spares" section shows "0" selected in a box with up/down arrows.

The "Stripe width" section shows "4" selected in a box with up/down arrows.

The "Array width" section shows "4" selected in a box with up/down arrows.

The "Size" section shows "2.45 TiB".

Below these settings, a line indicates "Pool Easy_Tier_Pool capacity: 2.45 TiB" with a refresh icon.

At the bottom, there is a "Summary" section and two buttons: "Cancel" and "Assign".

Figure 9-13 Add Storage - Advanced Internal Custom panel


The Quick Internal panel provides a recommended configuration that is based on the number and type of installed drives. You can use the Advanced Internal Custom panel to configure the specific drive types, Redundant Array of Independent Disks (RAID) levels, spares, stripe width, and array width.

In the following steps the Advanced Internal Custom option is used to create a single-tiered storage pool and then a multi-tiered storage pool by including another drive class to the single-tiered one. Each drive class needs to be included separately.


8. From the Advanced Internal Custom panel, select the required drive class, RAID type, number of spares, stripe width and array width. Click **Assign** to add the storage to the pool, as shown in Figure 9-14.

Assign Storage to Pool Easy_Tier_Pool

Quick


Internal

Advanced


Internal Custom

Drive Assignment

Drive Class:

837.86 GiB Enterprise Disk 10K

Drives:

0

9

6

RAID:

Distributed RAID-6

Spares:

1

Stripe width:

5

Array width:

6

Size:

2.42 TiB

Pool Easy_Tier_Pool capacity:

2.42 TiB

Summary

Cancel

Assign

Figure 9-14 Adding first drive class to a pool

9. Select the pool to which the storage was added. Select **Properties** from the Actions Menu and click **View more details** in the Properties panel. Although the Easy Tier status of a single-tiered pool is kept as **Balanced** as shown in Figure 9-15, the pool will not benefit from the Storage Pool Balancing feature with a single MDisk.

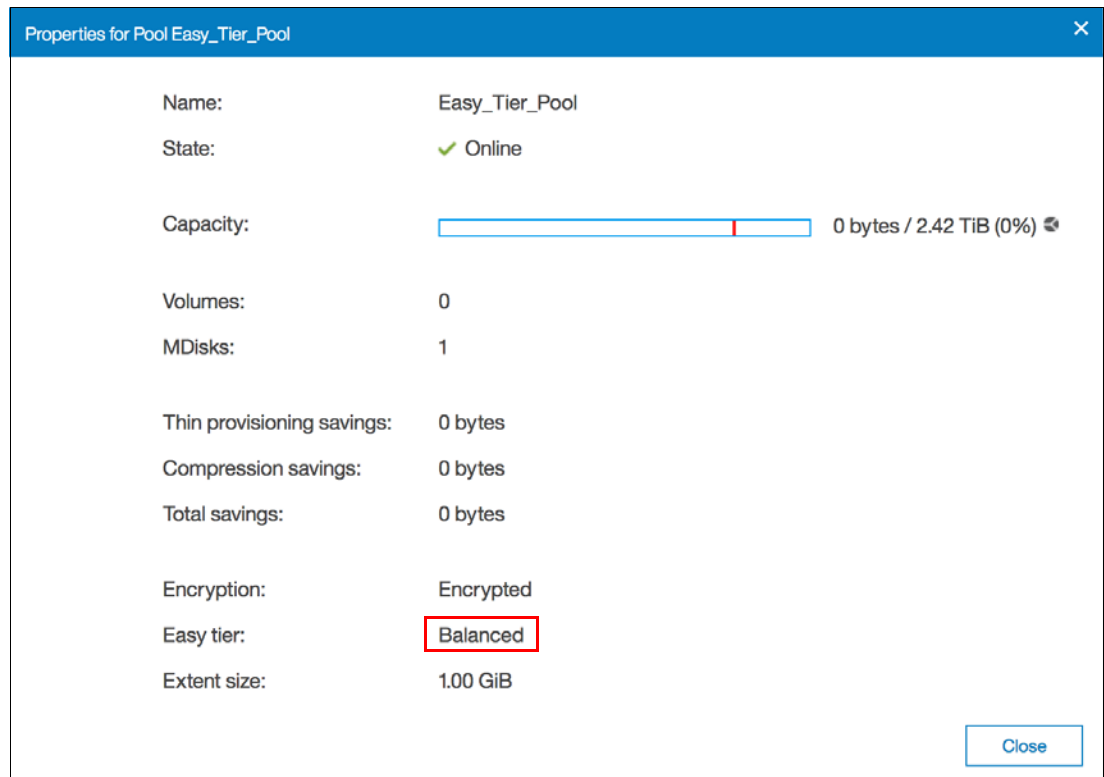


Figure 9-15 *Balanced easy tier status on Properties panel*

10.Repeat steps 7 and 8 to add a second drive class as shown in Figure 9-16.

Assign Storage to Pool Easy_Tier_Pool

Quick

Internal

Advanced

Internal Custom

Drive Assignment

Drive Class: 931.01 GiB Nearline Disk 7200

Drives: 0 3

RAID: RAID-5

Spares: 0

Stripe width: 3

Array width: 3

Size: 1.82 TiB

Pool Easy_Tier_Pool capacity: 4.24 TiB

Summary

Cancel Assign

Figure 9-16 Adding second drive class to a pool

11. Select the pool to which the second drive class was added. Select **Properties** from the Actions Menu and click **View more details** in the Properties panel. With two different tiers the Easy Tier status is automatically changed to Active (Figure 9-17 on page 424) and starts to manage the extents within the pool by promoting or demoting them.

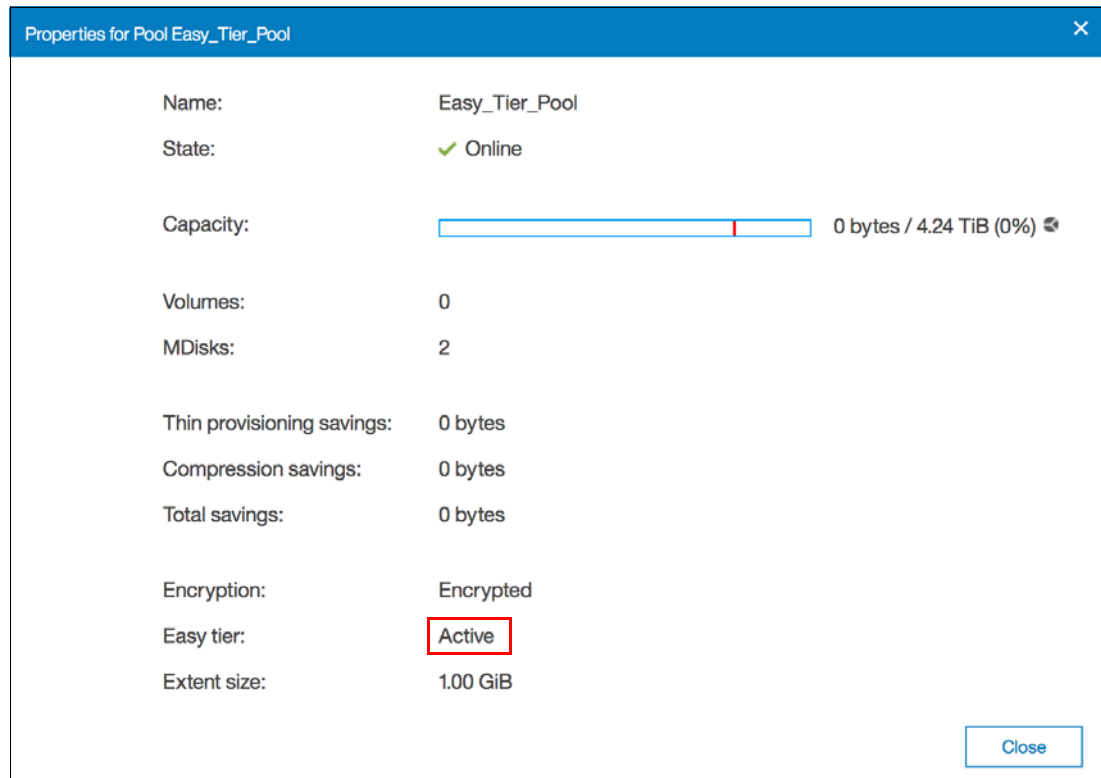


Figure 9-17 Pool properties that show that Easy Tier is active

Note: Adding multiple MDisks of the same drive class will result in a single-tiered pool with Balanced Easy Tier status, which will only benefit from the Storage Pool Balancing feature.

12. Navigate to **Pools** → **MDisks by Pools** to see the MDisks that were created within the pool with two different drive classes. The tier information is not a default column in the MDisks by Pools panel. To access the tier information right-click the gray header and select **Tier**. Each MDisk will display its tier class as shown in Figure 9-18.

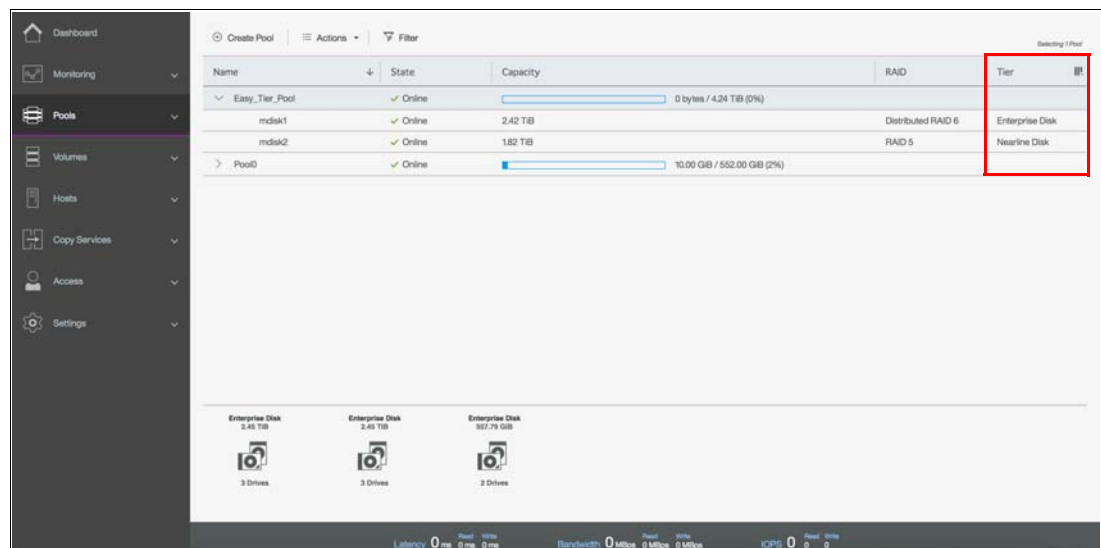


Figure 9-18 MDisks by Pools panel

If you create a volume within a multi-tiered storage pool and navigate to **Volumes** → **Volumes by Pool** panel, details such as the number of MDisks and the number of volumes within the selected pool are displayed. The pool icon for an Easy Tier Pool differs from the pools without Easy Tier enabled as shown in Figure 9-19.

If the Easy Tier Status column is enabled, the Easy Tier Status of each volume is displayed. Volumes inherit the Easy Tier state of their parent pool, but Easy Tier can be toggled on or toggled off at the volume level, if required. See “Enabling or disabling Easy Tier on single volumes” on page 429.

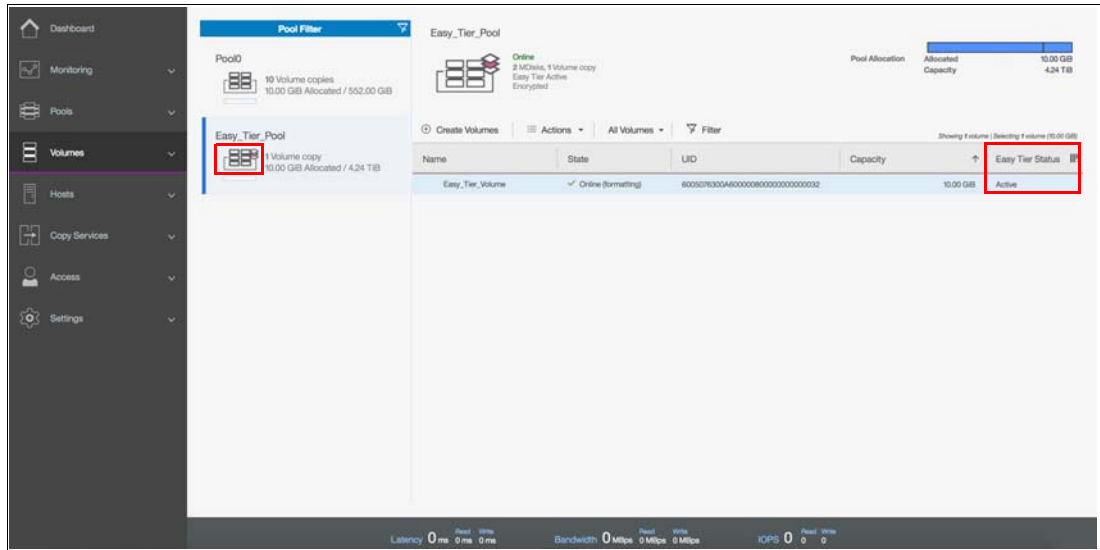


Figure 9-19 Volumes by pool panel

If external storage is used as a drive class, you must select the drive class type manually and add the external MDisks to a storage pool. If the internal storage and the external storage are in different drive classes, this action also changes the storage pool to a multi-tiered storage pool and enables Easy Tier on the pool and associated volumes.

9.2.14 Downloading Easy Tier I/O measurements

After enabling Easy Tier the Automatic Data Placement Mode is active. Extents are automatically migrated to or from disk tiers and the statistics summary collection is now active. The statistics log file can be downloaded to analyze how many extents were migrated and to monitor whether it makes sense to add more drives from an specific drive class to the multi-tiered storage pool.

Heat data files are produced approximately once a day (that is, roughly every 24 hours) when Easy Tier is active on one or more storage pools.

To download the statistics file, complete the following steps:

1. Navigate to **Settings** → **Support**. In the Support panel click Support Package as shown in Figure 9-20 on page 426. Click **Manual Upload Instructions** to display the **Download Support Package** button.

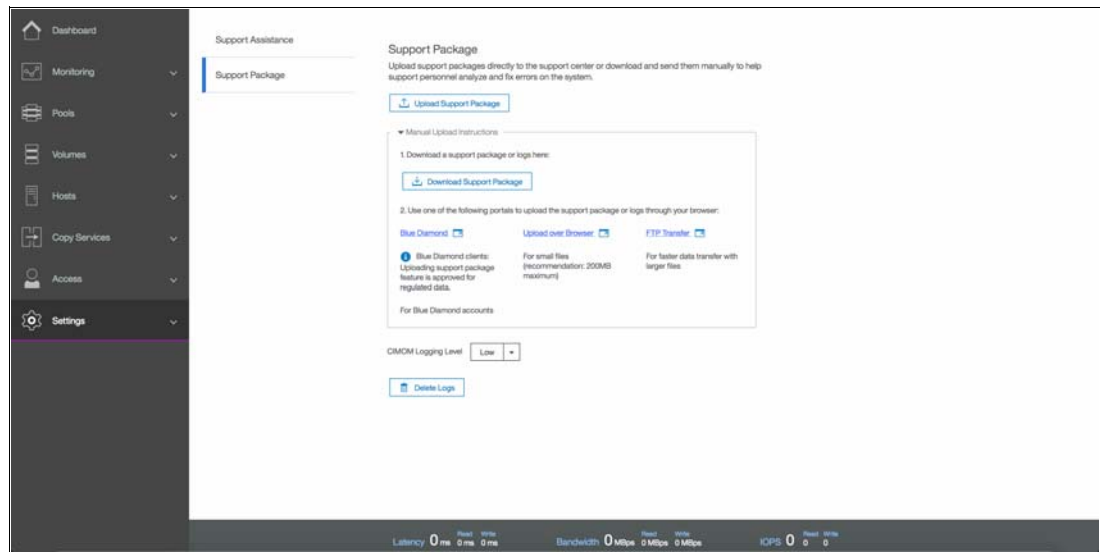


Figure 9-20 Support panel under Settings menu

Click **Download Support Package** to open the **Download New Support Package or Log File** panel, as shown in Figure 9-21.

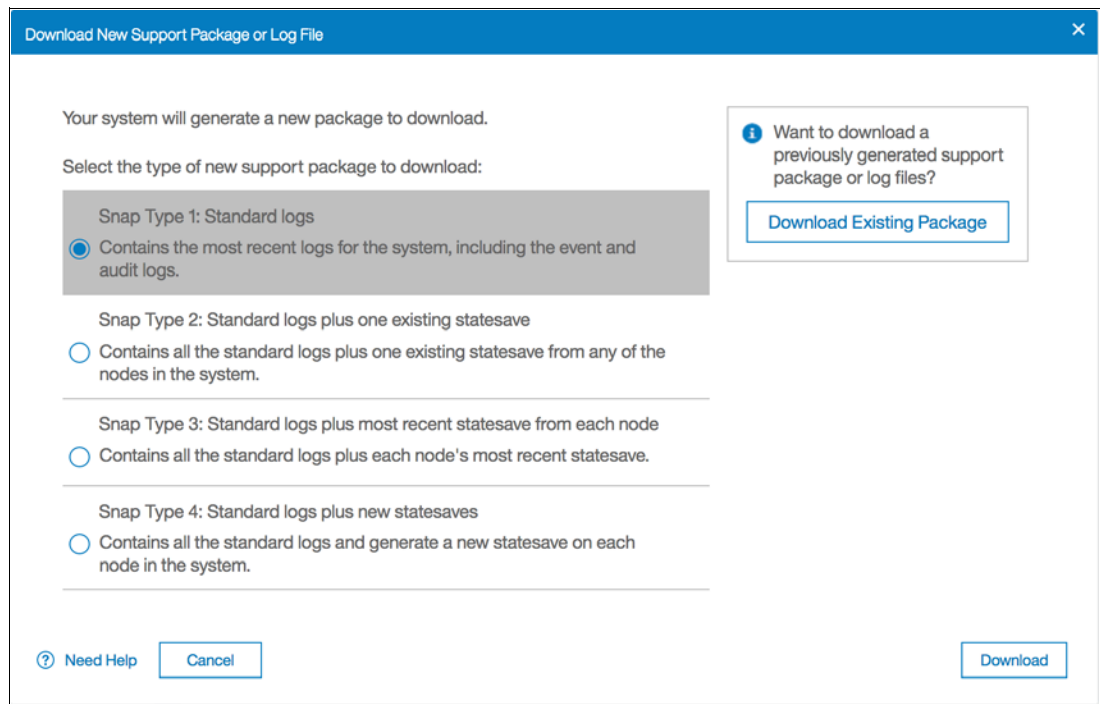


Figure 9-21 Download New Support Package or Log File panel

To download the Easy Tier log files you have two options:

- Choose one of the Snap Types shown in Figure 9-21 and click **Download**. The entire support package is downloaded and the Easy Tier log file is available within it.
- Click **Download Existing Package** to open the panel shown in Figure 9-22 on page 427. Select the required Easy Tier log file and click **Download**.

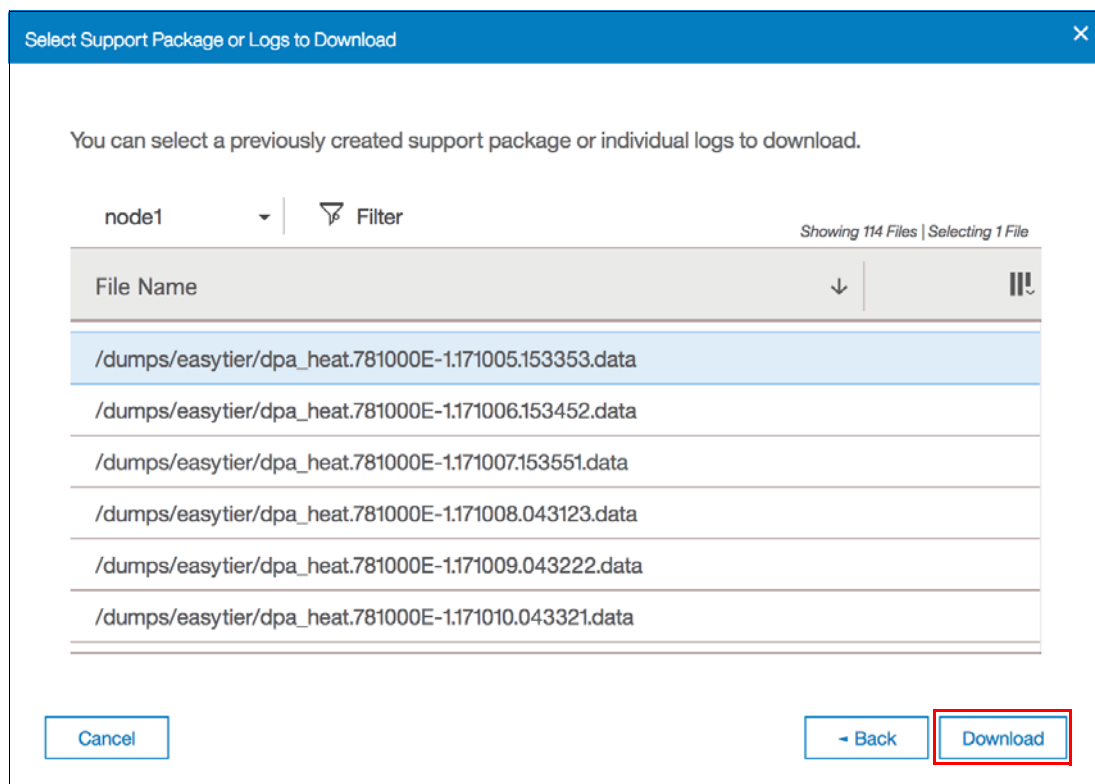


Figure 9-22 Select Support Package or Logs to Download panel

The Easy Tier log files are always named `dpa_heat.canister_name.date.time.data`.

If you run Easy Tier for a longer period, it generates a heat file at least every 24 hours. The time and date that the file was created is included in the file name.

9.2.15 Easy Tier I/O Measurement through the command-line interface

Easy Tier can also be configured through the command-line interface (CLI). For the advanced user, this method offers more options for Easy Tier configuration.

Before you use the CLI, you must configure CLI access, as described in Appendix A, “CLI setup and SAN Boot” on page 761.

Readability: In the examples that are shown in this section, we deleted many unrelated lines in the command output or responses so that you can focus on the information that relates to Easy Tier.

Enabling Easy Tier measured mode

You can enable Easy Tier in measured mode on either a single-tiered or multi-tiered storage pool. Connect to your Lenovo storage V-series system by using the CLI and run the **svcinfo lsmdiskgrp** command, as shown in Example 9-1. This command shows an overview of all configured storage pools and their Easy Tier status. In our example, two storage pools are listed: `Enterprise_Pool` with Easy Tier in auto status, and `Multi_Tier_Pool` with Easy Tier on.

Example 9-1 Show all configured storage pools

```
IBM_Storwize:ITS0_V5000:superuser>svcinfo lsmdiskgrp
id name                status mdisk_count easy_tier easy_tier_status type
```

```

0 Enterprise_Pool online 1          auto      balanced      parent
1 Multi_Tier_Pool online 3         on        active        parent
IBM_Storwize:ITS0_V5000:superuser>

```

To enable Easy Tier on a single-tiered storage pool in measure mode, run the **chmdiskgrp -easytier measure storage pool name** command, as shown in Example 9-2.

Example 9-2 Enable Easy Tier in measure mode on a single-tiered storage pool

```

IBM_Storwize:ITS0_V5000:superuser>chmdiskgrp -easytier measure Enterprise_Pool
IBM_Storwize:ITS0_V5000:superuser>

```

Check the status of the storage pool again by running the **lsmdiskgrp storage pool name** command again, as shown in Example 9-3.

Example 9-3 Storage pool details: Easy Tier measure status

```

IBM_Storwize:ITS0_V5000:superuser>lsmdiskgrp Enterprise_Pool
id 0
name Enterprise_Pool
status online
mdisk_count 1
vdisk_count 2
capacity 1.81TB
extent_size 1024
free_capacity 1.80TB
virtual_capacity 11.00GB
used_capacity 11.00GB
real_capacity 11.00GB
overallocation 0
warning 90
easy_tier measure
easy_tier_status measured
tier ssd
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier enterprise
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier nearline
tier_mdisk_count 1
tier_capacity 1.81TB
tier_free_capacity 1.80TB
parent_mdisk_grp_id 2
parent_mdisk_grp_name Enterprise_Pool
child_mdisk_grp_count 0
child_mdisk_grp_capacity 0.00MB
type parent
IBM_Storwize:ITS0_V5000:superuser>

```

Easy Tier measured mode does not place data. Easy Tier measured mode collects statistics for measurement only. For more information about downloading the I/O statistics, see 9.2.14, “Downloading Easy Tier I/O measurements” on page 425.

Enabling or disabling Easy Tier on single volumes

By default, enabling Easy Tier on a storage pool also enables it for the volume copies that are inside the selected pool. This setting applies to multi-tiered and single-tiered storage pools. It is also possible to turn on and turn off Easy Tier for single volume copies.

Before you disable Easy tier on a single volume, run the **svcinfn lsmdisgrp storage pool name** command to list all storage pools that are configured, as shown in Example 9-4. In our example, Multi_Tier_Pool is the storage pool that is used as a reference.

Example 9-4 Listing the storage pool

```
IBM_Storwize:ITS0_V5000:superuser>svcinfn lsmdisgrp
id name                status easy_tier easy_tier_status
.
1 Multi_Tier_Pool online on          active
.
.
IBM_Storwize:ITS0_V5000:superuser>
```

Run the **svcinfn lsvdisk** command to show all configured volumes within your Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, as shown in Example 9-5. For this example, we are only interested in a single volume.

Example 9-5 Show all configured volumes

```
IBM_Storwize:ITS0_V5000:superuser>svcinfn lsvdisk
id name                IO_group_id status mdisk_grp_id mdisk_grp_name capacity
0 Volume001            0          online 1           Multi_Tier_Pool 5.00GB
IBM_Storwize:ITS0_V5000:superuser>
```

To disable Easy Tier on single volumes, run the **svctask chvdisk -easytier off volume name** command, as shown in Example 9-6.

Example 9-6 Disable Easy Tier on a single volume

```
IBM_Storwize:ITS0_V5000:superuser>svctask chvdisk -easytier off Volume001
IBM_Storwize:ITS0_V5000:superuser>
```

This command disables Easy Tier on all copies of the volume. Example 9-7 shows Easy Tier turned off for copy 0 even if Easy Tier is still enabled on the storage pool. The status for copy 0 changed to measured because the pool is still actively measuring the I/O on the volume.

Example 9-7 Easy Tier that is disabled

```
IBM_Storwize:ITS0_V5000:superuser>svcinfn lsvdisk Volume001
id 0
name Volume001
IO_group_name io_grp0
status online
mdisk_grp_id 1
mdisk_grp_name Multi_Tier_Pool
capacity 5.00GB
type striped
throttling 0
preferred_node_id 2
parent_mdisk_grp_id 1
parent_mdisk_grp_name Multi_Tier_Pool
```

```

copy_id 0
status online
mdisk_grp_id 1
mdisk_grp_name Multi_Tier_Pool
fast_write_state empty
used_capacity 5.00GB
real_capacity 5.00GB
free_capacity 0.00MB
overallocation 100
easy_tier off
easy_tier_status measured
tier ssd
tier_capacity 1.00GB
tier enterprise
tier_capacity 4.00GB
tier nearline
tier_capacity 0.00MB
compressed_copy no
uncompressed_used_capacity 5.00GB
parent_mdisk_grp_id 1
parent_mdisk_grp_name Multi_Tier_Pool
IBM_Storwize:ITS0_V5000:superuser>

```

To enable Easy Tier on a volume, run the **svctask chvdisk -easytier on *volume name*** command (as shown in Example 9-8). Easy Tier changes back to on (as shown in Example 9-9). The copy 0 status also changed back to active.

Example 9-8 Easy Tier enabled

```

IBM_Storwize:ITS0_V5000:superuser>svctask chvdisk -easytier on Volume001
IBM_Storwize:ITS0_V5000:superuser>

```

Example 9-9 Easy Tier on single volume enabled

```

IBM_Storwize:ITS0_V5000:superuser>svcinfd lsvdisk Volume001
id 0
name Volume001
IO_group_id 0
IO_group_name io_grp0
status online
mdisk_grp_id 1
mdisk_grp_name Multi_Tier_Pool
capacity 5.00GB
parent_mdisk_grp_id 1
parent_mdisk_grp_name Multi_Tier_Pool

copy_id 0
status online
mdisk_grp_id 1
mdisk_grp_name Multi_Tier_Pool
type striped
mdisk_id
mdisk_name
used_capacity 5.00GB
real_capacity 5.00GB

```

```
free_capacity 0.00MB
overallocation 100
easy_tier on
easy_tier_status active
tier ssd
tier_capacity 1.00GB
tier enterprise
tier_capacity 4.00GB
tier nearline
tier_capacity 0.00MB
compressed_copy no
uncompressed_used_capacity 5.00GB
parent_mdisk_grp_id 1
parent_mdisk_grp_name Multi_Tier_Pool
IBM_Storwize:ITS0_V5000:superuser>
```

9.2.16 IBM Storage Tier Advisor Tool

IBM Storage Tier Advisor Tool (STAT) is a Microsoft Windows console tool. If you run Easy Tier in measure mode, the tool analyzes the extents and captures I/O profiles to estimate how much benefit you can derive from implementing Easy Tier Automatic Data Placement with additional MDisk tiers. If Automatic Data Placement Mode is already active, the analysis also includes an overview of migrated hot data and advice about whether you can derive any benefit from adding more Flash or Enterprise drives, for example.

The output provides a graphical representation of the performance data that is collected by Easy Tier over a 24-hour operational cycle.

The tool comes packaged as an International Organization for Standardization (ISO) file, which needs to be extracted to a temporary folder. The STAT can be downloaded from the following link:

<https://ibm.biz/BdEfrX>

9.2.17 Processing heat log files

IBM Storage Tier Advisor Tool takes input from the dpa_heat log file and produces an HTML file that contains the report. Download the heat_log file, as described in 9.2.14, “Downloading Easy Tier I/O measurements” on page 425, and save it to the hard disk drive (HDD) of a Windows system.

On Windows navigate to **Start** → **Run**, enter cmd, and then click **OK** to open a command prompt.

Typically, the tool is installed in the C:\Program Files\IBM\STAT directory. The command to create the index and other data files is has the following parameters:

```
C:\Program Files\IBM\STAT>STAT.exe -o c:\directory_where_you_want_the_output_to_go
c:\location_of_dpa_heat_data_file
```

Example 9-10 shows the command to create the report and the message that is displayed when it is successfully generated.

Example 9-10 Generate the HTML file

```
C:\EasyTier>STAT.exe -o C:\EasyTier C:\StorwizeV5000_Logs\dpa_heat.31G00KV-1.101
```

209.131801.data

CMUA00019I The STAT.exe command has completed.

C:\EasyTier>

If you do not specify -o c:\directory_where_you_want_the_output_to_go, the output goes to the directory of the STAT.exe file.

IBM Storage Tier Advisor Tool creates a set of HTML files. Browse to the directory where you directed the output file and locate the file that is named index.html. Open the file by using your browser to view the report.

9.3 Thin provisioning

In a shared storage environment, *thin provisioning* is a method for optimizing the usage of available storage. It relies on allocating blocks of data on demand versus the traditional method of allocating all of the blocks up front. This methodology eliminates almost all white space, which helps avoid the poor usage rates (often as low as 10%) that occur in the traditional storage allocation method. Traditionally, large pools of storage capacity are allocated to individual servers but remain unused (not written to).

Thin provisioning presents more storage space to the hosts or servers that are connected to the storage system than is available on the storage system. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support this capability for Fibre Channel (FC) and Internet Small Computer System Interface (iSCSI) provisioned volumes.

An example of thin provisioning is when a storage system contains 5000 GiB of usable storage capacity, but the storage administrator mapped volumes of 500 GiB each to 15 hosts. In this example, the storage administrator makes 7500 GiB of storage space visible to the hosts, even though the storage system has only 5000 GiB of usable space, as shown in Figure 9-23 on page 433. In this case, all 15 hosts cannot immediately use all 500 GiB that is provisioned to them. The storage administrator must monitor the system and add storage as needed.




Figure 9-23 Thin provisioning concept

9.3.1 Configuring a thin provisioned volume

Volumes can be configured as *thin-provisioned* or *fully allocated*. Thin-provisioned volumes are created with real and virtual capacities. You can still create volumes by using a striped, sequential, or image mode virtualization policy, as you can do with any other volume.

Real capacity defines how much disk space is allocated to a volume. *Virtual capacity* is the capacity of the volume that is reported to other Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 components (such as FlashCopy or remote copy) and to the hosts. For example, you can create a volume with real capacity of only 100 GiB, but virtual capacity of 1 terabyte (TiB). The actual space used by the volume on Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 is 100 GiB, but hosts see a 1 TiB volume.

A directory maps the virtual address space to the real address space. The directory and the user data share the real capacity.

Thin-provisioned volumes are available in two operating modes:

- Autoexpand
- Non-autoexpand

You can switch the mode at any time. If you select the autoexpand feature, the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 automatically add a fixed amount of more real capacity to the thin volume as required. Therefore, the autoexpand feature attempts to maintain a fixed amount of unused real capacity for the volume.

This amount is known as the *contingency capacity*. The contingency capacity is initially set to the real capacity that is assigned when the volume is created. If the user modifies the real capacity, the contingency capacity is reset to be the difference between the used capacity and real capacity.

A volume that is created without the autoexpand feature, and therefore has a zero contingency capacity, goes offline when the real capacity is used and the volume must expand.

Warning threshold: Enable the warning threshold, by using email or a Simple Network Management Protocol (SNMP) trap, when you work with thin-provisioned volumes. You can enable the warning threshold on the volume, and on the storage pool side, especially when you do not use the autoexpand mode. Otherwise, the thin volume goes offline if it runs out of space.

Autoexpand mode does not cause real capacity to grow much beyond the virtual capacity. The real capacity can be manually expanded to more than the maximum that is required by the current virtual capacity and the contingency capacity is recalculated.

A thin-provisioned volume can be converted non-disruptively to a fully allocated volume, or vice versa, by using the volume mirroring function. For example, you can add a thin-provisioned copy to a fully allocated primary volume, and then remove the fully allocated copy from the volume after they are synchronized.

The fully allocated to thin-provisioned migration procedure uses a zero-detection algorithm, so that grains that contain all zeros do not cause any real capacity to be used. Usually, Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are supposed to detect zeros on the volume, so you must use software on the host side to write zeros to all unused space on the disk or file system.

Space allocation

When a thin-provisioned volume is created, a small amount of the real capacity is used for initial metadata. Write I/Os to the grains of the thin volume (that were not previously written to) cause grains of the real capacity to be used to store metadata and user data. Write I/Os to the grains (that were previously written to) update the grain where data was previously written.

Grain definition: The grain is defined when the volume is created, and can be 32 KiB, 64 KiB, 128 KiB, or 256 KiB.

Smaller granularities can save more space, but they have larger directories. When you use thin-provisioning with FlashCopy, specify the same grain size for the thin-provisioned volume and FlashCopy.

To create a thin-provisioned volume from the dynamic menu, complete the following steps:

1. Navigate to **Volumes** → **Volumes** and click **Create Volumes** as shown in Figure 9-24 on page 435.

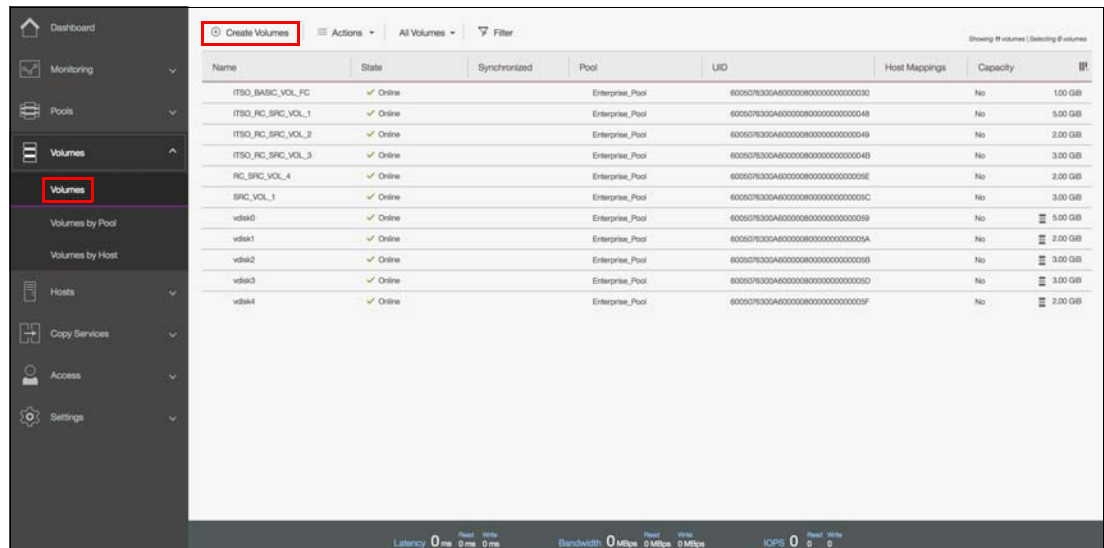


Figure 9-24 Volumes panel

2. Select the **Custom** tab. Specify the volume capacity, the type of the capacity saving and the name of the volume. By selecting **Thin-provisioned** as the capacity saving method, additional thin-provisioning parameters such as real capacity, autoexpand, warning threshold and grain size are displayed as shown in Figure 9-25 on page 436. Fill the necessary information for a single volume or click the **+** icon next to the volume name for multiple volumes and click **Create**.

Create Volumes

Basic Mirrored **Custom**

Volume Details

Quantity: Capacity: Capacity savings: Name:

Thin Provisioning

Real capacity: Automatically expand: ☒ Enabled Warning threshold: ☒ Enabled % of Virtual capacity Thin-Provisioned Grain Size:

General

Cache mode: OpenVMS UDID:

Figure 9-25 Creating a thin-provisioned volume

Note: A thin-provisioned volume can also be created by using the Basic or the Mirrored tabs within the Create Volumes panel, but you can only custom the thin-provisioning parameters through the customized volume creation. If your system uses HyperSwap topology the mirrored tab is replaced by the HyperSwap tab.

9.3.2 Performance considerations

Thin-provisioned volumes save capacity only if the host server does not write to whole volumes. Whether the thin-provisioned volume works well partly depends on how the file system allocated the space. Some file systems, for example, New Technology File System (NTFS), write to the whole volume before overwriting deleted files. Other file systems reuse space in preference to allocating new space.

File system problems can be moderated by tools, such as defrag, or by managing storage by using host Logical Volume Managers (LVMs). The thin-provisioned volume also depends on how applications use the file system. For example, some applications delete log files only when the file system is nearly full.

Important: Do not use defrag on thin-provisioned volumes. The defragmentation process can write data to different areas of a volume, which can cause a thin-provisioned volume to grow up to its virtual size.

There is no recommendation for thin-provisioned volumes. As explained previously, the performance of thin-provisioned volumes depends on what is used in the particular environment. For the best performance, use fully allocated volumes rather than thin-provisioned volumes.

9.3.3 Limitations of virtual capacity

A few factors (extent and grain size) limit the virtual capacity of thin-provisioned volumes beyond the factors that limit the capacity of regular volumes. Table 9-2 shows the maximum thin provisioned volume virtual capacities for an extent size.

Table 9-2 Maximum thin provisioned volume virtual capacities for an extent size

Extent size in megabytes (MB)	Maximum volume real capacity in gigabytes (GB)	Maximum thin virtual capacity in GB
16	2,048	2,000
32	4,096	4,000
64	8,192	8,000
128	16,384	16,000
256	32,768	32,000
512	65,536	65,000
1,024	131,072	130,000
2,048	262,144	260,000
4,096	262,144	262,144
8,192	262,144	262,144

Table 9-3 shows the maximum thin-provisioned volume virtual capacities for a grain size.

Table 9-3 Maximum thin volume virtual capacities for a grain size

Grain size in KiB	Maximum thin virtual capacity in GiB
32	260,000
64	520,000
128	1,040,000
256	2,080,000

9.4 Real-time Compression Software

The IBM Real-time Compression Software that is embedded in the controller firmware addresses the requirements for primary storage data reduction, including performance. It does so by using a purpose-built technology, called *Real-time Compression*, that uses the Random Access Compression Engine (RACE). It offers the following benefits:

- Compression for active primary data

IBM Real-time Compression can be used with active primary data. Therefore, it supports workloads that are not candidates for compression in other solutions. The solution supports online compression of existing data. Storage administrators can regain free disk space in an existing storage system without requiring users to clean up or archive data.

This configuration significantly enhances the value of existing storage assets and the benefits to the business are immediate. The capital expense of upgrading or expanding the storage system is delayed.

- Compression for replicated or mirrored data

Remote volume copies can be compressed, in addition to the volumes at the primary storage tier. This process reduces storage requirements in Metro Mirror and Global Mirror destination volumes as well.

- No changes to the existing environment are required

IBM Real-time Compression is part of the storage system. It was designed to be implemented without changes to applications, hosts, networks, fabrics, or external storage systems. The solution is not apparent to hosts, so users and applications continue to work non-disruptively.

- Overall savings in operational expenses

More data is stored in a rack space, so fewer storage expansion enclosures are required to store a data set. A reduced rack space has the following benefits:

- Reduced power and cooling requirements. More data is stored in a system, which requires less power and cooling per gigabyte or used capacity.
- Reduced software licensing for more functions in the system. More data that is stored per enclosure reduces the overall spending on licensing.

- Disk space savings are immediate

The space reduction occurs when the host writes the data. This process is unlike other compression solutions in which some or all of the reduction is performed only after running a post-process compression batch job.

9.4.1 Common use cases

This section addresses the most common use cases for implementing compression:

- General-purpose volumes
- Databases
- Virtualized infrastructures
- Log server data stores

For additional information on how to estimate compression ratios for each of the listed items, see 9.4.9, “Comprestimator” on page 448.

General-purpose volumes

Most general-purpose volumes are used for highly compressible data types, such as home directories, CAD/CAM, oil and gas geo-seismic data and log data. Storing such types of data in compressed volumes provides immediate capacity reduction to the overall used space. More space can be provided to users without any change to the environment.

Many file types can be stored in general-purpose servers. However, for practical information, the estimated compression ratios are based on actual field experience.

File systems that contain audio, video files, and compressed files are not good candidates for compression. The overall capacity savings on these file types are minimal.

Databases

Database information is stored in table space files. It is common to observe high compression ratios in database volumes. Examples of databases that can greatly benefit from Real-Time Compression are IBM DB2, Oracle and Microsoft SQL Server.

Important: Some databases offer optional built-in compression. Generally, do not compress already compressed database files.

Virtualized infrastructures

The proliferation of open systems virtualization in the market has increased the use of storage space, with more virtual server images and backups kept online. The use of compression reduces the storage requirements at the source.

Examples of virtualization solutions that can greatly benefit from Real-time Compression are VMware, Microsoft Hyper-V, and KVM.

Tip: Virtual machines with file systems that contain compressed files are not good candidates for compression, as described in “Databases”.

Log server data stores

Logs are a critical part for any information technology (IT) department in any organization. Log aggregates or syslog servers are a central point for the administrators, and immediate access and a smooth work process is necessary. Log server data stores are good candidates for Real-time Compression.

9.4.2 Real-time Compression concepts

RACE technology is based on over 70 patents that are not primarily about compression. Instead, they define how to make industry-standard Lempel-Ziv (LZ) compression of primary storage operate in real-time and allow random access. The primary intellectual property behind this is the RACE engine.

At a high level, the IBM RACE component compresses data that is written into the storage system dynamically. This compression occurs transparently, so Fibre Channel and iSCSI connected hosts are not aware of the compression. RACE is an online compression technology, which means that each host write is compressed as it passes to the disks. This technique has a clear benefit over other compression technologies that are post-processing based.

Those technologies do not provide immediate capacity savings. Therefore, they are not a good fit for primary storage workloads, such as databases and active data set applications.

RACE is based on the Lempel-Ziv lossless data compression algorithm and operates using a real-time method. When a host sends a write request, it is acknowledged by the write cache of the system and then staged to the storage pool. As part of its staging, it passes through the compression engine and is then stored in compressed format into the storage pool. Therefore, writes are acknowledged immediately after they are received by the write cache, with compression occurring as part of the staging to internal or external physical storage.

Capacity is saved when the data is written by the host because the host writes are smaller when they are written to the storage pool. IBM Real-time Compression is a self-tuning solution. It is adapting to the workload that runs on the system at any particular moment.

9.4.3 Random Access Compression Engine

To understand why RACE is unique, you need to review the traditional compression techniques. This description is not about the compression algorithm itself, that is, how the data structure is reduced in size mathematically. Rather, the description is about how the data is laid out within the resulting compressed output.

Compression utilities

Compression is probably most known to users because of the widespread use of compression utilities. At a high level, these utilities take a file as their input and parse the data by using a sliding window technique. Repetitions of data are detected within the sliding window history, most often 32 KiB. Repetitions outside of the window cannot be referenced. Therefore, the file cannot be reduced in size unless data is repeated when the window “slides” to the next 32 KiB slot.

Figure 9-26 shows compression that uses a sliding window, where the first two repetitions of the string “ABCD” fall within the same compression window, and can therefore be compressed by using the same dictionary. The third repetition of the string falls outside of this window, and therefore cannot be compressed by using the same compression dictionary as the first two repetitions, reducing the overall achieved compression ratio.




Figure 9-26 Compression that uses a sliding window

Traditional data compression in storage systems

The traditional approach taken to implement data compression in storage systems is an extension of how compression works in the previously mentioned compression utilities. Similar to compression utilities, the incoming data is broken into fixed chunks, and then each chunk is compressed and extracted independently.

However, there are drawbacks to this approach. An update to a chunk requires a read of the chunk followed by a recompression of the chunk to include the update. The larger the chunk size chosen, the heavier the I/O penalty to recompress the chunk. If a small chunk size is chosen, the compression ratio is reduced because the repetition detection potential is reduced.

Figure 9-27 shows an example of how the data is broken into fixed-size chunks (in the upper-left side of the figure). It also shows how each chunk gets compressed independently into variable length compressed chunks (in the upper-right side of the figure). The resulting compressed chunks are stored sequentially in the compressed output.

Although this approach is an evolution from compression utilities, it is limited to low-performance use cases. This limitation is mainly because it does not provide real random access to the data.

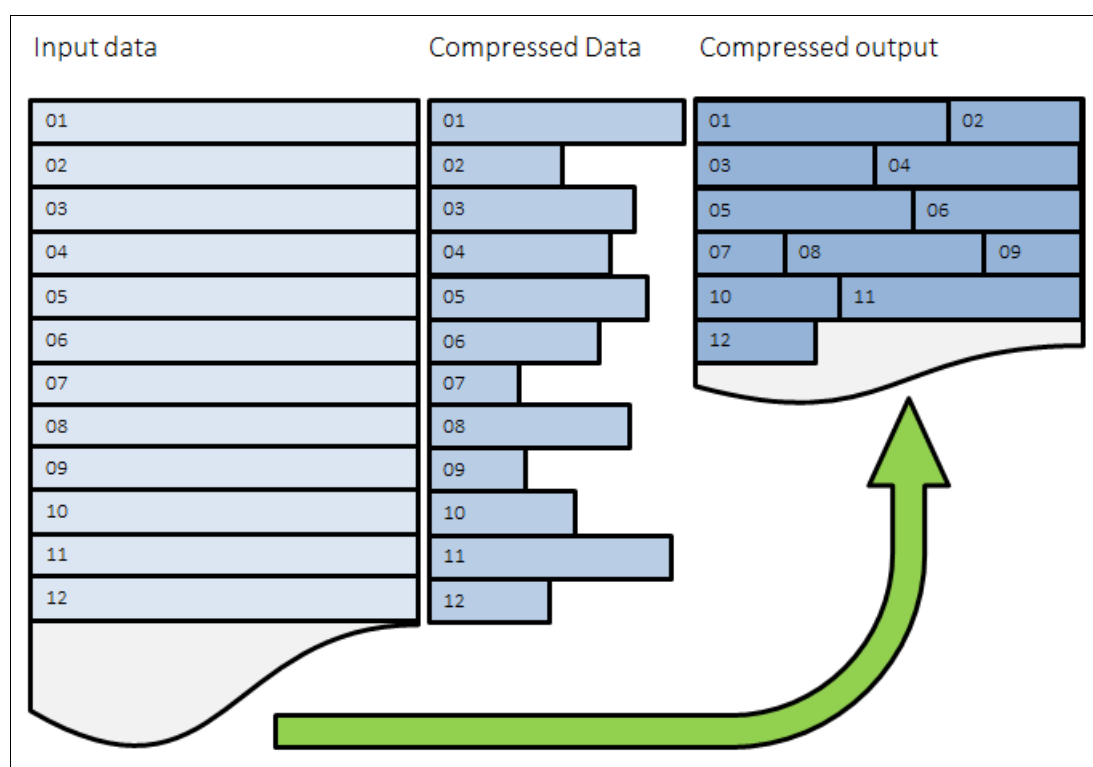


Figure 9-27 Traditional data compression in storage systems

Random Access Compression Engine

The IBM patented RACE implements an inverted approach when compared to traditional approaches to compression. RACE uses variable-size chunks for the input, and produces fixed-size chunks for the output.

This method enables an efficient and consistent way to index the compressed data because it is stored in fixed-size containers (Figure 9-28 on page 442).

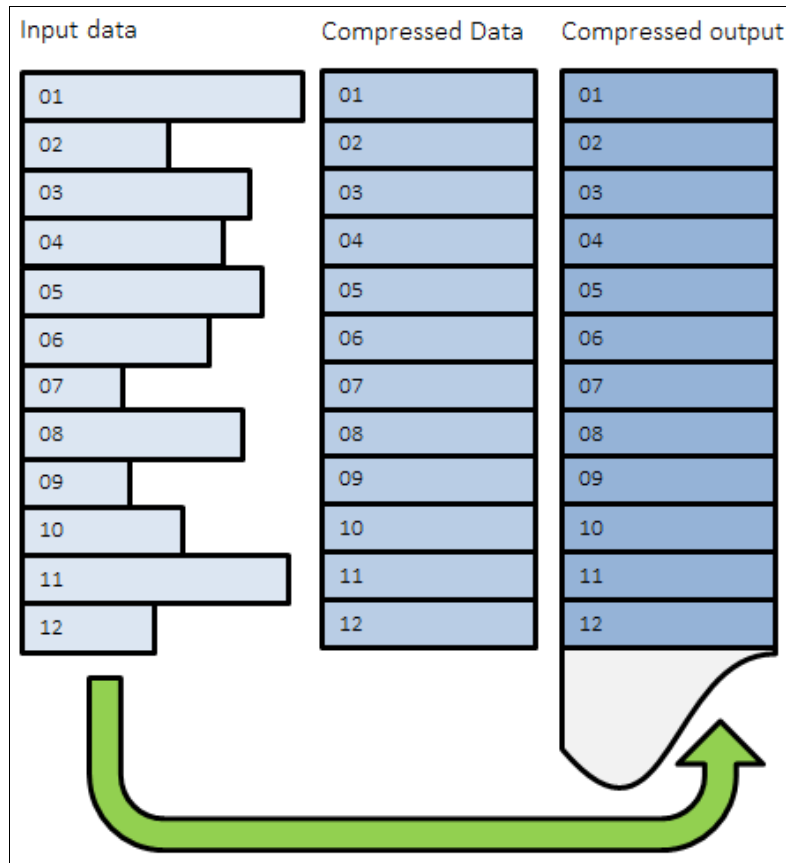


Figure 9-28 Random Access Compression

Location-based compression

Both compression utilities and traditional storage systems compression compress data by finding repetitions of bytes within the chunk that is being compressed. The compression ratio of this chunk depends on how many repetitions can be detected within it. The number of repetitions is affected by how much the bytes stored in the chunk are related to each other. The relation between bytes is driven by the format of the object. For example, an office document might contain textual information and an embedded drawing.

Because the chunking of the file is arbitrary, it has no concept of how the data is laid out within the document. Therefore, a compressed chunk can be a mixture of the textual information and part of the drawing. This process yields a lower compression ratio because the different data types mixed together cause a suboptimal dictionary of repetitions. Which means that fewer repetitions can be detected, because a repetition of bytes in a text object is unlikely to be found in a drawing.

This traditional approach to data compression is also called *location-based compression*. The data repetition detection is based on the location of data within the same chunk.

Predecide mechanism

Some data chunks have a higher compression ratio than others. Compressing some of the chunks saves little space, but still requires resources, such as processor (CPU) and memory. To avoid spending resources on incompressible data, and to provide the ability to use a different, more effective compression algorithm.

The chunks that are below a given compression ratio are skipped by the compression engine, saving CPU time and memory processing. Chunks that are not compressed with the main compression algorithm, but that still can be compressed well with the other, are marked and processed accordingly. The result might vary because predecide does not check the entire block, only a sample of it.

Figure 9-29 shows how the detection mechanism works.

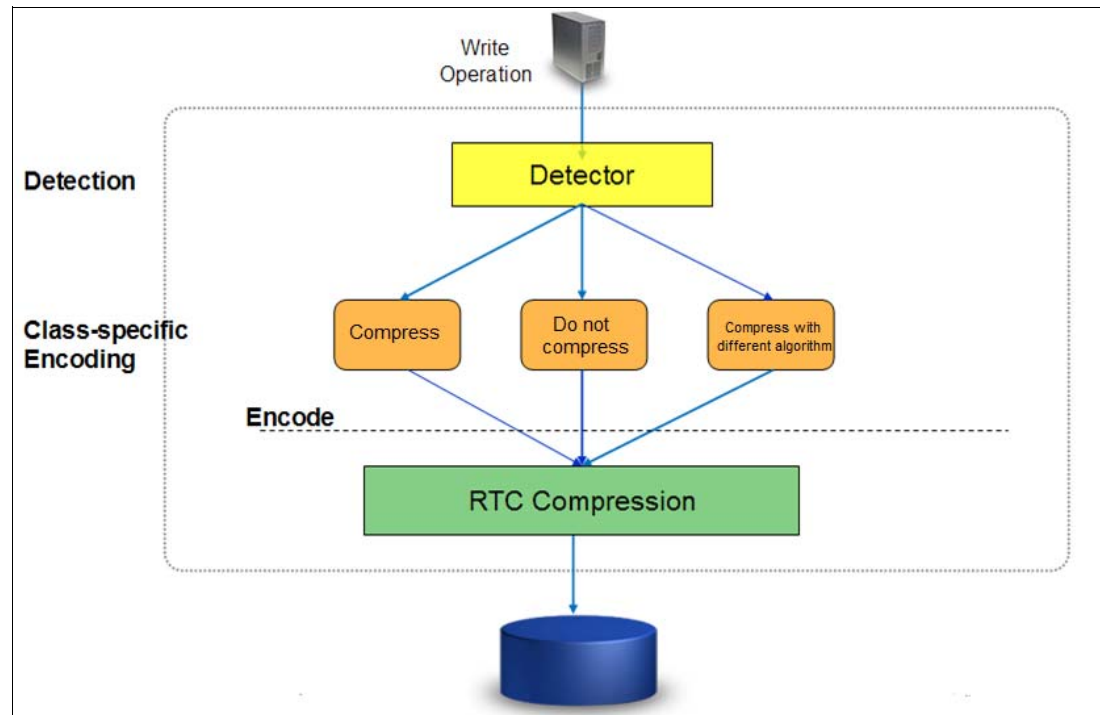


Figure 9-29 Detection mechanism

Temporal compression

RACE offers a technology leap beyond location-based compression, called *temporal compression*. When host writes arrive to RACE, they are compressed and fill up fixed size chunks, also named as *compressed blocks*. Multiple compressed writes can be aggregated into a single compressed block. A dictionary of the detected repetitions is stored within the compressed block.

When applications write new data or update existing data, it is typically sent from the host to the storage system as a series of writes. Because these writes are likely to originate from the same application and be of the same data type, more repetitions are usually detected by the compression algorithm. This type of data compression is called *temporal compression* because the data repetition detection is based on the time the data was written into the same compressed block.

Temporal compression adds the time dimension that is not available to other compression algorithms. It offers a higher compression ratio because the compressed data in a block represents a more homogeneous set of input data.

Figure 9-30 shows how three writes sent one after the other by a host end up in different chunks. They get compressed in different chunks because their location in the volume is not adjacent. This process yields a lower compression ratio because the same data must be compressed non-natively by using three separate dictionaries.




Figure 9-30 Location-based compression

When the same three writes are sent through RACE, as shown on Figure 9-31, the writes are compressed together by using a single dictionary. This process yields a higher compression ratio than location-based compression (Figure 9-31).




Figure 9-31 Temporal compression

9.4.4 Random Access Compression Engine in stack

RACE technology is implemented into the Storwize thin provisioning layer and is an organic part of the stack. Compression is transparently integrated with existing system management design. All of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 advanced features are supported on compressed volumes. You can create, delete, migrate, map (assign), and unmap (unassign) a compressed volume as though it were a fully allocated volume.

In addition, you can use Real-time Compression along with Easy Tier on the same volumes. This compression method provides non disruptive conversion between compressed and decompressed volumes. This conversion provides a uniform user-experience and eliminates the need for special procedures when dealing with compressed volumes.

9.4.5 Data write flow

When a host sends a write request to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, it reaches the upper cache layer. The host is immediately sent an acknowledgment of its I/Os.

When the upper cache layer destages to the RACE, the I/Os are sent to the thin-provisioning layer. They are then sent to RACE, and if necessary, to the original host write or writes. The metadata that holds the index of the compressed volume is updated if needed, and is compressed as well.

9.4.6 Data read flow

When a host sends a read request to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 for compressed data, it is forwarded directly to the Real-time Compression component:

- ▶ If the Real-time Compression component contains the requested data, cache replies to the host with the requested data without having to read the data from the lower-level cache or disk.
- ▶ If the Real-time Compression component does not contain the requested data, the request is forwarded to the Lenovo Storage V-series system lower-level cache.
- ▶ If the lower-level cache contains the requested data, it is sent up the stack and returned to the host without accessing the storage.
- ▶ If the lower-level cache does not contain the requested data, it sends a read request to the storage for the requested data.

9.4.7 Compression of existing data

In addition to compressing data in real time, you can also compress existing data sets (convert volume to compressed). To do so, you must change the capacity savings settings of the volume:

1. Right-click a particular volume and select **Modify Capacity Settings**, as shown in Figure 9-32.

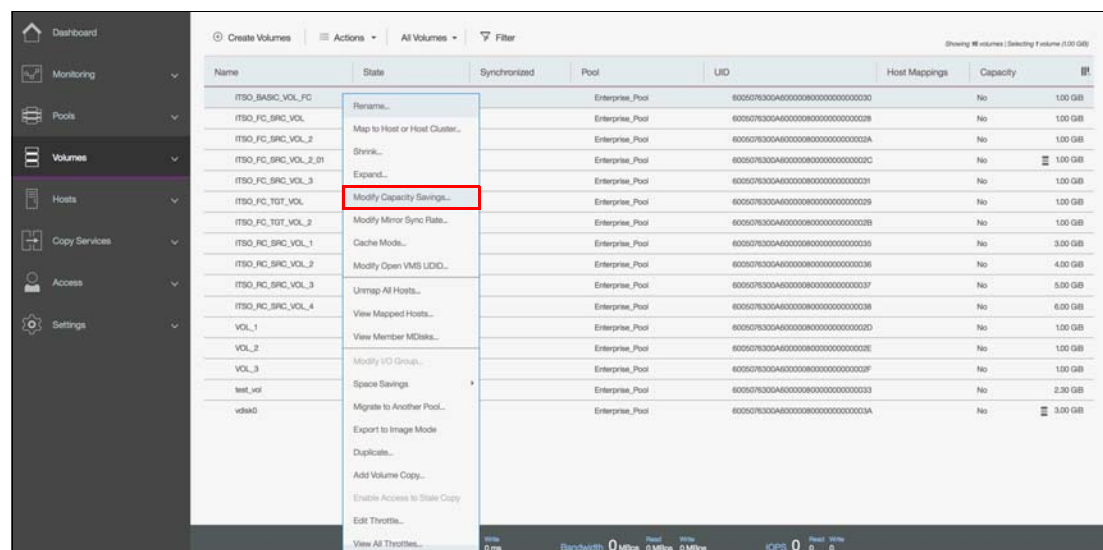


Figure 9-32 Modifying Capacity Settings

2. In the menu select **Compression** as the Capacity Savings option, as shown in Figure 9-33.

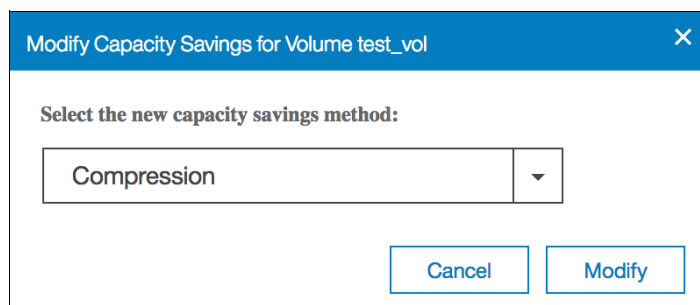


Figure 9-33 Selecting Capacity Setting

3. After the copies are fully synchronized, the original volume copy is deleted automatically.

As a result, you have compressed data on the existing volume. This process is non disruptive, so the data remains online and accessible by applications and users.

With virtualization of external storage systems, the ability to compress already stored data significantly enhances and accelerates the benefit to users. It enables them to see a tremendous return on their Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 investment. On initial purchase of a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 with Real-time Compression, customers can defer their purchase of new storage. As new storage needs to be acquired, IT purchases a lower amount of the required storage before compression.

Important: Remember that Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 will reserve some of its resources like CPU cores and RAM after you create just one compressed volume or volume copy. This setting can affect your system performance if you do not plan accordingly in advance.

9.4.8 Configuring compressed volumes

To use compression on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, licensing is required. With the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, Real-time Compression is licensed by enclosure. Every enclosure that works with compression needs to be licensed.

There are two ways of creating a compressed volume: Basic and Custom.

To create a compressed volume using Basic option, navigate to **Volumes** → **Volumes** and click **Create Volumes**. Select the **Basic** tab and fill the required information as shown in Figure 9-34. Click **Create** to finish the compressed volume creation.

The screenshot shows the 'Create Volumes' dialog box with the 'Basic' tab selected. The dialog has three tabs: 'Basic', 'Mirrored', and 'Custom'. Below the tabs, it says 'Create a preset volume with all the basic features.'

Pool: A dropdown menu shows 'Multi_Tier_Pool'. To its right is a progress bar labeled 'Total 4.24 TiB'.

Volume Details: This section contains four fields:

- Quantity:** A spinner box set to '1'.
- Capacity:** A text box with '10', a unit dropdown set to 'GiB', and a small 'v' icon.
- Capacity savings:** A dropdown menu set to 'Compressed'.
- Name:** A text box containing 'ITSO_Volume' and a '+' icon.

I/O group: A dropdown menu set to 'Automatic'.

Summary: A section with a list of details:

- 1 volume
- Volume name: ITSO_Volume
- 1 volume in pool Multi_Tier_Pool
- Caching I/O group: Automatic
- Accessible I/O group: Automatic
- Total real capacity: 204.80 MiB
- Total virtual capacity: 10.00 GiB

At the bottom right, there are three buttons: 'Cancel', 'Create and Map', and 'Create'. The 'Create' button is highlighted with a red rectangle.

Figure 9-34 Creating Basic compressed volume

To create a compressed volume using the Custom option, navigate to **Volumes** → **Volumes** and click **Create Volumes**. Select the **Custom** tab and fill the required information under as shown in Figure 9-35. Click **Create** to finish the compressed volume creation

The screenshot shows the 'Create Volumes' dialog box with the 'Custom' tab selected. The 'Volume copy type' is set to 'None'. The 'Pool' is set to 'Multi_Tier_Pool'. The 'Caching I/O group' is set to 'Automatic'. The 'Preferred node' is set to 'Automatic'. The 'Accessible I/O groups' is set to 'Only the caching I/O group'. The 'Volume Details' section shows 'Quantity' as 1, 'Capacity' as 10 GiB, 'Capacity savings' as Compressed, and 'Name' as ITSO_Volume. The 'Compressed' section shows 'Real capacity' as 2, '% of Virtual capacity' selected, 'Automatically expand' as Enabled, and 'Warning threshold' as 80 % of Virtual capacity. The 'Create' button is highlighted with a red box.

Figure 9-35 Creating Advanced compressed volume

If the volume being created through the Custom tab is the first compressed volume in the environment, a warning message will be displayed as shown in Figure 9-36.

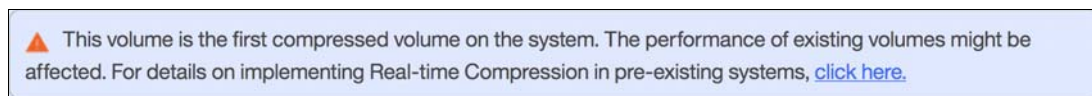


Figure 9-36 First compressed volume warning

9.4.9 Comprestimator

The built-in Comprestimator is a command-line utility that can be used to estimate an expected compression rate for a specific volume.

Comprestimator uses advanced mathematical and statistical algorithms to perform the sampling and analysis process in a short and efficient way. The utility also displays its

accuracy level by showing the maximum error range of the results achieved based on the formulas that it uses.

The following commands are available:

- The **analyzevdisk** command provides an option to analyze a single volume.

Usage: **analyzevdisk** <volume ID>

Example: **analyzevdisk 0**

This command can be canceled by running the **analyzevdisk <volume ID> -cancel** command.

- The **lsvdiskanalysis** command provides a list and the status of the volumes. Some of them can be analyzed already, some of them not yet. The command can either be used for all volumes on the system or it can be used per volume, similar to **lsvdisk**. See Example 9-11.

Example 9-11 Example of the command run over one volume with ID 0

```
IBM_2078:ITS0 Gen2:superuser>lsvdiskanalysis 0
id 0
name SQL_Data0
state estimated
started_time 151012104343
analysis_time 151012104353
capacity 300.00GB
thin_size 290.85GB
thin_savings 9.15GB
thin_savings_ratio 3.05
compressed_size 141.58GB
compression_savings 149.26GB
compression_savings_ratio 51.32
total_savings 158.42GB
total_savings_ratio 52.80
accuracy 4.97
```

The **state** parameter can have the following values:

- **idle**. Was never estimated and not currently scheduled.
 - **scheduled**. Volume is queued for estimation, and will be processed based on lowest volume ID first.
 - **active**. Volume is being analyzed.
 - **canceled**. Volume was requested to cancel an active analysis, but the analysis was not canceled yet.
 - **estimated**. Volume was analyzed and results show the expected savings of thin provisioning and compression.
 - **sparse**. Volume was analyzed but Comprestimator could not find enough nonzero samples to establish a good estimation.
 - **compression_savings_ratio**. The compression saving ratio is the estimated amount of space that can be saved on the storage in the frame of this specific volume expressed as a percentage.
- The **analyzevdiskbysystem** command provides an option to run Comprestimator on all volumes within the system. The analyzing process is nondisruptive and should not affect the system significantly. Analysis speed might vary due to the fullness of the volume, but should not take more than a few minutes per volume.

This command can be canceled by running the **analyzevdiskbysystem -cancel** command.

- The **lsvdiskanalysisprogress** command shows the progress of the Comprestimator analysis as shown in Example 9-12.

Example 9-12 Comprestimator progress

id	vdisk_count	pending_analysis	estimated_completion_time
0	45	12	151012154400

Copy services

In this chapter, we describe the copy services functions that are provided by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 storage systems, including FlashCopy, Remote Copy and HyperSwap. Copy services functions are useful for making data copies for backup, application test, recovery, and so on. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems make it easy to apply these functions to your environment through its intuitive graphical user interface (GUI).

This chapter includes the following topics:

- ▶ 10.1, “FlashCopy” on page 452
- ▶ 10.2, “FlashCopy functional overview” on page 456
- ▶ 10.3, “Implementing FlashCopy” on page 457
- ▶ 10.4, “Managing FlashCopy by using the GUI” on page 480
- ▶ 10.5, “Volume mirroring and migration options” on page 513
- ▶ 10.6, “Native IP replication” on page 515
- ▶ 10.7, “Remote Copy” on page 523
- ▶ 10.8, “Consistency protection for Remote and Global mirror” on page 549
- ▶ 10.9, “Remote Copy commands” on page 551
- ▶ 10.10, “Managing Remote Copy using the GUI” on page 559
- ▶ 10.11, “Troubleshooting remote copy” on page 590
- ▶ 10.12, “HyperSwap” on page 593

10.1 FlashCopy

By using the FlashCopy function, you can perform a *point-in-time copy* of one or more volumes. This section describes the inner workings of FlashCopy and provides details about its configuration and use.

You can use FlashCopy to help you solve critical and challenging business needs that require duplication of data of your source volume. Volumes can remain online and active while you create consistent copies of the data sets. Because the copy is performed at the block level, it operates below the host operating system and its cache. Therefore, the copy is not apparent to the host.

Important: Because FlashCopy operates at the block level below the host operating system and cache, those levels do need to be flushed for consistent FlashCopies.

While the FlashCopy operation is performed, the source volume is briefly halted to initialize the FlashCopy bitmap, and then input/output (I/O) can resume. Although several FlashCopy options require the data to be copied from the source to the target in the background, which can take time to complete, the resulting data on the target volume is presented so that the copy appears to complete immediately.

This process is performed by using a bitmap (or bit array), which tracks changes to the data after the FlashCopy is started, and an indirection layer, which enables data to be read from the source volume transparently.

10.1.1 Business requirements for FlashCopy

When you are deciding whether FlashCopy addresses your needs, you must adopt a combined business and technical view of the problems that you want to solve. First, determine the needs from a business perspective. Then, determine whether FlashCopy can address the technical needs of those business requirements.

The business applications for FlashCopy are wide-ranging. Common use cases for FlashCopy include, but are not limited to, the following examples:

- ▶ Rapidly creating consistent backups of dynamically changing data
- ▶ Rapidly creating consistent copies of production data to facilitate data movement or migration between hosts
- ▶ Rapidly creating copies of production data sets for application development and testing
- ▶ Rapidly creating copies of production data sets for auditing purposes and data mining
- ▶ Rapidly creating copies of production data sets for quality assurance

Regardless of your business needs, FlashCopy is flexible and offers a broad feature set, which makes it applicable to many scenarios.

10.1.2 Backup improvements with FlashCopy

FlashCopy does not reduce the time that it takes to perform a backup to traditional backup infrastructure. However, it can be used to minimize and, under certain conditions, eliminate application downtime that is associated with performing backups. FlashCopy can also transfer the resource usage of performing intensive backups from production systems.

After the FlashCopy is performed, the resulting image of the data can be backed up to tape, as though it were the source system. After the copy to tape is complete, the image data is redundant and the target volumes can be discarded. For time-limited applications, such as these examples, “no copy” or incremental FlashCopy is used most often. The use of these methods puts less load on your infrastructure.

When FlashCopy is used for backup purposes, the target data usually is managed as read-only at the operating system level. This approach provides extra security by ensuring that your target data was not modified and remains true to the source.

10.1.3 Restore with FlashCopy

FlashCopy can perform a restore from any existing FlashCopy mapping. Therefore, you can restore (or copy) from the target to the source of your regular FlashCopy relationships. When restoring data from FlashCopy, this method can be qualified as reversing the direction of the FlashCopy mappings.

This capability has the following benefits:

- ▶ There is no need to worry about pairing mistakes; you trigger a restore.
- ▶ The process appears instantaneous.
- ▶ You can maintain a pristine image of your data while you are restoring what was the primary data.

This approach can be used for various applications, such as recovering your production database application after an errant batch process that caused extensive damage.

Preferred practices: Although restoring from a FlashCopy is quicker than a traditional tape media restore, you must not use restoring from a FlashCopy as a substitute for good archiving practices. Instead, keep one to several iterations of your FlashCopies so that you can near-instantly recover your data from the most recent history, and keep your long-term archive as appropriate for your business.

In addition to the restore option, which copies the original blocks from the target volume to modified blocks on the source volume, the target can be used to perform a restore of individual files. To do that you need to make the target available on a host. We suggest that you do not make the target available to the source host, because seeing duplicates of disks causes problems for most host operating systems. Copy the files to the source using normal host data copy methods for your environment.

10.1.4 Moving and migrating data with FlashCopy

FlashCopy can be used to facilitate the movement or migration of data between hosts while minimizing downtime for applications. By using FlashCopy, application data can be copied from source volumes to new target volumes while applications remain online. After the volumes are fully copied and synchronized, the application can be brought down and then immediately brought back up on the new server that is accessing the new FlashCopy target volumes.

This method differs from the other migration methods, which are described later in this chapter. Common uses for this capability are host and back-end storage hardware refreshes.

10.1.5 Application testing with FlashCopy

It is often important to test a new version of an application or operating system that is using actual production data. This testing ensures the highest quality possible for your environment. FlashCopy makes this type of testing easy to accomplish without putting the production data at risk or requiring downtime to create a constant copy.

You create a FlashCopy of your source and use that for your testing. This copy is a duplicate of your production data down to the block level so that even physical disk identifiers are copied. Therefore, it is impossible for your applications to tell the difference.

10.1.6 Host and application considerations to ensure FlashCopy integrity

Because FlashCopy is at the block level, it is necessary to understand the interaction between your application and the host operating system. From a logical standpoint, it is easiest to think of these objects as “layers” that sit on top of one another. The application is the topmost layer, and beneath it is the operating system layer.

Both of these layers have various levels and methods of caching data to provide better speed. Because the FlashCopy sits below these layers, they are unaware of the cache at the application or operating system layers.

To ensure the integrity of the copy that is made, it is necessary to flush the host operating system and application cache for any outstanding reads or writes before the FlashCopy operation is performed. Failing to flush the host operating system and application cache produces what is referred to as a *crash consistent* copy.

The resulting copy requires the same type of recovery procedure, such as log replay and file system checks, that is required following a host crash. FlashCopies that are crash consistent often can be used following file system and application recovery procedures.

Various operating systems and applications provide facilities to stop I/O operations and ensure that all data is flushed from host cache. If these facilities are available, they can be used to prepare for a FlashCopy operation. When this type of facility is unavailable, the host cache must be flushed manually by quiescing the application and unmounting the file system or drives.

Preferred practice: From a practical standpoint, when you have an application that is backed by a database and you want to make a FlashCopy of that application’s data, it is sufficient in most cases to use the write-suspend method that is available in most modern databases, because the database maintains strict control over I/O.

This method is as opposed to flushing data from both the application and the backing database, which is always the suggested method because it is safer. However, this method can be used when facilities do not exist or your environment includes time sensitivity.

10.1.7 FlashCopy attributes

The FlashCopy function features the following attributes:

- ▶ The target is the time-zero copy of the source, which is known as *FlashCopy mapping targets*.
- ▶ FlashCopy produces an exact copy of the source volume, including any metadata that was written by the host operating system, logical volume manager, and applications.

- ▶ The source volume and target volume are available (almost) immediately following the FlashCopy operation.
- ▶ The source and target volumes must be the same “virtual” size.
- ▶ The source and target volumes must be on the same IBM Storwize for Lenovo system.
- ▶ The source and target volumes do not need to be in the same I/O Group or storage pool.
- ▶ The storage pool extent sizes can differ between the source and target.
- ▶ The source volumes can have up to 256 target volumes (Multiple Target FlashCopy).
- ▶ The target volumes can be the source volumes for other FlashCopy relationships (*cascaded FlashCopy*).
- ▶ Consistency groups are supported to enable FlashCopy across multiple volumes at the same time.
- ▶ Up to 255 FlashCopy consistency groups are supported per system.
- ▶ Up to 512 FlashCopy mappings can be placed in one consistency group.
- ▶ The target volume can be updated independently of the source volume.
- ▶ Bitmaps that are governing I/O redirection (I/O indirection layer) are maintained in both node canisters of the IBM Storwize for Lenovo I/O Group to prevent a single point of failure.
- ▶ FlashCopy mapping and Consistency Groups can be automatically withdrawn after the completion of the background copy.
- ▶ Thin-provisioned FlashCopy (or Snapshot in the graphical user interface (GUI)) use disk space only when updates are made to the source or target data, and not for the entire capacity of a volume copy.
- ▶ FlashCopy licensing is based on the virtual capacity of the source volumes.
- ▶ Incremental FlashCopy copies all of the data when you first start FlashCopy and then only the changes when you stop and start FlashCopy mapping again. Incremental FlashCopy can substantially reduce the time that is required to re-create an independent image.
- ▶ Reverse FlashCopy enables FlashCopy targets to become restore points for the source without breaking the FlashCopy relationship, and without having to wait for the original copy operation to complete.
- ▶ The maximum number of supported FlashCopy mappings is 4096 per clustered system.
- ▶ The size of the source and target volumes cannot be altered (increased or decreased) while a FlashCopy mapping is defined.

10.1.8 Reverse FlashCopy

Reverse FlashCopy enables FlashCopy targets to become restore points for the source without breaking the FlashCopy relationship, and without having to wait for the original copy operation to complete. It supports multiple targets (up to 256) and therefore multiple rollback points.

A key advantage of the Multiple Target Reverse FlashCopy function is that the reverse FlashCopy does not destroy the original target, which enables processes that are using the target, such as a tape backup, to continue uninterrupted.

Lenovo Storage V series system also provides the ability to create an optional copy of the source volume to be made before the reverse copy operation starts. This ability to restore back to the original source data can be useful for diagnostic purposes.

Complete the following steps to restore from an on-disk backup:

1. (Optional) Create a target volume (volume Z) and use FlashCopy to copy the production volume (volume X) onto the new target for later problem analysis.
2. Create a FlashCopy map with the backup to be restored (volume Y) or (volume W) as the source volume and volume X as the target volume, if this map does not exist.
3. Start the FlashCopy map (volume Y → volume X) with the -restore option to copy the backup data onto the production disk. If the -restore option is specified and no FlashCopy mapping exists, the command is ignored, which preserves your data integrity.

The production disk is instantly available with the backup data. Figure 10-1 shows an example of Reverse FlashCopy.

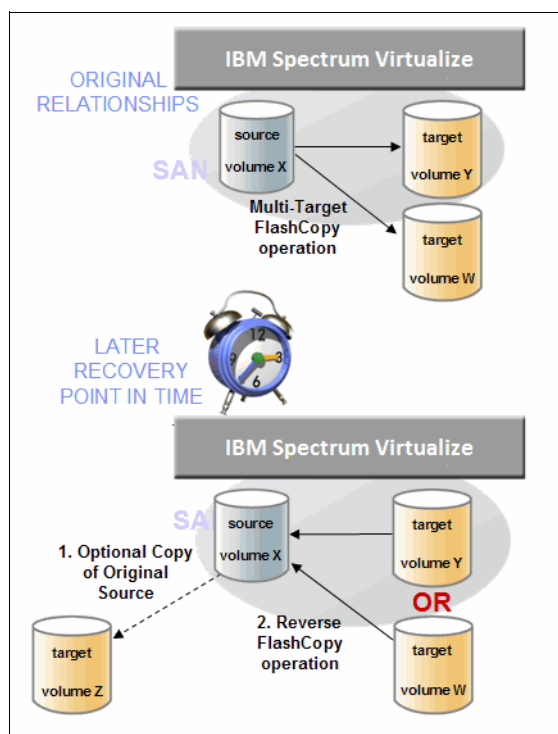


Figure 10-1 Reverse FlashCopy

Regardless of whether the initial FlashCopy map (volume X → volume Y) is incremental, the Reverse FlashCopy operation copies the modified data only.

Consistency Groups are reversed by creating a set of new reverse FlashCopy maps and adding them to a new reverse Consistency Group. Consistency Groups cannot contain more than one FlashCopy map with the same target volume.

10.2 FlashCopy functional overview

FlashCopy works by defining a FlashCopy mapping that consists of one source volume with one target volume. Multiple FlashCopy mappings (source-to-target relationships) can be defined, and point-in-time consistency can be maintained across multiple individual mappings by using Consistency Groups. For more information, see “Consistency Group with Multiple Target FlashCopy” on page 461.

Before you start a FlashCopy (regardless of the type and options specified), you must issue a **prestartfcmap** or **prestartfcconsistgrp** command, which puts the cache into write-through mode and provides a flushing of the I/O currently bound for your volume. After FlashCopy is started, an effective copy of a source volume to a target volume is created.

The content of the source volume is presented immediately on the target volume and the original content of the target volume is lost. This FlashCopy operation is also referred to as a *time-zero copy* (T0).

Tip: Rather than using **prestartfcmap** or **prestartfcconsistgrp**, you can also use the **-prep** parameter in the **startfcmap** or **startfcconsistgrp** command to prepare and start FlashCopy in one step.

The source and target volumes are available for use immediately after the FlashCopy operation. The FlashCopy operation creates a bitmap that is referenced and maintained to direct I/O requests within the source and target relationship. This bitmap is updated to reflect the active block locations as data is copied in the background from the source to the target, and updates are made to the source.

For more information about background copy, see 10.3.5, “Grains and the FlashCopy bitmap” on page 462.

Figure 10-2 shows the redirection of the host I/O toward the source volume and the target volume.

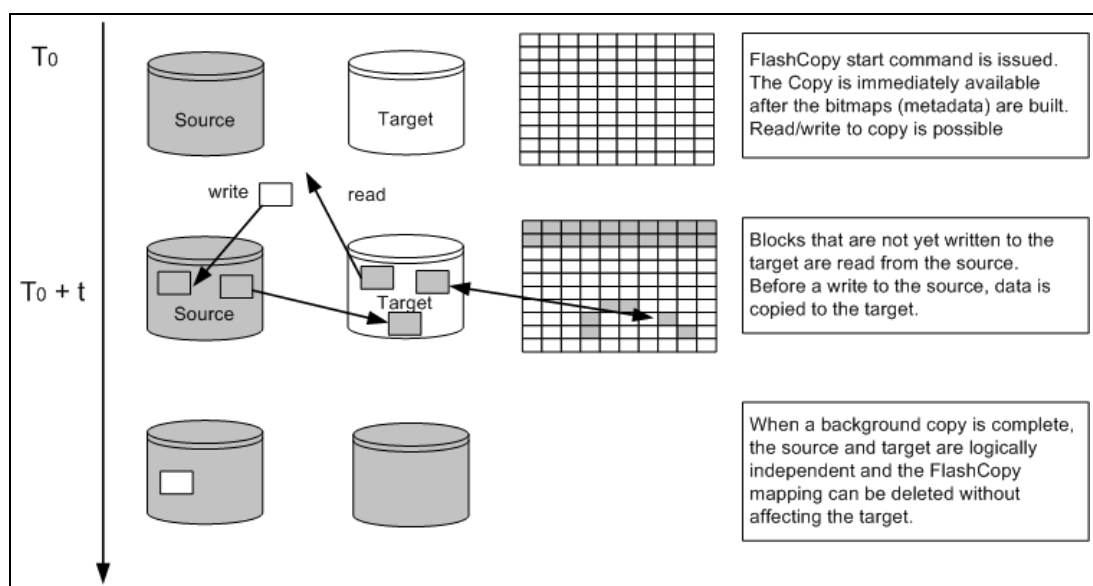


Figure 10-2 Redirection of host I/O

10.3 Implementing FlashCopy

This section describes how FlashCopy is implemented in the controller firmware running on Lenovo storage V series systems.

10.3.1 FlashCopy mappings

FlashCopy occurs between a source volume and a target volume. The source and target volumes must be the same size. The minimum granularity supports for FlashCopy is an entire volume. It is not possible to use FlashCopy to copy only part of a volume.

Important: As with any point-in-time copy technology, you are bound by operating system and application requirements for interdependent data and the restriction to an entire volume.

The source and target volumes must belong to the same Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, but they do not have to be in the same I/O Group or storage pool. FlashCopy associates a source volume to a target volume through FlashCopy mapping.

To become members of a FlashCopy mapping, source and target volumes must be the same size. Volumes that are members of a FlashCopy mapping cannot have their size increased or decreased while they are members of the FlashCopy mapping.

A *FlashCopy mapping* is the act of creating a relationship between a source volume and a target volume. FlashCopy mappings can be stand-alone or a member of a Consistency Group. You can perform the actions of preparing, starting, or stopping FlashCopy on either a stand-alone mapping or a Consistency Group.

Figure 10-3 shows the concept of FlashCopy mapping.




Figure 10-3 FlashCopy mapping

10.3.2 Multiple Target FlashCopy

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems support up to 256 target volumes from a single source volume. Each copy is managed by a unique mapping. Figure 10-4 on page 459 shows the Multiple Target FlashCopy implementation.

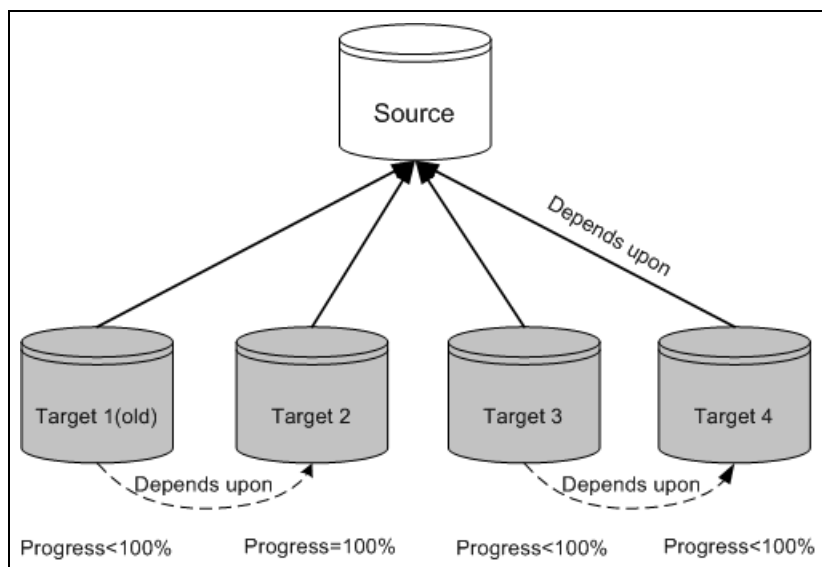


Figure 10-4 Multiple Target FlashCopy implementation

Figure 10-4 also shows four targets and mappings that are taken from a single source, along with their interdependencies. In this example, Target 1 is the oldest (as measured from the time that it was started) through to Target 4, which is the newest. The ordering is important because of how the data is copied when multiple target volumes are defined and because of the dependency chain that results.

A write to the source volume does not cause its data to be copied to all of the targets. Instead, it is copied to the newest target volume only (Target 4 in Figure 10-4). The older targets refer to new targets first before referring to the source.

From the point of view of an intermediate target disk (not the oldest or the newest), it treats the set of newer target volumes and the true source volume as a type of composite source. It treats all older volumes as a kind of target (and behaves like a source to them).

If the mapping for an intermediate target volume shows 100% progress, its target volume contains a complete set of data. In this case, mappings treat the set of newer target volumes (up to and including the 100% progress target) as a form of composite source. A dependency relationship exists between a particular target and all newer targets (up to and including a target that shows 100% progress) that share the source until all data is copied to this target and all older targets.

For more information about Multiple Target FlashCopy, see 10.3.6, “Interaction and dependency between multiple target FlashCopy mappings” on page 464.

10.3.3 Consistency Groups

Consistency Groups address the requirement to preserve point-in-time data consistency across multiple volumes for applications that include related data that spans multiple volumes. For these volumes, Consistency Groups maintain the integrity of the FlashCopy by ensuring that “dependent writes” are run in the application’s intended sequence.

When Consistency Groups are used, the FlashCopy commands are issued to the FlashCopy Consistency Group, which performs the operation on all FlashCopy mappings that are contained within the Consistency Group at the same time.

Figure 10-5 shows a Consistency Group that includes two FlashCopy mappings.

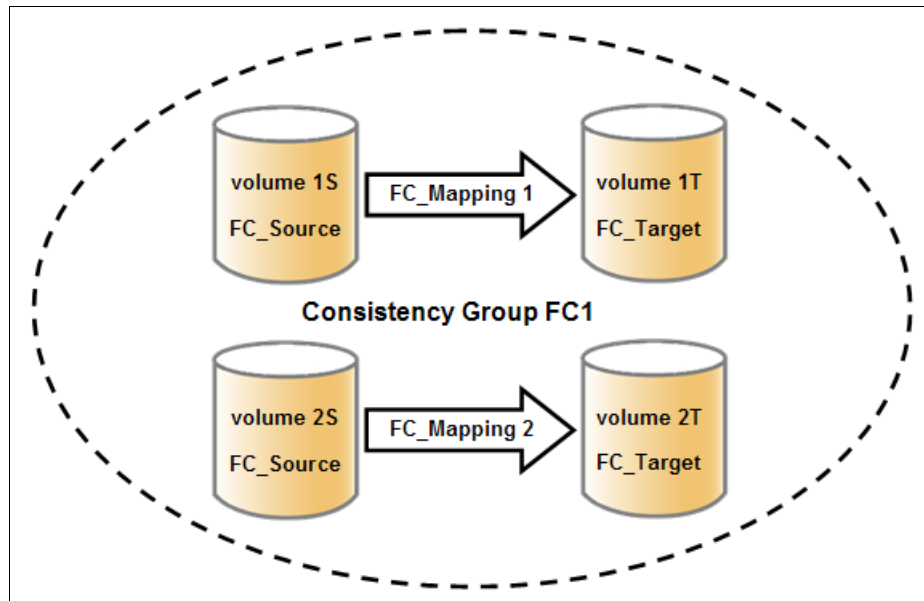


Figure 10-5 FlashCopy Consistency Group

Important: After an individual FlashCopy mapping is added to a Consistency Group, it can be managed as part of the group only. Operations, such as prepare, start, and stop, are no longer allowed on the individual mapping.

Dependent writes

To show why it is crucial to use Consistency Groups when a data set spans multiple volumes, consider the following typical sequence of writes for a database update transaction:

1. A write is run to update the database log, which indicates that a database update is about to be performed.
2. A second write is run to perform the actual update to the database.
3. A third write is run to update the database log, which indicates that the database update completed successfully.

The database ensures the correct ordering of these writes by waiting for each step to complete before the next step is started. However, if the database log (updates 1 and 3) and the database (update 2) are on separate volumes, it is possible for the FlashCopy of the database volume to occur before the FlashCopy of the database log. This sequence can result in the target volumes seeing writes 1 and 3 but not 2 because the FlashCopy of the database volume occurred before the write was completed.

In this case, if the database was restarted by using the backup that was made from the FlashCopy target volumes, the database log indicates that the transaction completed successfully. In fact, it did not complete successfully because the FlashCopy of the volume with the database file was started (the bitmap was created) before the write completed to the volume. Therefore, the transaction is lost and the integrity of the database is in question.

To overcome the issue of dependent writes across volumes and to create a consistent image of the client data, a FlashCopy operation must be performed on multiple volumes as an atomic operation. To accomplish this method, the concept of *Consistency Groups* is supported.

A FlashCopy Consistency Group can contain up to 512 FlashCopy mappings. The maximum number of FlashCopy mappings that is supported by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems V8.1.0 is 4096. FlashCopy commands can then be issued to the FlashCopy Consistency Group and, therefore, simultaneously for all of the FlashCopy mappings that are defined in the Consistency Group.

For example, when a FlashCopy **start** command is issued to the Consistency Group, all of the FlashCopy mappings in the Consistency Group are started at the same time. This simultaneous start results in a point-in-time copy that is consistent across all of the FlashCopy mappings that are contained in the Consistency Group.

Consistency Group with Multiple Target FlashCopy

A Consistency Group aggregates FlashCopy mappings, not volumes. Therefore, where a source volume has multiple FlashCopy mappings, they can be in the same or separate Consistency Groups.

If a particular volume is the source volume for multiple FlashCopy mappings, you might want to create separate Consistency Groups to separate each mapping of the same source volume. Regardless of whether the source volume with multiple target volumes is in the same consistency group or in separate consistency groups, the resulting FlashCopy produces multiple identical copies of the source data.

Maximum configurations

Table 10-1 lists the FlashCopy properties and maximum configurations.

Table 10-1 FlashCopy properties and maximum configurations

FlashCopy property	Maximum	Comment
FlashCopy targets per source	256	This maximum is the number of FlashCopy mappings that can exist with the same source volume.
FlashCopy mappings per system	4096	The number of mappings is no longer limited by the number of volumes in the system, so the FlashCopy component limit applies.
FlashCopy Consistency Groups per system	255	This maximum is an arbitrary limit that is policed by the software.
FlashCopy volume capacity per I/O Group	4 pebibytes (PiB)	This maximum is a limit on the quantity of FlashCopy mappings that are using bitmap space from this I/O Group. This maximum configuration uses all 4 gibibytes (GiB) of bitmap space for the I/O Group and allows no Metro or Global Mirror bitmap space. The default is 40 tebibytes (TiB).
FlashCopy mappings per Consistency Group	512	This limit is because of the time that is taken to prepare a Consistency Group with many mappings.

10.3.4 FlashCopy indirection layer

The *FlashCopy indirection layer* governs the I/O to the source and target volumes when a FlashCopy mapping is started, which is done by using a FlashCopy bitmap. The purpose of the FlashCopy indirection layer is to enable the source and target volumes for read and write I/O immediately after the FlashCopy is started.

To show how the FlashCopy indirection layer works, we examine what happens when a FlashCopy mapping is prepared and then started.

When a FlashCopy mapping is prepared and started, the following sequence is applied:

1. Flush the write cache to the source volume or volumes that are part of a Consistency Group.
2. Put cache into write-through mode on the source volumes.
3. Discard cache for the target volumes.
4. Establish a sync point on all of the source volumes in the Consistency Group (which creates the FlashCopy bitmap).
5. Ensure that the indirection layer governs all of the I/O to the source volumes and target volumes.
6. Enable cache on the source volumes and target volumes.

FlashCopy provides the semantics of a point-in-time copy by using the indirection layer, which intercepts I/O that is directed at the source or target volumes. The act of starting a FlashCopy mapping causes this indirection layer to become active in the I/O path, which occurs automatically across all FlashCopy mappings in the Consistency Group.

The indirection layer then determines how each I/O is to be routed, based on the following factors:

- ▶ The volume and the logical block address (LBA) to which the I/O is addressed
- ▶ Its direction (read or write)
- ▶ The state of an internal data structure, the FlashCopy bitmap

The indirection layer allows the I/O to go through to the underlying volume, redirects the I/O from the target volume to the source volume, or queues the I/O while it arranges for data to be copied from the source volume to the target volume. To explain in more detail which action is applied for each I/O, we first look at the FlashCopy bitmap.

10.3.5 Grains and the FlashCopy bitmap

When data is copied between volumes, it is copied in units of address space that are known as *grains*. Grains are units of data that are grouped to optimize the use of the bitmap that tracks changes to the data between the source and target volume. You can use 64 kibibytes (KiB) or 256 KiB grain sizes (256 KiB is the default). The FlashCopy bitmap contains 1 bit for each grain, and is used to show whether the source grain was copied to the target. The 64 KiB grain size uses bitmap space at a rate of four times the default 256 KiB size.

The FlashCopy bitmap dictates read and write behavior for the source and target volumes.

Source reads

Reads are performed from the source volume, which is the same as for non-FlashCopy volumes.

Source writes

Writes to the source cause one of the following actions:

- ▶ If the grain was not copied to the target yet, the grain is copied before the actual write is performed to the source. The bitmap is updated to indicate that this grain is already copied to the target.
- ▶ If the grain was already copied, the write is performed to the source as usual.

Target reads

Reads are performed from the target if the grain was copied. Otherwise, the read is performed from the source and no copy is performed.

Target writes

Writes to the target cause one of the following actions:

- ▶ If the grain was not copied from the source to the target, the grain is copied from the source to the target before the actual write is performed to the source. The bitmap is updated to indicate that this grain is already copied to the target.
- ▶ If the entire grain is being updated on the target, the target is marked as split with the source (if there is no I/O error during the write) and the write goes directly to the target.
- ▶ If the grain in question was already copied from the source to the target, the write goes directly to the target.

The FlashCopy indirection layer algorithm

Imagine the FlashCopy indirection layer as the I/O traffic director when a FlashCopy mapping is active. The I/O is intercepted and handled according to whether it is directed at the source volume or at the target volume, depending on the nature of the I/O (read or write) and the state of the grain (whether it was copied).

Figure 10-6 on page 464 shows how the background copy runs while I/Os are handled according to the indirection layer algorithm.

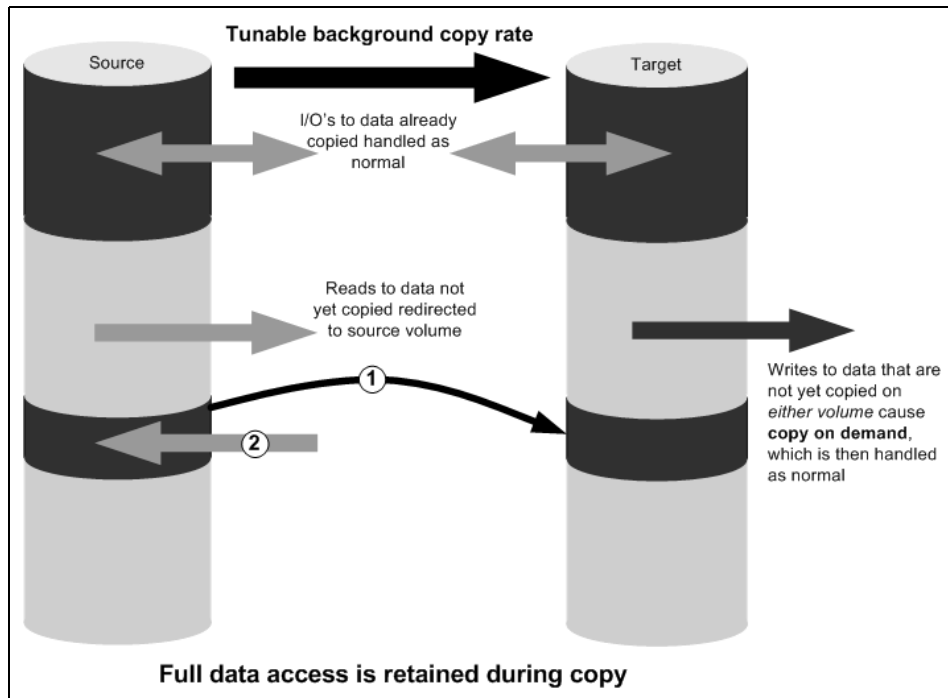


Figure 10-6 I/O processing with FlashCopy

10.3.6 Interaction and dependency between multiple target FlashCopy mappings

Figure 10-7 shows a set of four FlashCopy mappings that share a common source. The FlashCopy mappings target volumes Target 0, Target 1, Target 2, and Target 3.

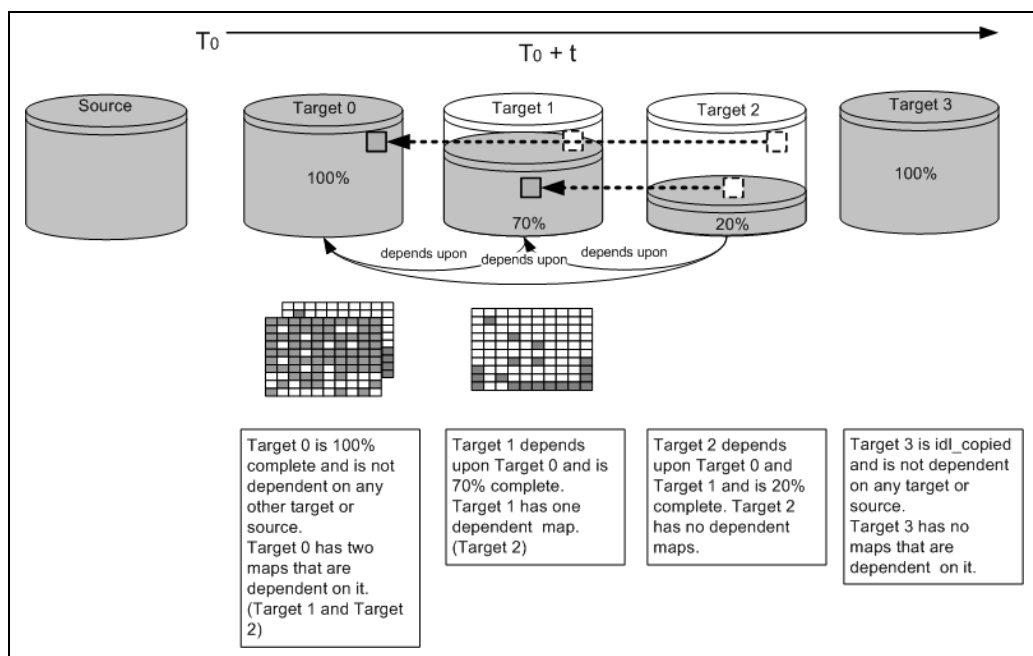


Figure 10-7 Interactions among multiple target FlashCopy mappings

In Figure 10-7 on page 464, Target 0 is not dependent on a source because it completed copying. Target 0 has two dependent mappings (Target 1 and Target 2).

Target 1 depends on Target 0. It remains dependent until all of Target 1 is copied. Target 2 depends on it because Target 2 is 20% copy complete. After all of Target 1 is copied, it can then move to the `idle_copied` state.

Target 2 is dependent upon Target 0 and Target 1 and remains dependent until all of Target 2 is copied. No target depends on Target 2; therefore, when all of the data is copied to Target 2, it can move to the `idle_copied` state.

Target 3 completed copying, so it is not dependent on any other maps.

Target writes with Multiple Target FlashCopy

A write to an intermediate or the newest target volume must consider the state of the grain within its own mapping, and the state of the grain of the next oldest mapping.

If the grain of the next oldest mapping is not yet copied, it must be copied before the write can proceed, to preserve the contents of the next oldest mapping. The data that is written to the next oldest mapping comes from a target or source.

If the grain in the target that is being written is not yet copied, the grain is copied from the oldest copied grain in the mappings that are newer than the target, or from the source if none is copied. After this copy is done, the write can be applied to the target.

Target reads with Multiple Target FlashCopy

If the grain being read is copied from the source to the target, the read returns data from the target that is being read. If the grain is not yet copied, each of the newer mappings is examined in turn, and the read is performed from the first copy that is found. If none is found, the read is performed from the source.

Stopping the copy process

When a **stop** command is issued to a mapping that contains a target that has dependent mappings, the mapping enters the stopping state and begins copying all grains that are uniquely held on the target volume of the mapping that is being stopped to the next oldest mapping that is in the Copying state. The mapping remains in the stopping state until all grains are copied, and then enters the stopped state.

Note: The stopping copy process can be ongoing for several mappings that share the source at the same time. At the completion of this process, the mapping automatically makes an asynchronous state transition to the stopped state, or the `idle_copied` state if the mapping was in the copying state with `progress = 100%`.

For example, if the mapping that is associated with Target 0 was issued a **stopfcmap** or **stopfcconsistgrp** command, Target 0 enters the stopping state while a process copies the data of Target 0 to Target 1. After all of the data is copied, Target 0 enters the stopped state, and Target 1 is no longer dependent upon Target 0; however, Target 1 remains dependent on Target 2.

10.3.7 Summary of the FlashCopy indirection layer algorithm

Table 10-2 on page 466 summarizes the indirection layer algorithm.

Table 10-2 Summary table of the FlashCopy indirection layer algorithm

Accessed volume	Was the grain copied?	Host I/O operation	
		Read	Write
Source	No	Read from the source volume.	Copy grain to most recently started target for this source, then write to the source.
	Yes	Read from the source volume.	Write to the source volume.
Target	No	If any newer targets exist for this source in which this grain was copied, read from the oldest of these targets. Otherwise, read from the source.	Hold the write. Check the dependency target volumes to see whether the grain was copied. If the grain is not copied to the next oldest target for this source, copy the grain to the next oldest target. Then, write to the target.
	Yes	Read from the target volume.	Write to the target volume.

10.3.8 Interaction with the cache

Starting with V7.3, the entire cache subsystem was redesigned and changed. Cache has been divided into upper and lower cache. Upper cache serves mostly as write cache and hides the write latency from the hosts and application. Lower cache is a read/write cache and optimizes I/O to and from disks. Figure 10-8 shows the new controller firmware cache architecture.

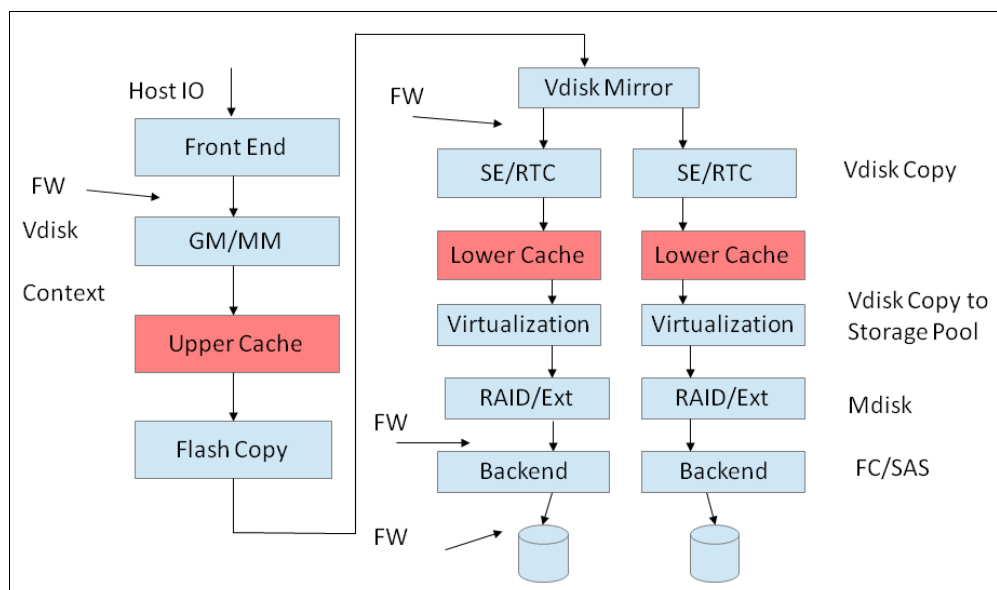


Figure 10-8 New cache architecture

This copy-on-write process introduces significant latency into write operations. To isolate the active application from this additional latency, the FlashCopy indirection layer is placed logically between upper and lower cache. Therefore, the additional latency that is introduced

by the copy-on-write process is encountered only by the internal cache operations and not by the application.

Figure 10-9 shows the logical placement of the FlashCopy indirection layer.

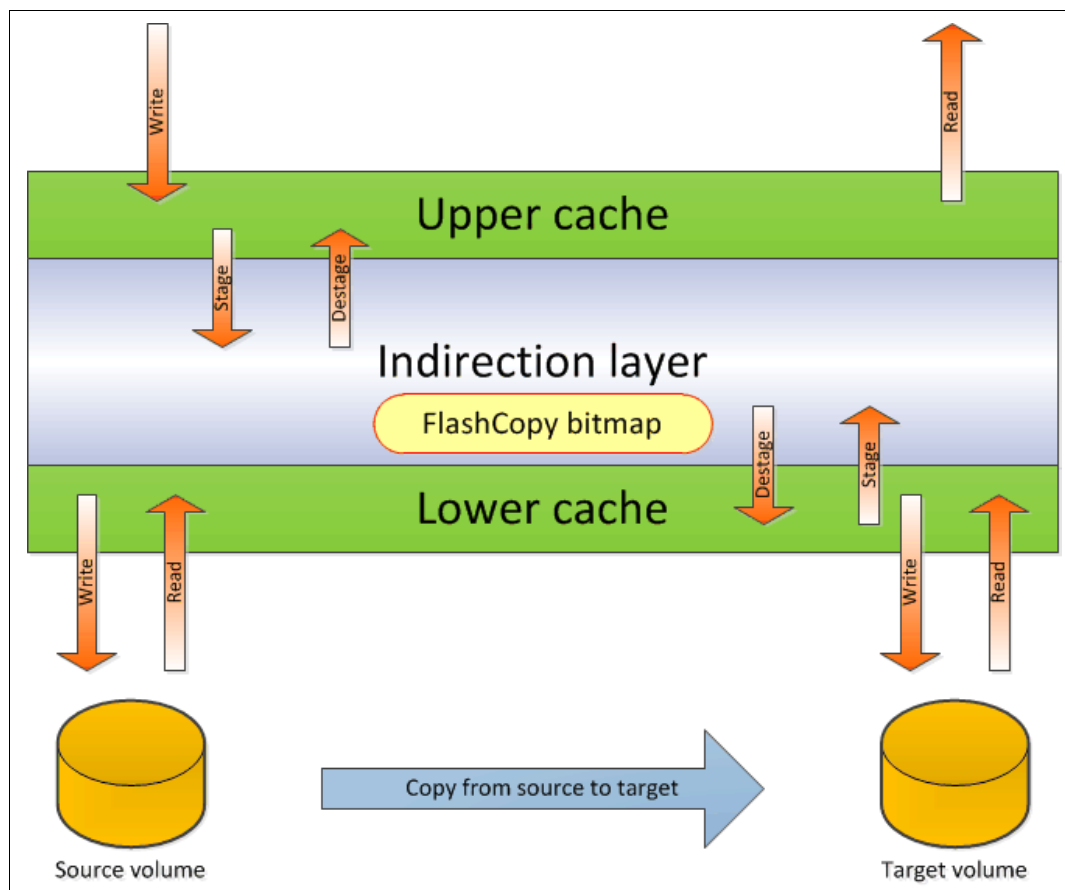


Figure 10-9 Logical placement of the FlashCopy indirection layer

Introduction of the two-level cache provides additional performance improvements to the FlashCopy mechanism. Because now the FlashCopy layer is above lower cache in the Lenovo storage V-series software stack, it can benefit from read prefetching and coalescing writes to backend storage. Also, preparing FlashCopy is much faster, because upper cache write data does not have to go directly to backend storage but to lower cache layer.

Additionally, in the multitarget FlashCopy the target volumes of the same image share cache data. This design is opposite to previous code versions, where each volume had its own copy of cached data.

10.3.9 FlashCopy and image mode volumes

FlashCopy can be used with image mode volumes. Because the source and target volumes must be the same size, you must create a target volume with the same size as the image mode volume when you are creating a FlashCopy mapping. To accomplish this task, use the `svcinfo lsvdisk -bytes volumeName` command. The size in bytes is then used to create the volume that is used in the FlashCopy mapping.

This method provides an exact number of bytes because image mode volumes might not line up one-to-one on other measurement unit boundaries. Example 10-1 on page 468 lists the

size of the test_image_vol_1 volume. The test_image_vol_copy_1 volume is then created, which specifies the same size.

Example 10-1 Listing the size of a volume in bytes and creating a volume of equal size

```
IBM_V5000_Gen2:superuser>lsvdisk -bytes test_image_vol_1
```

```
id 12
name test_image_vol_1
IO_group_id 0
IO_group_name io_grp0
status online
mdisk_grp_id 3
mdisk_grp_name temp_migration_pool
capacity 21474836480
type image
formatted no
formatting no
mdisk_id 5
mdisk_name mdisk3
FC_id
FC_name
RC_id
RC_name
vdisk_UID 600507680283818B3000000000000000E
throttling 0
preferred_node_id 2
fast_write_state empty
cache readwrite
udid
fc_map_count 0
sync_rate 50
copy_count 1
se_copy_count 0
filesystem
mirror_write_priority latency
RC_change no
compressed_copy_count 0
access_IO_group_count 1
last_access_time
parent_mdisk_grp_id 3
parent_mdisk_grp_name temp_migration_pool
owner_type none
owner_id
owner_name
encrypt no
volume_id 12
volume_name test_image_vol_1
function

copy_id 0
status online
sync yes
auto_delete no
primary yes
mdisk_grp_id 3
mdisk_grp_name temp_migration_pool
```

```

type image
mdisk_id 5
mdisk_name mdisk3
fast_write_state empty
used_capacity 21474836480
real_capacity 21474836480
free_capacity 0
overallocation 100
autoexpand
warning
grainsize
se_copy no
easy_tier on
easy_tier_status measured
tier ssd
tier_capacity 0
tier enterprise
tier_capacity 21474836480
tier nearline
tier_capacity 0
compressed_copy no
uncompressed_used_capacity 21474836480
parent_mdisk_grp_id 3
parent_mdisk_grp_name temp_migration_pool
encrypt no

```

```

IBM_V5000_Gen2:superuser>mkvdisk -mdiskgrp test_pool_1 -iogrp 0 -size 21474836480
-unit b -name test_image_vol_copy_1
Virtual Disk, id [13], successfully created

```

```

IBM_V5000_Gen2:superuser>lsvdisk -delim " "
12 test_image_vol_1 0 io_grp0 online 3 temp_migration_pool 20.00GB image
600507680283818B3000000000000000E 0 1 empty 0 no 0 3 temp_migration_pool no no 12
test_image_vol_1
13 test_image_vol_copy_1 0 io_grp0 online 0 test_pool_1 20.00GB striped
600507680283818B3000000000000000F 0 1 not_empty 0 no 0 0 test_pool_1 yes no 13
test_image_vol_copy_1

```

Tip: Alternatively, you can use the **expandvolumesize** and **shrinkvolumesize** volume commands to modify the size of the volume.

These actions must be performed before a mapping is created.

You can use an image mode volume as a FlashCopy source volume or target volume.

10.3.10 FlashCopy mapping events

In this section, we describe the events that modify the states of a FlashCopy. We also describe the mapping events that are listed in Table 10-3.

Overview of a FlashCopy sequence of events: The following tasks show the FlashCopy sequence:

1. Associate the source data set with a target location (one or more source and target volumes).
2. Create a FlashCopy mapping for each source volume to the corresponding target volume. The target volume must be equal in size to the source volume.
3. Discontinue access to the target (application dependent).
4. Prepare (pre-trigger) the FlashCopy:
 - a. Flush the cache for the source.
 - b. Discard the cache for the target.
5. Start (trigger) the FlashCopy:
 - a. Pause I/O (briefly) on the source.
 - b. Resume I/O on the source.
 - c. Start I/O on the target.

Table 10-3 Mapping events

Mapping event	Description
Create	A FlashCopy mapping is created between the specified source volume and the specified target volume. The operation fails if any one of the following conditions is true: <ul style="list-style-type: none">▶ The source volume is a member of 256 FlashCopy mappings.▶ The node has insufficient bitmap memory.▶ The source and target volumes are different sizes.
Prepare	The prestartfcmap or prestartfcconsistgrp command is directed to a Consistency Group for FlashCopy mappings that are members of a normal Consistency Group or to the mapping name for FlashCopy mappings that are stand-alone mappings. The prestartfcmap or prestartfcconsistgrp command places the FlashCopy mapping into the Preparing state. The prestartfcmap or prestartfcconsistgrp command can corrupt any data that was on the target volume because cached writes are discarded. Even if the FlashCopy mapping is never started, the data from the target might be changed logically during the act of preparing to start the FlashCopy mapping.
Flush done	The FlashCopy mapping automatically moves from the preparing state to the prepared state after all cached data for the source is flushed and all cached data for the target is no longer valid.

Mapping event	Description
Start	<p>When all of the FlashCopy mappings in a Consistency Group are in the prepared state, the FlashCopy mappings can be started. To preserve the cross-volume Consistency Group, the start of all of the FlashCopy mappings in the Consistency Group must be synchronized correctly concerning I/Os that are directed at the volumes by using the startfcmap or startfcconsistgrp command.</p> <p>The following actions occur during the running of the startfcmap command or the startfcconsistgrp command:</p> <ul style="list-style-type: none"> ▶ New reads and writes to all source volumes in the Consistency Group are paused in the cache layer until all ongoing reads and writes beneath the cache layer are completed. ▶ After all FlashCopy mappings in the Consistency Group are paused, the internal cluster state is set to enable FlashCopy operations. ▶ After the cluster state is set for all FlashCopy mappings in the Consistency Group, read and write operations continue on the source volumes. ▶ The target volumes are brought online. <p>As part of the startfcmap or startfcconsistgrp command, read and write caching is enabled for the source and target volumes.</p>
Modify	<p>The following FlashCopy mapping properties can be modified:</p> <ul style="list-style-type: none"> ▶ FlashCopy mapping name ▶ Clean rate ▶ Consistency group ▶ Copy rate (for background copy or stopping copy priority) ▶ Automatic deletion of the mapping when the background copy is complete
Stop	<p>The following separate mechanisms can be used to stop a FlashCopy mapping:</p> <ul style="list-style-type: none"> ▶ Issue a command ▶ An I/O error occurred
Delete	<p>This command requests that the specified FlashCopy mapping is deleted. If the FlashCopy mapping is in the copying state, the force flag must be used.</p>
Flush failed	<p>If the flush of data from the cache cannot be completed, the FlashCopy mapping enters the stopped state.</p>
Copy complete	<p>After all of the source data is copied to the target and there are no dependent mappings, the state is set to copied. If the option to automatically delete the mapping after the background copy completes is specified, the FlashCopy mapping is deleted automatically. If this option is not specified, the FlashCopy mapping is not deleted automatically and can be reactivated by preparing and starting again.</p>
Bitmap online/offline	<p>The node failed.</p>

10.3.11 FlashCopy mapping states

This section describes the states of a FlashCopy mapping.

Idle_or_copied

The source and target volumes act as independent volumes even if a mapping exists between the two. Read and write caching is enabled for the source and the target volumes.

If the mapping is incremental and the background copy is complete, the mapping records the differences between the source and target volumes only. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target volumes are offline.

Copying

The copy is in progress. Read and write caching is enabled on the source and the target volumes.

Prepared

The mapping is ready to start. The target volume is online, but is not accessible. The target volume cannot perform read or write caching. Read and write caching is failed by the Small Computer System Interface (SCSI) front end as a hardware error. If the mapping is incremental and a previous mapping is completed, the mapping records the differences between the source and target volumes only. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target volumes go offline.

Preparing

The target volume is online, but not accessible. The target volume cannot perform read or write caching. Read and write caching is failed by the SCSI front end as a hardware error. Any changed write data for the source volume is flushed from the cache. Any read or write data for the target volume is discarded from the cache.

If the mapping is incremental and a previous mapping is completed, the mapping records the differences between the source and target volumes only. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target volumes go offline.

Performing the cache flush that is required as part of the **startfcmap** or **startfcconsistgrp** command causes I/Os to be delayed while they are waiting for the cache flush to complete. To overcome this problem, FlashCopy supports the **prestartfcmap** or **prestartfcconsistgrp** commands, which prepare for a FlashCopy start while still allowing I/Os to continue to the source volume.

In the Preparing state, the FlashCopy mapping is prepared by completing the following steps:

1. Flushing any modified write data that is associated with the source volume from the cache. Read data for the source is left in the cache.
2. Placing the cache for the source volume into write-through mode so that subsequent writes wait until data is written to disk before the **write** command that is received from the host is complete.
3. Discarding any read or write data that is associated with the target volume from the cache.

Stopped

The mapping is stopped because you issued a **stop** command or an I/O error occurred. The target volume is offline and its data is lost. To access the target volume, you must restart or delete the mapping. The source volume is accessible and the read and write cache is enabled. If the mapping is incremental, the mapping is recording write operations to the source volume. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target volumes go offline.

Stopping

The mapping is copying data to another mapping.

If the background copy process is complete, the target volume is online while the stopping copy process completes.

If the background copy process is not complete, data is discarded from the target volume cache. The target volume is offline while the stopping copy process runs.

The source volume is accessible for I/O operations.

Suspended

The mapping started, but it did not complete. Access to the metadata is lost, which causes the source and target volume to go offline. When access to the metadata is restored, the mapping returns to the copying or stopping state and the source and target volumes return online. The background copy process resumes. Any data that was not flushed and was written to the source or target volume before the suspension is in cache until the mapping leaves the suspended state.

Summary of FlashCopy mapping states

Table 10-4 lists the various FlashCopy mapping states, and the corresponding states of the source and target volumes.

Table 10-4 *FlashCopy mapping state summary*

State	Source		Target	
	Online/Offline	Cache state	Online/Offline	Cache state
Idling/Copied	Online	Write-back	Online	Write-back
Copying	Online	Write-back	Online	Write-back
Stopped	Online	Write-back	Offline	N/A
Stopping	Online	Write-back	► Online if copy complete ► Offline if copy incomplete	N/A
Suspended	Offline	Write-back	Offline	N/A
Preparing	Online	Write-through	Online but not accessible	N/A
Prepared	Online	Write-through	Online but not accessible	N/A

10.3.12 Thin provisioned FlashCopy

FlashCopy source and target volumes can be thin-provisioned.

Source or target thin-provisioned

The most common configuration is a fully allocated source and a thin-provisioned target. By using this configuration, the target uses a smaller amount of real storage than the source. With this configuration, use the NOCOPY (background copy rate = 0%) option only. Although the COPY option is supported, this option creates a fully allocated target, which defeats the purpose of thin provisioning.

Source and target thin-provisioned

When the source and target volumes are thin-provisioned, only the data that is allocated to the source is copied to the target. In this configuration, the background copy option has no effect.

Performance: The best performance is obtained when the grain size of the thin-provisioned volume is the same as the grain size of the FlashCopy mapping.

Thin-provisioned incremental FlashCopy

The implementation of thin-provisioned volumes does not preclude the use of incremental FlashCopy on the same volumes. It does not make sense to have a fully allocated source volume and then use incremental FlashCopy (which is always a full copy the first time) to copy this fully allocated source volume to a thin-provisioned target volume. However, this action is not prohibited.

Consider the following optional configurations:

- ▶ A thin-provisioned source volume can be copied incrementally by using FlashCopy to a thin-provisioned target volume. Whenever the FlashCopy is performed, only data that was modified is recopied to the target. If space is allocated on the target because of I/O to the target volume, this space is not reclaimed with subsequent FlashCopy operations.
- ▶ A fully allocated source volume can be copied incrementally by using FlashCopy to another fully allocated volume at the same time as it is being copied to multiple thin-provisioned targets (taken at separate points in time). By using this combination, a single full backup can be kept for recovery purposes, and the backup workload is separated from the production workload. At the same time, older thin-provisioned backups can be retained.

10.3.13 Background copy

With FlashCopy background copy enabled, the source volume data is copied to the corresponding target volume. With the FlashCopy background copy disabled, only data that changed on the source volume is copied to the target volume.

The benefit of using a FlashCopy mapping with background copy enabled is that the target volume becomes a real clone (independent from the source volume) of the FlashCopy mapping source volume after the copy is complete. When the background copy function is not performed, the target volume remains a valid copy of the source data only while the FlashCopy mapping remains in place.

The *background copy rate* is a property of a FlashCopy mapping that is defined as a value 0 - 100. The background copy rate can be defined and changed dynamically for individual FlashCopy mappings. A value of 0 disables the background copy.

Table 10-5 shows the relationship of the background copy rate value to the attempted number of grains to be copied per second.

Table 10-5 Background copy rate

Value	Data copied per second	Grains per second (256 KB grain)	Grains per second (64 KB grain)
01 - 10	128 KiB	0.5	2
11 - 20	256 KiB	1	4
21 - 30	512 KiB	2	8
31 - 40	1 mebibyte (MiB)	4	16
41 - 50	2 MiB	8	32
51 - 60	4 MiB	16	64
61 - 70	8 MiB	32	128
71 - 80	16 MiB	64	256
81 - 90	32 MiB	128	512
91 - 100	64 MiB	256	1024
101-110	128 MiB	512	2048
111-120	256 MiB	1024	4096
121 - 130	512 MiB	2048	8192
131 - 140	1 GiB	4096	16384
141 - 150	2 GiB	8192	32768

The *grains per second* numbers represent the maximum number of grains that the IBM Storwize V5000 for Lenovo copies per second, assuming that the bandwidth to the managed disks (MDisks) can accommodate this rate.

If the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems cannot achieve these copy rates because of insufficient width from the nodes to the MDisk, the background copy I/O contends for resources on an equal basis with the I/O that is arriving from the hosts. Background copy I/O and I/O that is arriving from the hosts tend to see an increase in latency and a consequential reduction in throughput.

Background copy and foreground I/O continue to make progress, and do not stop, hang, or cause the node to fail. The background copy is performed by both nodes of the I/O Group in which the source volume is found.

10.3.14 Serialization of I/O by FlashCopy

In general, the FlashCopy function in the introduces no explicit serialization into the I/O path. Therefore, many concurrent I/Os are allowed to the source and target volumes.

However, there is a lock for each grain. The lock can be in shared or exclusive mode. For multiple targets, a common lock is shared, and the mappings are derived from a particular source volume. The lock is used in the following modes under the following conditions:

- ▶ The lock is held in shared mode during a read from the target volume, which touches a grain that was not copied from the source.
- ▶ The lock is held in exclusive mode while a grain is being copied from the source to the target.

If the lock is held in shared mode and another process wants to use the lock in shared mode, this request is granted unless a process is already waiting to use the lock in exclusive mode.

If the lock is held in shared mode and it is requested to be exclusive, the requesting process must wait until all holders of the shared lock free it.

Similarly, if the lock is held in exclusive mode, a process that is wanting to use the lock in shared or exclusive mode must wait for it to be freed.

10.3.15 Event handling

When a FlashCopy mapping is not copying or stopping, the FlashCopy function does not affect the handling or reporting of events for error conditions that are encountered in the I/O path. Event handling and reporting are affected only by FlashCopy when a FlashCopy mapping is copying or stopping; that is, actively moving data.

We describe these scenarios next.

Node failure

Normally, two copies of the FlashCopy bitmap are maintained. One copy of the FlashCopy bitmap is on each of the two nodes that make up the I/O Group of the source volume. When a node fails, one copy of the bitmap for all FlashCopy mappings whose source volume is a member of the failing node's I/O Group becomes inaccessible.

FlashCopy continues with a single copy of the FlashCopy bitmap that is stored as non-volatile in the remaining node in the source I/O Group. The system metadata is updated to indicate that the missing node no longer holds a current bitmap. When the failing node recovers or a replacement node is added to the I/O Group, the bitmap redundancy is restored.

Path failure (Path Offline state)

In a fully functioning system, all of the nodes have a software representation of every volume in the system within their application hierarchy.

Because the storage area network (SAN) that links Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems node canisters to each other and to the MDisks is made up of many independent links, it is possible for a subset of the nodes to be temporarily isolated from several of the MDisks. When this situation happens, the managed disks are said to be *path offline* on certain nodes.

Other nodes: Other nodes might see the managed disks as Online because their connection to the managed disks is still functioning.

Path Offline for the source volume

If a FlashCopy mapping is in the copying state and the source volume goes path offline, this path offline state is propagated to all target volumes up to, but not including, the target volume

for the newest mapping that is 100% copied but remains in the copying state. If no mappings are 100% copied, all of the target volumes are taken offline. Path offline is a state that exists on a per-node basis. Other nodes might not be affected. If the source volume comes online, the target and source volumes are brought back online.

Path Offline for the target volume

If a target volume goes path offline but the source volume is still online, and if there are any dependent mappings, those target volumes also go path offline. The source volume remains online.

10.3.16 Asynchronous notifications

FlashCopy raises informational event log entries for certain mapping and Consistency Group state transitions. These state transitions occur as a result of configuration events that complete asynchronously. The informational events can be used to generate Simple Network Management Protocol (SNMP) traps to notify the user.

Other configuration events complete synchronously, and no informational events are logged as a result of the following events:

- ▶ **PREPARE_COMPLETED**

This state transition is logged when the FlashCopy mapping or Consistency Group enters the prepared state as a result of a user request to prepare. The user can now start (or stop) the mapping or Consistency Group.

- ▶ **COPY_COMPLETED**

This state transition is logged when the FlashCopy mapping or Consistency Group enters the idle_or_copied state when it was in the copying or stopping state. This state transition indicates that the target disk now contains a complete copy and no longer depends on the source.

- ▶ **STOP_COMPLETED**

This state transition is logged when the FlashCopy mapping or Consistency Group enters the stopped state as a result of a user request to stop. It is logged after the automatic copy process completes. This state transition includes mappings where no copying needed to be performed. This state transition differs from the event that is logged when a mapping or group enters the stopped state as a result of an I/O error.

10.3.17 Interoperation with Metro Mirror and Global Mirror

A volume can be part of any copy relationship (FlashCopy, Metro Mirror, or Remote Mirror). Therefore, FlashCopy can work with Metro Mirror and Global Mirror to provide better protection of the data.

For example, we can perform a Metro Mirror copy to duplicate data from Site_A to Site_B and then perform a daily FlashCopy to back up the data to another location.

Note: If a volume is set to Transparent Cloud Tiering function, it cannot be part of FlashCopy, Metro Mirror, or Remote Mirror.

Table 10-6 on page 478 lists the supported combinations of FlashCopy and remote copy. In the table, *remote copy* refers to Metro Mirror and Global Mirror.

Table 10-6 FlashCopy and remote copy interaction

Component	Remote copy primary site	Remote copy secondary site
FlashCopy Source	Supported	Supported latency: When the FlashCopy relationship is in the preparing and prepared states, the cache at the remote copy secondary site operates in write-through mode. This process adds latency to the latent remote copy relationship.
FlashCopy Target	This is a supported combination and has the following restrictions: <ul style="list-style-type: none"> ▶ Issuing a stop -force might cause the remote copy relationship to be fully resynchronized. ▶ Code level must be 6.2.x or later. ▶ I/O Group must be the same. 	This is a supported combination with the major restriction that the FlashCopy mapping cannot be copying, stopping, or suspended. Otherwise, the restrictions are the same as at the remote copy primary site.

10.3.18 FlashCopy presets

The controller firmware GUI interface provides three FlashCopy presets (Snapshot, Clone, and Backup) to simplify the more common FlashCopy operations.

Although these presets meet most FlashCopy requirements, they do not support all possible FlashCopy options. If more specialized options are required that are not supported by the presets, the options must be performed by using CLI commands.

This section describes the preset options and their use cases.

Snapshot

This preset creates a copy-on-write point-in-time copy. The snapshot is not intended to be an independent copy. Instead, the copy is used to maintain a view of the production data at the time that the snapshot is created. Therefore, the snapshot holds only the data from regions of the production volume that changed since the snapshot was created. Because the snapshot preset uses thin provisioning, only the capacity that is required for the changes is used.

Snapshot uses the following preset parameters:

- ▶ Background copy: None
- ▶ Incremental: No
- ▶ Delete after completion: No
- ▶ Cleaning rate: No
- ▶ Primary copy source pool: Target pool

Use case

The user wants to produce a copy of a volume without affecting the availability of the volume. The user does not anticipate many changes to be made to the source or target volume; a significant proportion of the volumes remains unchanged.

By ensuring that only changes require a copy of data to be made, the total amount of disk space that is required for the copy is reduced. Therefore, many Snapshot copies can be used in the environment.

Snapshots are useful for providing protection against corruption or similar issues with the validity of the data, but they do not provide protection from physical controller failures. Snapshots can also provide a vehicle for performing repeatable testing (including “what-if” modeling that is based on production data) without requiring a full copy of the data to be provisioned.

Clone

The clone preset creates a replica of the volume, which can be changed without affecting the original volume. After the copy completes, the mapping that was created by the preset is automatically deleted.

Clone uses the following preset parameters:

- ▶ Background copy rate: 50
- ▶ Incremental: No
- ▶ Delete after completion: Yes
- ▶ Cleaning rate: 50
- ▶ Primary copy source pool: Target pool

Use case

Users want a copy of the volume that they can modify without affecting the original volume. After the clone is established, there is no expectation that it is refreshed or that there is any further need to reference the original production data again. If the source is thin-provisioned, the target is thin-provisioned for the auto-create target.

Backup

The backup preset creates a point-in-time replica of the production data. After the copy completes, the backup view can be refreshed from the production data, with minimal copying of data from the production volume to the backup volume.

Backup uses the following preset parameters:

- ▶ Background Copy rate: 50
- ▶ Incremental: Yes
- ▶ Delete after completion: No
- ▶ Cleaning rate: 50
- ▶ Primary copy source pool: Target pool

Use case

The user wants to create a copy of the volume that can be used as a backup if the source becomes unavailable, as in the case of loss of the underlying physical controller. The user plans to periodically update the secondary copy, and does not want to suffer from the resource demands of creating a new copy each time (and incremental FlashCopy times are faster than full copy, which helps to reduce the window where the new backup is not yet fully effective). If the source is thin-provisioned, the target is also thin-provisioned in this option for the auto-create target.

Another use case, which is not supported by the name, is to create and maintain (periodically refresh) an independent image that can be subjected to intensive I/O (for example, data mining) without affecting the source volume's performance.

10.4 Managing FlashCopy by using the GUI

It is often easier to work with the FlashCopy function from the GUI if you have a reasonable number of host mappings. However, in enterprise data centers with many host mappings, we suggest that you use the CLI to run your FlashCopy commands.

This section describes the tasks that you can perform at a FlashCopy level using the GUI.

The following methods can be used to visualize and manage your FlashCopy:

- Use the main pane. Move the mouse pointer over Copy Services in the dynamic menu and click **FlashCopy**, as shown in Figure 10-10.

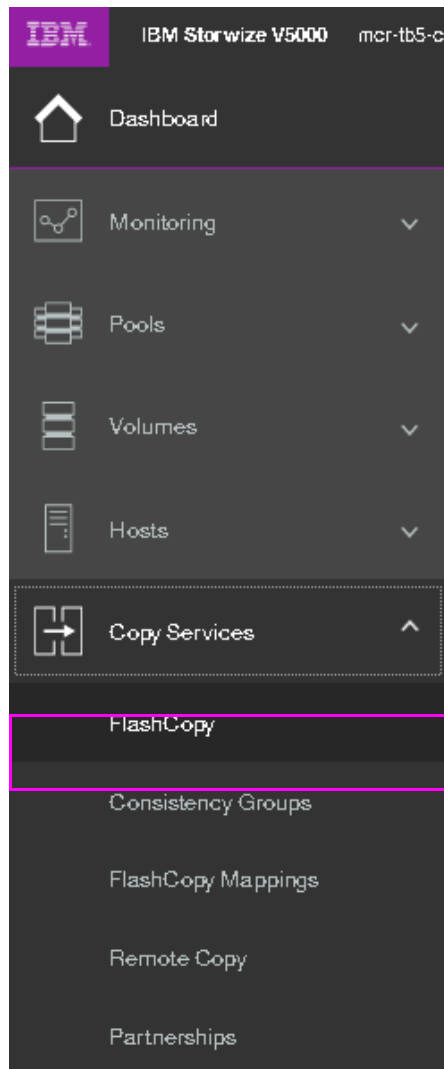


Figure 10-10 FlashCopy pane

In its basic mode, the FlashCopy function copies the contents of a source volume to a target volume. Any data that existed on the target volume is lost, and that data is replaced by the copied data.

- From the Copy Services option on the main panel, use the Consistency Groups option, as shown in Figure 10-11. A *Consistency Group* is a container for mappings. You can add many mappings to a Consistency Group.



Figure 10-11 Consistency Groups pane

- From the Copy Services option on the main panel, use the FlashCopy Mappings pane, as shown in Figure 10-12. A *FlashCopy mapping* defines the relationship between a source volume and a target volume.



Figure 10-12 FlashCopy Mappings pane

10.4.1 Creating a FlashCopy mapping

In this section, we create FlashCopy mappings for volumes and their targets.

Complete the following steps:

1. From the main pane, move the mouse pointer over Copy Services click **FlashCopy**. The FlashCopy pane opens, as shown in Figure 10-13.



Figure 10-13 FlashCopy pane

2. Select the volume for which you want to create the FlashCopy relationship, as shown in Figure 10-14 on page 482.

Multiple FlashCopy mappings: To create multiple FlashCopy mappings at one time, select multiple volumes by holding down Ctrl and clicking the entries that you want.

Volume Name	Status	Progress	Capacity	Group	Flash Time
ITSO_BASIC_VOL_FC	Idle		100 GiB		
ITSO_FC_SRC_VOL	Idle		100 GiB		
ITSO_FC_SRC_VOL_2	Idle		100 GiB		
ITSO_FC_SRC_VOL_2_3	Idle		100 GiB		
ITSO_FC_SRC_VOL_3	Idle		100 GiB		
ITSO_FC_TGT_VOL	Idle		100 GiB		
ITSO_FC_TGT_VOL_2	Idle		100 GiB		
VOL_1	Idle		100 GiB		
VOL_2	Idle		100 GiB		
VOL_3	Idle		100 GiB		

Figure 10-14 FlashCopy mapping: Select the volume (or volumes)

Depending on whether you created the target volumes for your FlashCopy mappings or you want the system to create the target volumes for you, the following options are available:

- ▶ If you created the target volumes, see “Using existing target volumes” on page 482.
- ▶ If you want the system to create the target volumes for you, see “Creating target volumes” on page 487.

Using existing target volumes

Complete the following steps to use existing target volumes for the FlashCopy mappings:

1. Select the source volume that you want to use. Then, click **Actions** → **Advanced FlashCopy** → **Use Existing Target Volumes**, as shown in Figure 10-15.

Volume Name	Status	Progress	Capacity	Group	Flash Time
Encrypted_Copies			16.00 GiB		
ITSO_BASIC_VOL			1.00 GiB		
ITSO_COMPR_VOL			1.00 GiB		
ITSO_CUST_MIR_VOL			2.00 GiB		
ITSO_CUST_TP_MIR_VOL			10.00 GiB		
ITSO_FC_1_COPY_VOL			5.00 GiB		
ITSO_FC_1_VOL			5.00 GiB		
ITSO_FC_2_COPY_VOL			3.00 GiB		
ITSO_FC_2_VOL			3.00 GiB		
ITSO_FC_2_COPY_VOL	Idle	0%		ITSO_FC_CS_GRP	
ITSO_FC_SRC_VOL_1			4.00 GiB		
ITSO_FC_SRC_VOL_2			5.00 GiB		
ITSO_FC_SRC_VOL_3			6.00 GiB		
ITSO_FC_SRC_VOL_4			2.00 GiB		
ITSO_FC_TGT_VOL_1			4.00 GiB		
ITSO_FC_TGT_VOL_3			6.00 GiB		
ITSO_MIG_BY_VOL_COPY_VOL			3.00 GiB		
ITSO_THROTTLE_VOL			5.00 GiB		
ITSO_TP_VOL			1.00 GiB		
basic_1_pool_1			1.00 GiB		
test_vol_1			1.00 GiB		

Figure 10-15 Using existing target volumes

2. The Create FlashCopy Mapping window opens (Figure 10-16). In this window, you must create the relationship between the source volume (the disk that is copied) and the target volume (the disk that receives the copy). A mapping can be created between any two volumes that are managed by the same clustered system. Select a source volume and a target volume for your FlashCopy mapping, and then click **Add**. If you must create other copies, repeat this step.

Important: The source volume and the target volume must be of equal size. Therefore, only targets of the same size are shown in the list for a source volume.

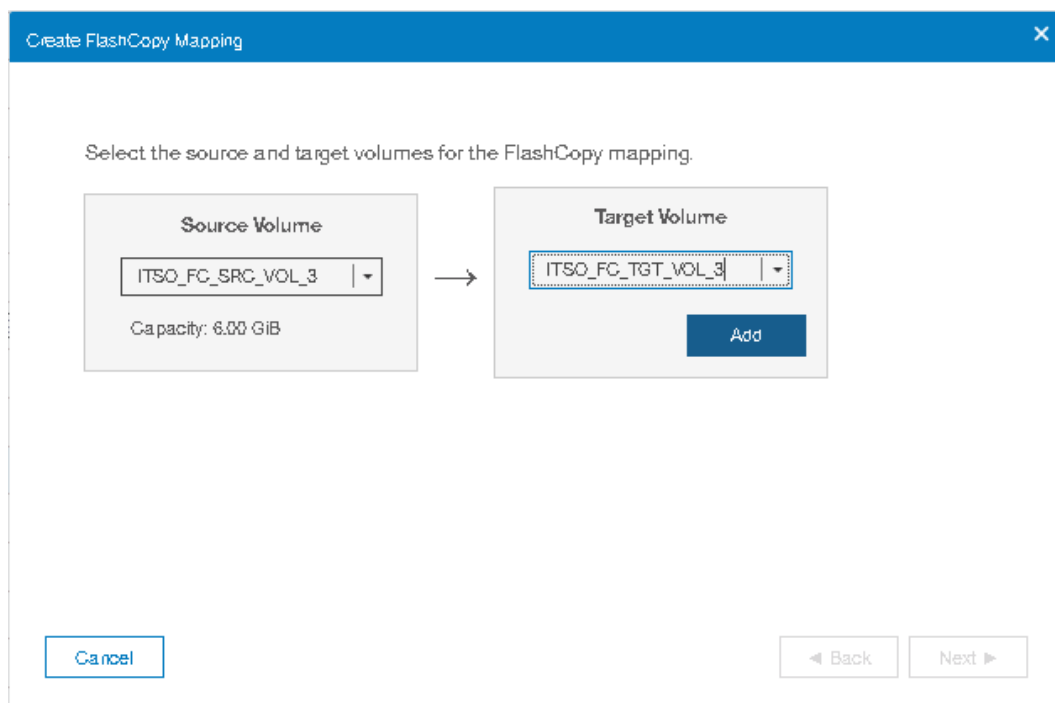


Figure 10-16 Create a FlashCopy Mapping by using an existing target volume

To remove a relationship that was created, click **X**, as shown in Figure 10-17 on page 484.

Volumes: The volumes do not have to be in the same I/O Group or storage pool.

3. Click **Next** after you create all of the relationships that you need, as shown in Figure 10-17 on page 484.

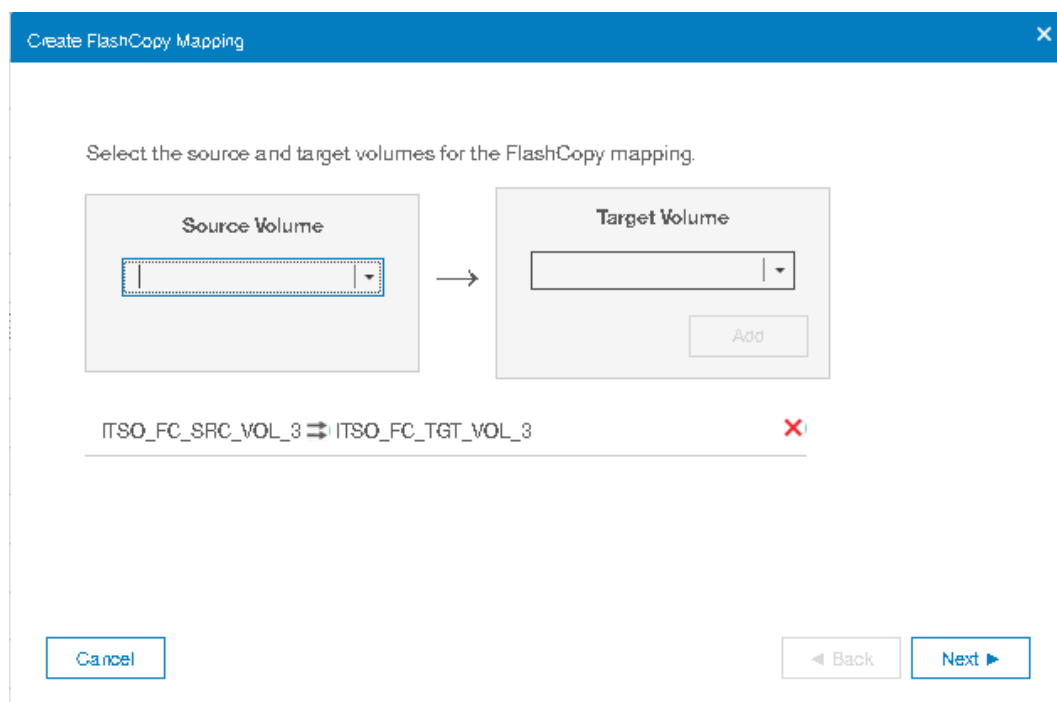


Figure 10-17 Create FlashCopy Mapping window

4. In the next window, select one FlashCopy preset. The GUI provides the following presets to simplify common FlashCopy operations.
 - Snapshot: Creates a copy-on-write point-in-time copy.
 - Clone: Creates a replica of the source volume on a target volume. The copy can be changed without affecting the original volume.
 - Backup: Creates a FlashCopy mapping that can be used to recover data or objects if the system experiences data loss. These backups can be copied multiple times from source and target volumes.

For each preset, you can customize various advanced options. You can access these settings by clicking on the preset. The preset options are shown in Figure 10-18 on page 485.

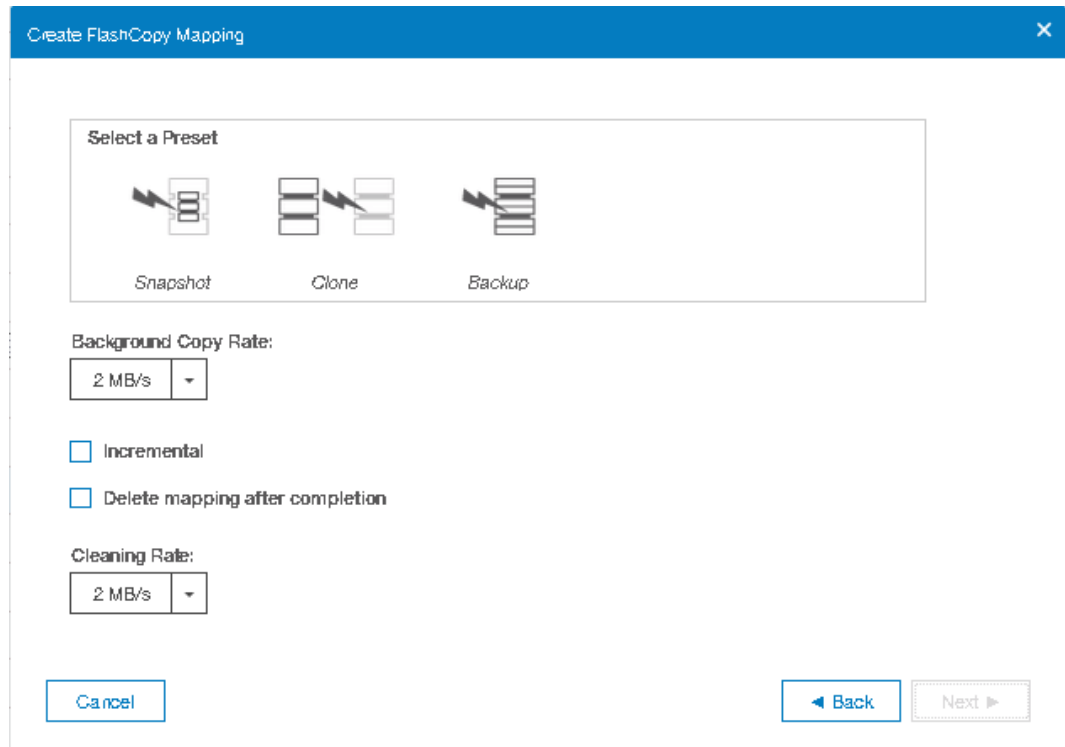


Figure 10-18 Create FlashCopy Mapping Presets

If you prefer not to customize these settings, go directly to step 5.

You can customize the following advanced setting options, as shown in Figure 10-18:

- Background Copy Rate
- Incremental

Incremental FlashCopy mapping: Even if the type of the FlashCopy mapping is incremental, the first copy process copies all of the data from the source volume to the target volume.

- Delete mapping after completion
- Cleaning Rate

After you complete your modifications, click **Next**.

5. You can choose whether to add the mappings to a Consistency Group.

If you want to include this FlashCopy mapping in a Consistency Group, select **Yes, add the mappings to a consistency group** in the window that is shown in Figure 10-19 on page 486. You also can select the Consistency Group from the drop-down list.

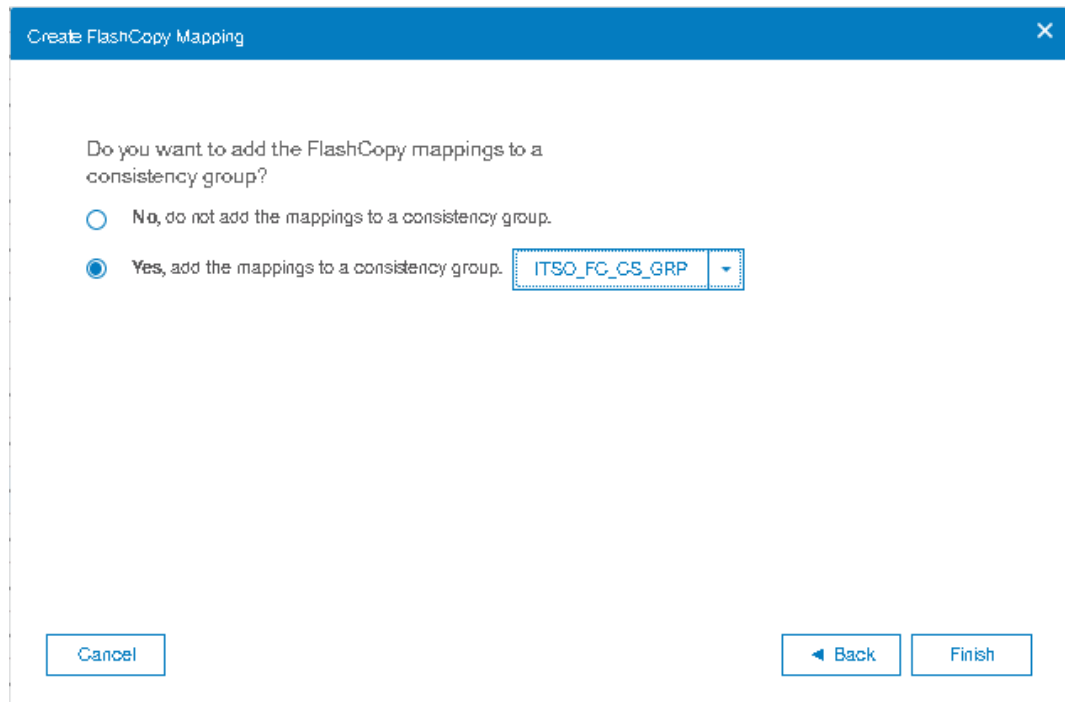
A screenshot of a 'Create FlashCopy Mapping' dialog box. The title bar is blue with the text 'Create FlashCopy Mapping' and a close button (X). The main area is white and contains the question 'Do you want to add the FlashCopy mappings to a consistency group?'. Below the question are two radio button options: 'No, do not add the mappings to a consistency group.' and 'Yes, add the mappings to a consistency group.'. The 'Yes' option is selected. To the right of the 'Yes' option is a dropdown menu showing 'ITSQ_FC_CS_GRP'. At the bottom of the dialog are three buttons: 'Cancel', 'Back', and 'Finish'.

Figure 10-19 Add the mappings to a Consistency Group

Alternatively, if you do not want to include this FlashCopy mapping in a Consistency Group, select **No, do not add the mappings to a consistency group.**

6. Click **Finish**, as shown in Figure 10-20.

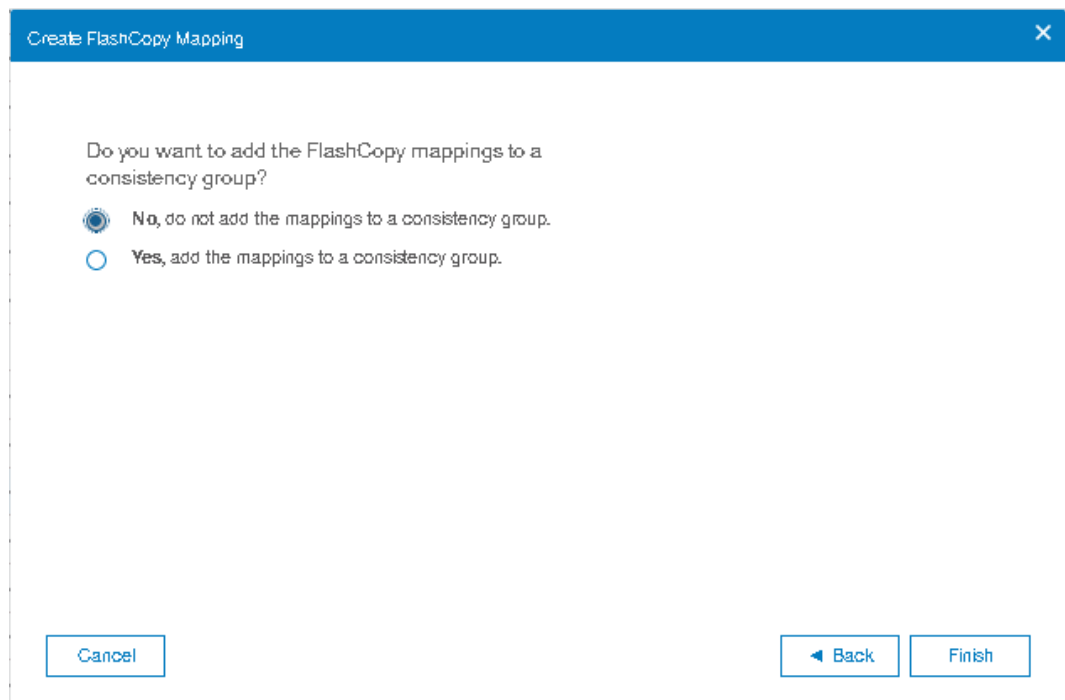
A screenshot of a 'Create FlashCopy Mapping' dialog box, identical to Figure 10-19. The title bar is blue with the text 'Create FlashCopy Mapping' and a close button (X). The main area is white and contains the question 'Do you want to add the FlashCopy mappings to a consistency group?'. Below the question are two radio button options: 'No, do not add the mappings to a consistency group.' and 'Yes, add the mappings to a consistency group.'. The 'No' option is selected. At the bottom of the dialog are three buttons: 'Cancel', 'Back', and 'Finish'.

Figure 10-20 Do not add the mappings to a Consistency Group

7. Check the result of this FlashCopy mapping. From the main panel, click **Copy Services** → **FlashCopy Mappings** as shown in Figure 10-21.

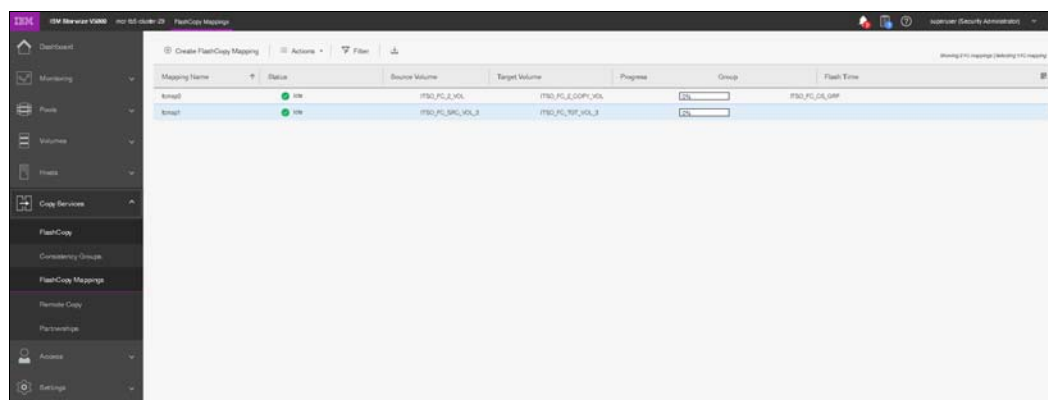


Figure 10-21 FlashCopy maps

8. For each FlashCopy mapping relationship that was created, a mapping name is automatically generated that starts with `fmapX`, where `X` is the next available number. If needed, you can rename these mappings, as shown in 10.4.11, “Renaming FlashCopy mapping” on page 509.
9. The FlashCopy mapping is now ready for use as shown in Figure 10-22.



Figure 10-22 fmap ready to use

Tip: You can start FlashCopy from the GUI. However, the use of the GUI might be impractical if you plan to handle many FlashCopy mappings or Consistency Groups periodically or at varying times. In these cases, creating a script by using the CLI might be more convenient.

Creating target volumes

Complete the following steps to create target volumes for FlashCopy mapping:

1. If you did not create a target volume for this source volume, then select the source volume via left-click, then click **Actions** → **Advanced FlashCopy** → **Create New Target Volumes**, as shown in Figure 10-23 on page 488.

Target volume naming: If the target volume does not exist, the target volume is created. The target volume name is based on its source volume and a generated number at the end, for example, `source_volume_name_XX`, where `XX` is a number that was generated dynamically.

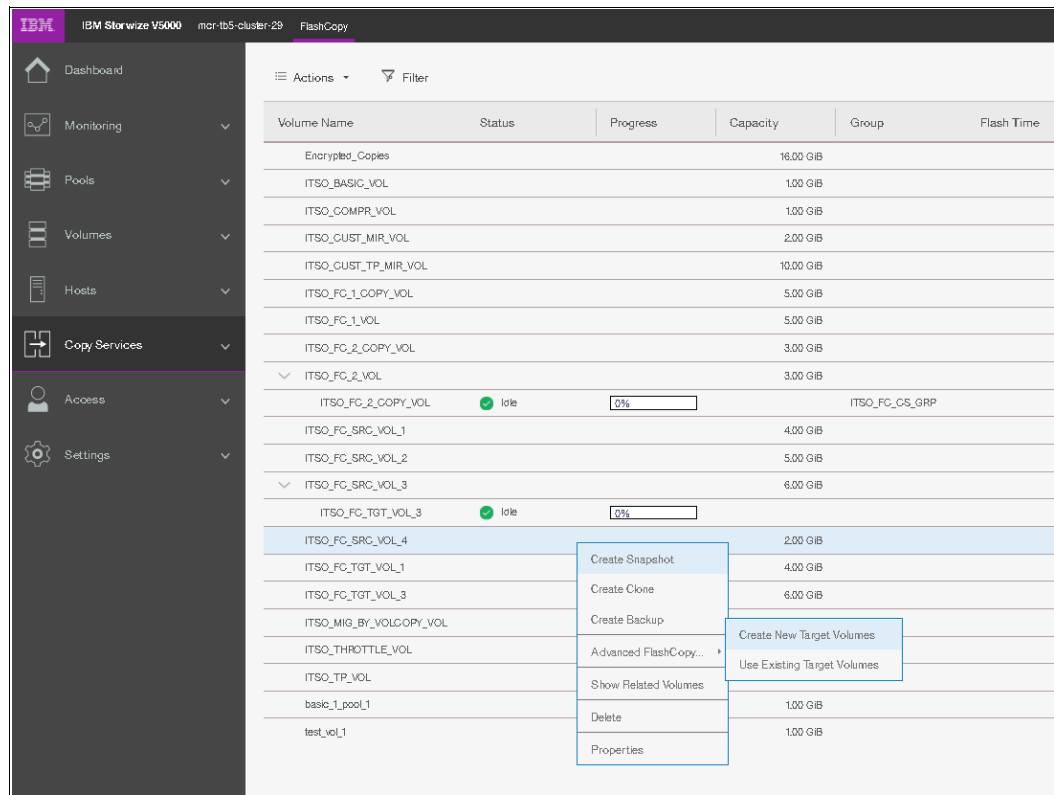


Figure 10-23 Selecting Create New Target Volumes

- In the Create FlashCopy Mapping window (Figure 10-24 on page 489), you must select one FlashCopy preset. The GUI provides the following presets to simplify common FlashCopy operations:

- Snapshot
- Clone
- Backup

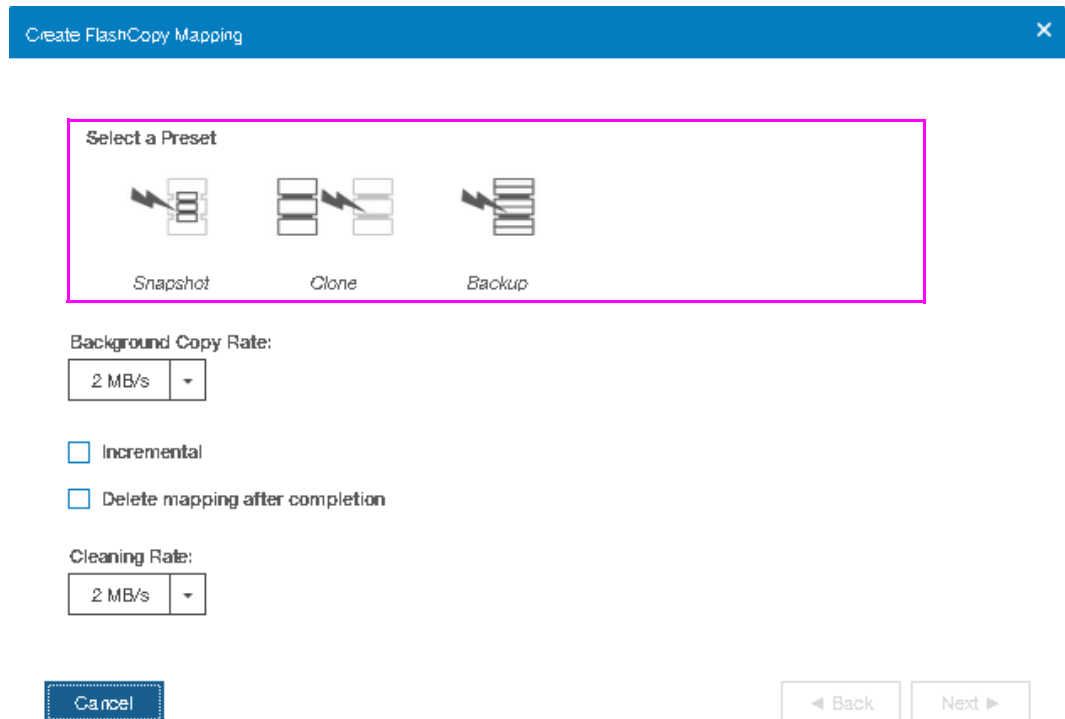


Figure 10-24 Create FlashCopy Mapping window

3. For each preset, you can customize various advanced options as shown in Figure 10-25.




Figure 10-25 Create FlashCopy Mapping Preset customization

If you prefer not to customize these advanced settings, go directly to step 4 on page 490.

You can customize the advanced setting options that are shown in Figure 10-25 on page 489:

- Background Copy Rate
- Incremental

Incremental FlashCopy mapping: Even if the type of the FlashCopy mapping is incremental, the first copy process copies all of the data from the source volume to the target volume.

- Delete mapping after completion (This option automatically deletes a FlashCopy mapping after the background copy is completed. Do not use this option when the background copy rate is set to zero).
- Cleaning Rate

4. You can choose whether to add this FlashCopy mapping to a Consistency Group.

If you want to include this FlashCopy mapping in a Consistency Group, select **Yes, add the mappings to a consistency group** in the next window (Figure 10-26). Select the Consistency Group from the drop-down list.

If you do not want to include this FlashCopy mapping in a Consistency Group, select **No, do not add the mappings to a consistency group**.

5. Click **Next**.

Figure 10-26 Selecting the option to add the mappings to a Consistency Group

6. The next window will show capacity management options about the new target volume as shown in Figure 10-27 on page 491.

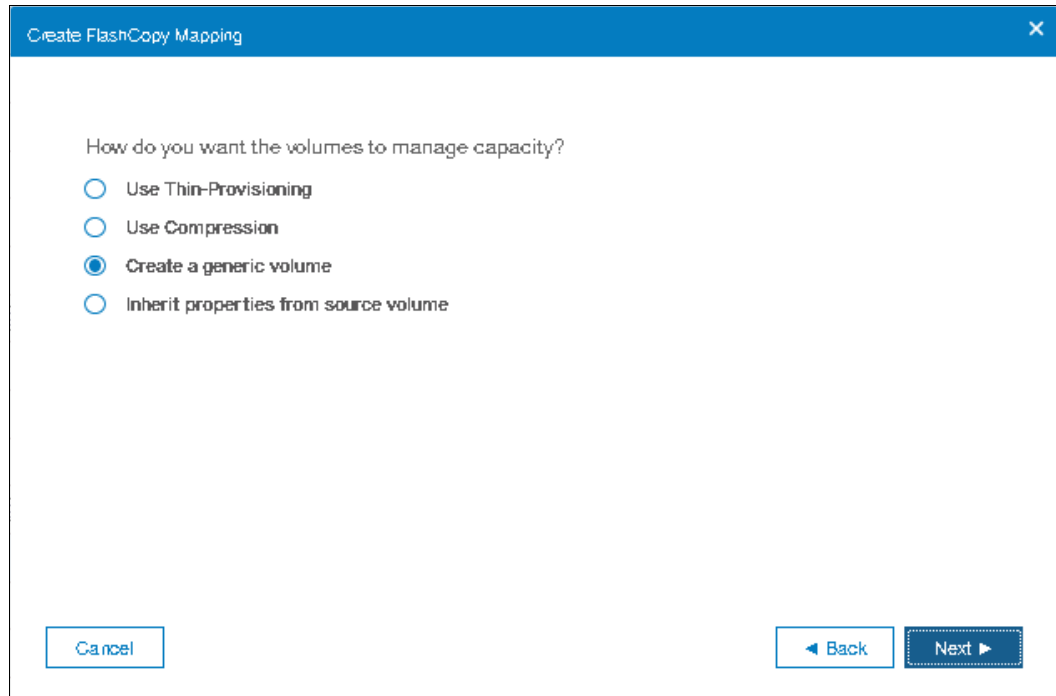


Figure 10-27 FlashCopy target volume capacity management options

7. Click **Next** and you will be asked to select the desired storage pool in which the new FlashCopy target volume needs to be created as shown in Figure 10-28.

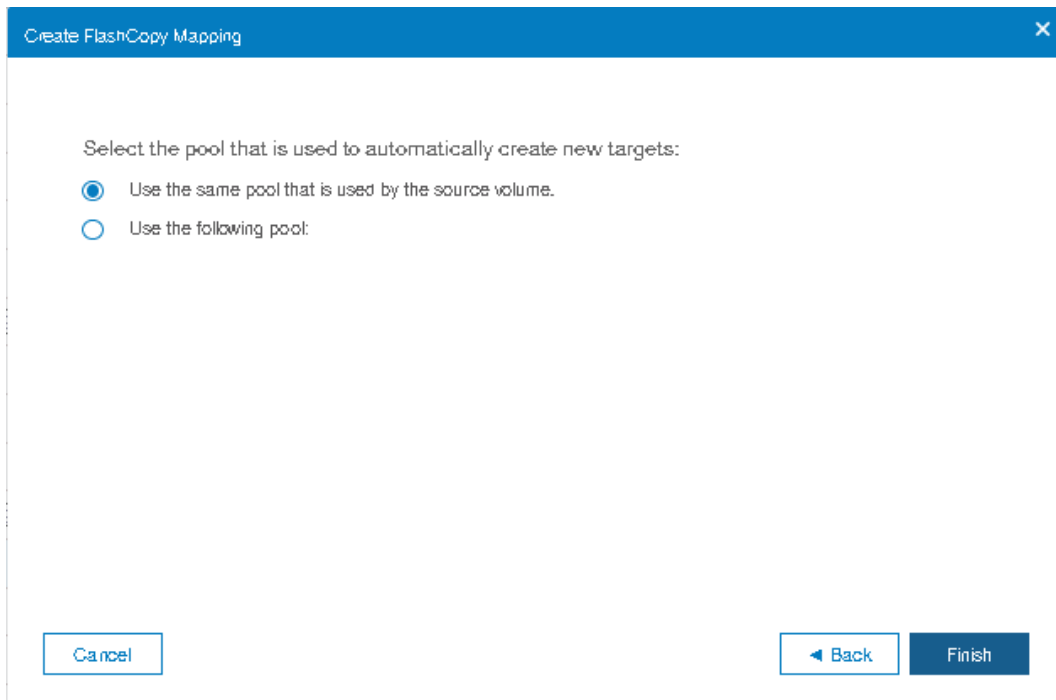


Figure 10-28 Pool selection

8. Click **Finish** and you will see a window indicating the status of the operation as shown in Figure 10-29 on page 492.

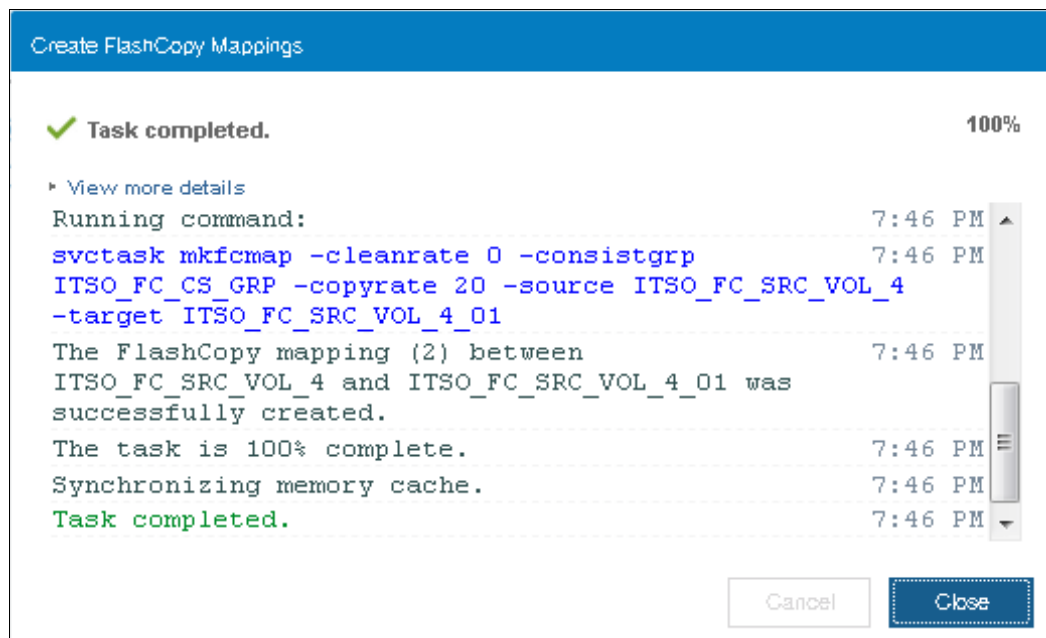


Figure 10-29 FlashCopy map and target volume created

9. Check the result of this FlashCopy mapping, as shown in Figure 10-30. For each FlashCopy mapping relationship that is created, a mapping name is automatically generated that starts with fcmapX where X is the next available number. If necessary, you can rename these mappings, as shown in Figure 10-30. For more information, see 10.4.11, “Renaming FlashCopy mapping” on page 509.

Note: If the FlashCopy target volume is a generic volume and is not ready, then the volume may be getting formatted. Check the running tasks in the GUI.

Volume Name	Status	Progress	Capacity	Group	Flash Time
Examples_Copied			10.00 GB		
ITSO_BM500_VOL			100.00 GB		
ITSO_COMP1_VOL			100.00 GB		
ITSO_GUST_APP_VOL			2.00 GB		
ITSO_GUST_TP_APP_VOL			10.00 GB		
ITSO_FC_1_COPY_VOL			5.00 GB		
ITSO_FC_1_VOL			5.00 GB		
ITSO_FC_2_COPY_VOL			5.00 GB		
ITSO_FC_2_VOL			5.00 GB		
ITSO_FC_2_COPY_VOL			5.00 GB		
ITSO_FC_2_COPY_VOL	Copying	10%	5.00 GB	ITSO_FC_CS_GRP	Oct 9, 2016 8:04:29 PM
ITSO_FC_SRC_VOL_1			4.00 GB		
ITSO_FC_SRC_VOL_2			5.00 GB		
ITSO_FC_SRC_VOL_3			4.00 GB		
ITSO_FC_SRC_VOL_4	New	10%	2.00 GB		
ITSO_FC_SRC_VOL_4_01	Copying	10%	2.00 GB	ITSO_FC_CS_GRP	Oct 9, 2016 8:04:29 PM
ITSO_FC_SRC_VOL_4_01			2.00 GB		
ITSO_FC_TST_VOL_1			4.00 GB		
ITSO_FC_TST_VOL_3			6.00 GB		
ITSO_MQ_BF_VOL_COPY_VOL			3.00 GB		
ITSO_THRDTITLE_VOL			5.00 GB		
ITSO_TP_VOL			100.00 GB		
base1_image1			100.00 GB		

Figure 10-30 FlashCopy mapping

The FlashCopy mapping is ready for use.

Tip: You can start FlashCopy from the GUI. However, the use of the GUI might be impractical if you plan to handle many FlashCopy mappings or Consistency Groups periodically or at varying times. In these cases, creating a script by using the CLI might be more convenient.

10.4.2 Single-click snapshot

The *snapshot* creates a point-in-time backup of production data. The snapshot is not intended to be an independent copy. Instead, it is used to maintain a view of the production data at the time that the snapshot is created. Therefore, the snapshot holds only the data from regions of the production volume that changed since the snapshot was created. Because the snapshot preset uses thin provisioning, only the capacity that is required for the changes is used.

Snapshot uses the following preset parameters:

- ▶ Background copy: No
- ▶ Incremental: No
- ▶ Delete after completion: No
- ▶ Cleaning rate: No
- ▶ Primary copy source pool: Target pool

To create and start a snapshot, complete the following steps:

1. From the main panel, click **Copy Services** → **FlashCopy**.
2. Select the volume that you want to create a snapshot of and click **Actions** → **Create Snapshot**, as shown in Figure 10-31 on page 494.

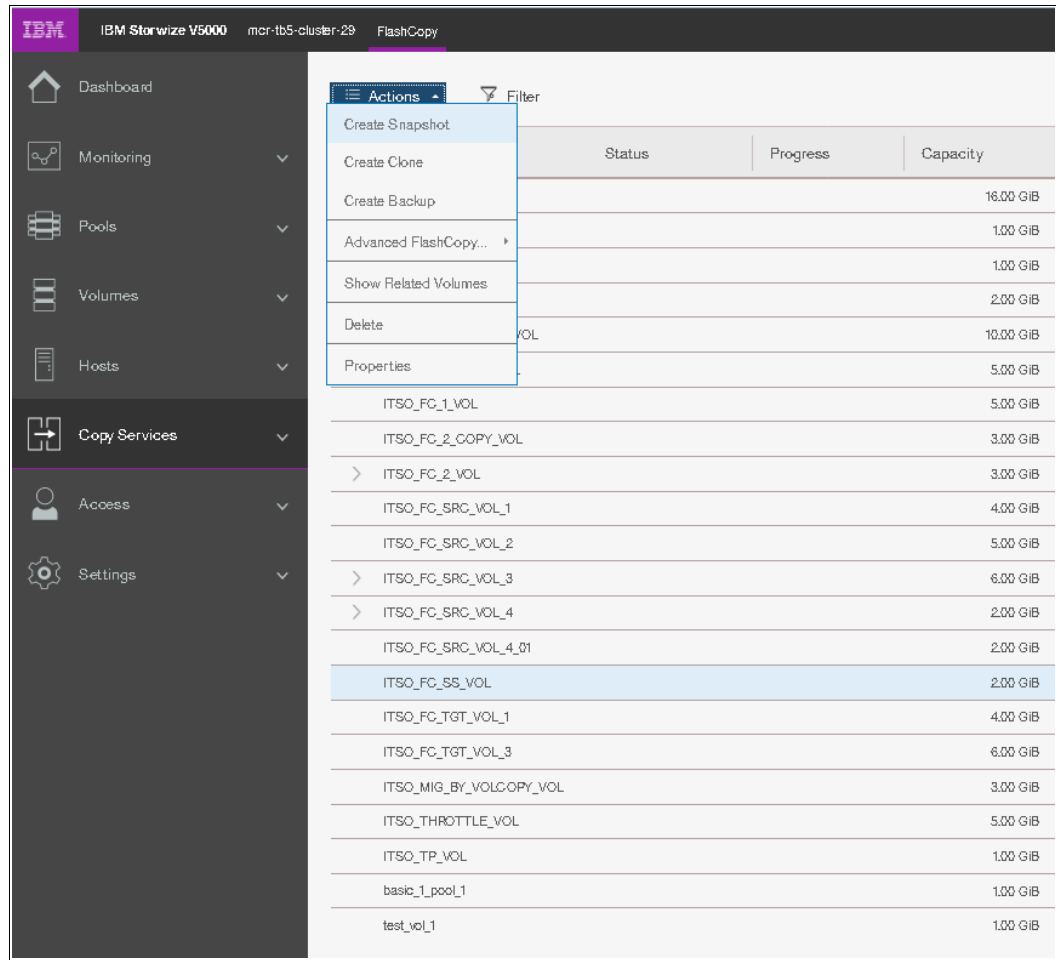


Figure 10-31 Create Snapshot option

3. A volume is created as a target volume for this snapshot in the same pool as the source volume. The FlashCopy mapping is created and started.

You can check the FlashCopy progress in the Progress column Status area, as shown in Figure 10-32 on page 495.

Actions ▾		Filter			
Volume Name	Status	Progress	Capacity	Group	Flash Time
Encrypted_Copies			16.00 GiB		
ITSO_BASIC_VOL			1.00 GiB		
ITSO_COMPR_VOL			1.00 GiB		
ITSO_CUST_MIR_VOL			2.00 GiB		
ITSO_CUST_TP_MIR_VOL			10.00 GiB		
ITSO_FC_1_COPY_VOL			5.00 GiB		
ITSO_FC_1_VOL			5.00 GiB		
ITSO_FC_2_COPY_VOL			3.00 GiB		
> ITSO_FC_2_VOL			3.00 GiB		
ITSO_FC_SRC_VOL_1			4.00 GiB		
ITSO_FC_SRC_VOL_2			5.00 GiB		
> ITSO_FC_SRC_VOL_3			6.00 GiB		
> ITSO_FC_SRC_VOL_4			2.00 GiB		
ITSO_FC_SRC_VOL_4_01			2.00 GiB		
▽ ITSO_FC_SS_VOL			2.00 GiB		
ITSO_FC_SS_VOL_01	Copying	0%			Oct 10, 2017, 12:09:58 PM
ITSO_FC_SS_VOL_01			2.00 GiB		
ITSO_FC_TGT_VOL_1			4.00 GiB		
ITSO_FC_TGT_VOL_3			6.00 GiB		
ITSO_MIG_BY_VOLCOPY_VOL			3.00 GiB		
ITSO_THROTTLE_VOL			5.00 GiB		
ITSO_TP_VOL			1.00 GiB		
basic_1_pool_1			1.00 GiB		

Figure 10-32 Snapshot created and started

10.4.3 Single-click clone

The *clone preset* creates an exact replica of the volume, which can be changed without affecting the original volume. After the copy completes, the mapping that was created by the preset is automatically deleted.

The clone preset uses the following parameters:

- ▶ Background copy rate: 50
- ▶ Incremental: No
- ▶ Delete after completion: Yes
- ▶ Cleaning rate: 50
- ▶ Primary copy source pool: Target pool

To create and start a clone, complete the following steps:

1. From the main panel, click **Copy Services** → **FlashCopy**.
2. Select the volume that you want to clone.
3. Click **Actions** → **Create Clone**, as shown in Figure 10-33 on page 496.

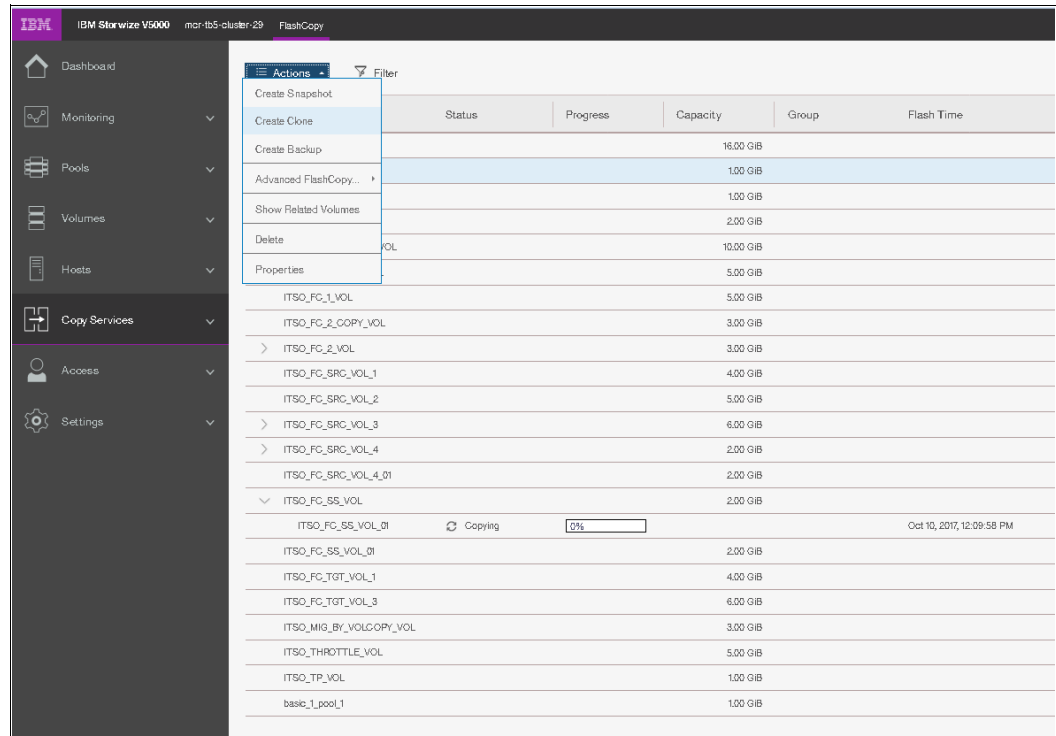


Figure 10-33 Create Clone option

A volume is created as a target volume for this clone in the same pool as the source volume. The FlashCopy mapping is created and started. You can check the FlashCopy progress in the FlashCopy mappings option or by clicking on to the **Tasks** → **Running Tasks** from the main pane. After the FlashCopy clone is created, the mapping is removed and the new cloned volume becomes available, as shown in Figure 10-34.

Actions		Filter			
Volume Name	Status	Progress	Capacity	Group	Flash Time
Encrypted_Copies			16.00 GiB		
ITSO_BASIC_VOL			1.00 GiB		
ITSO_BASIC_VOL_01			1.00 GiB		
ITSO_COMPR_VOL			1.00 GiB		

Figure 10-34 Clone created and FlashCopy relationship removed

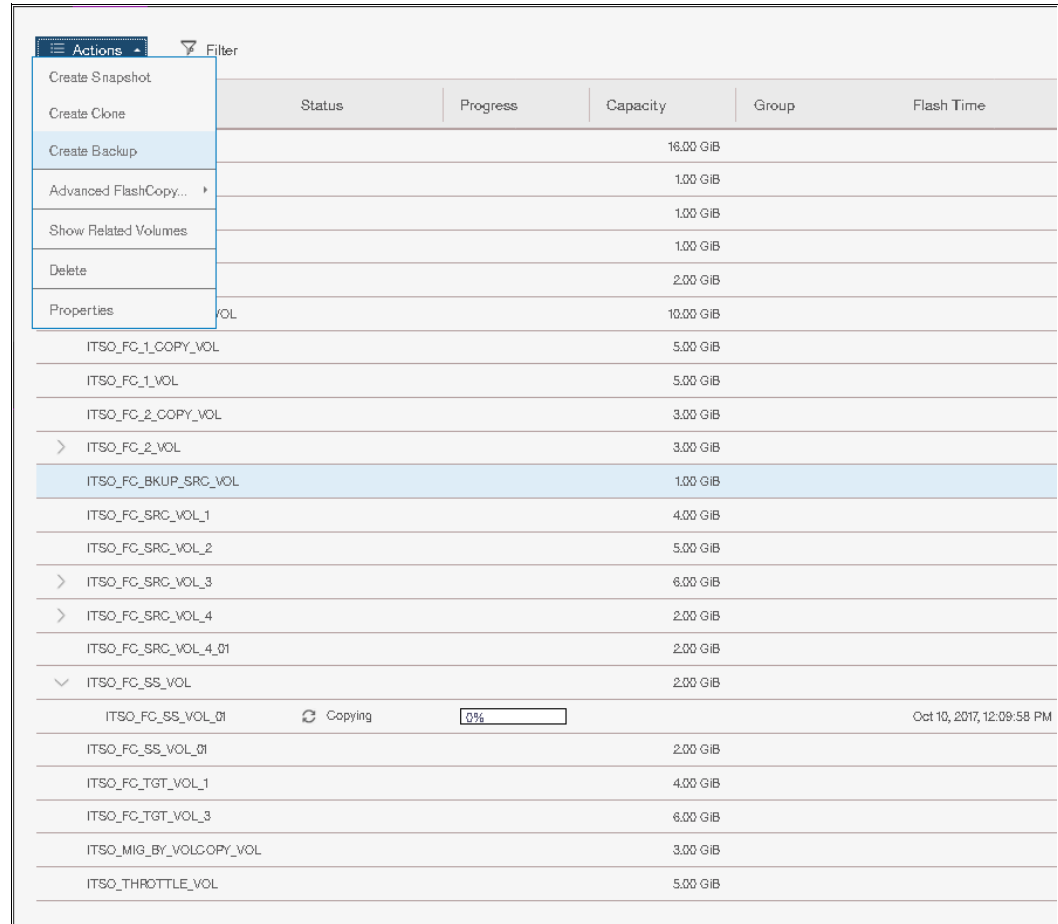
10.4.4 Single-click backup

The backup creates a point-in-time replica of the production data. After the copy completes, the backup view can be refreshed from the production data, with minimal data copied from the production volume to the backup volume. The backup preset uses the following parameters:

- ▶ Background Copy rate: 50
- ▶ Incremental: Yes
- ▶ Delete after completion: No
- ▶ Cleaning rate: 50
- ▶ Primary copy source pool: Target pool

To create and start a backup, complete the following steps:

1. From the main panel, click **Copy Services** → **FlashCopy**.
2. Select the volume that you want to back up, and click **Actions** → **Create Backup**, as shown in Figure 10-35.



	Status	Progress	Capacity	Group	Flash Time
ITSO_FC_1_COPY_VOL			16.00 GiB		
ITSO_FC_1_VOL			1.00 GiB		
ITSO_FC_2_COPY_VOL			1.00 GiB		
ITSO_FC_2_VOL			1.00 GiB		
ITSO_FC_BKUP_SRC_VOL			2.00 GiB		
ITSO_FC_SRC_VOL_1			10.00 GiB		
ITSO_FC_SRC_VOL_2			5.00 GiB		
ITSO_FC_SRC_VOL_3			3.00 GiB		
ITSO_FC_SRC_VOL_4			3.00 GiB		
ITSO_FC_SRC_VOL_4_01			4.00 GiB		
ITSO_FC_SS_VOL			5.00 GiB		
ITSO_FC_SS_VOL_01	Copying	0%	2.00 GiB		Oct 10, 2017, 12:09:58 PM
ITSO_FC_TGT_VOL_1			2.00 GiB		
ITSO_FC_TGT_VOL_3			4.00 GiB		
ITSO_MIG_BY_VOLCOPY_VOL			6.00 GiB		
ITSO_THROTTLE_VOL			3.00 GiB		
			5.00 GiB		

Figure 10-35 Create Backup option

3. A volume is created as a target volume for this backup in the same pool as the source volume. The FlashCopy mapping is created and started.

You can check the FlashCopy progress in the Progress column, as shown in Figure 10-36 on page 498, or in the Running Tasks from the main panel.

Actions ▾		Filter			
Volume Name	Status	Progress	Capacity	Group	Flash Time
Encrypted_Copies			16.00 GiB		
ITSO_BASIC_VOL			1.00 GiB		
ITSO_BASIC_VOL_01			1.00 GiB		
ITSO_COMPR_VOL			1.00 GiB		
ITSO_CUST_MIR_VOL			2.00 GiB		
ITSO_CUST_TP_MIR_VOL			10.00 GiB		
ITSO_FC_1_COPY_VOL			5.00 GiB		
ITSO_FC_1_VOL			5.00 GiB		
ITSO_FC_2_COPY_VOL			3.00 GiB		
> ITSO_FC_2_VOL			3.00 GiB		
▽ ITSO_FC_BKUP_SRC_VOL			1.00 GiB		
ITSO_FC_BKUP_SRC_VO...	Copying	1%			Oct 10, 2017, 12:49:00 PM
ITSO_FC_BKUP_SRC_VOL_01			1.00 GiB		
ITSO_FC_SRC_VOL_1			4.00 GiB		
ITSO_FC_SRC_VOL_2			5.00 GiB		
> ITSO_FC_SRC_VOL_3			6.00 GiB		
> ITSO_FC_SRC_VOL_4			2.00 GiB		
ITSO_FC_SRC_VOL_4_01			2.00 GiB		
> ITSO_FC_SS_VOL			2.00 GiB		
ITSO_FC_SS_VOL_01			2.00 GiB		
ITSO_FC_TGT_VOL_1			4.00 GiB		
ITSO_FC_TGT_VOL_3			6.00 GiB		
ITSO_MIG_BY_VOLCOPY_VOL			3.00 GiB		

Figure 10-36 Backup created and started

10.4.5 Creating a FlashCopy Consistency Group

To create a FlashCopy Consistency Group in the GUI, complete the following steps:

1. From the main panel, click **Copy Services** → **Consistency Group**. The Consistency Groups pane opens, as shown in Figure 10-37.

Create Consistency Group		Actions ▾		Filter	
Mapping Name	Status	Source Volume	Target Volume	Progress	Flash Time
Not in a Group					

Figure 10-37 Consistency Groups pane

2. Click **Create Consistency Group** and enter the FlashCopy Consistency Group name that you want to use and click **Create** (Figure 10-38 on page 499).

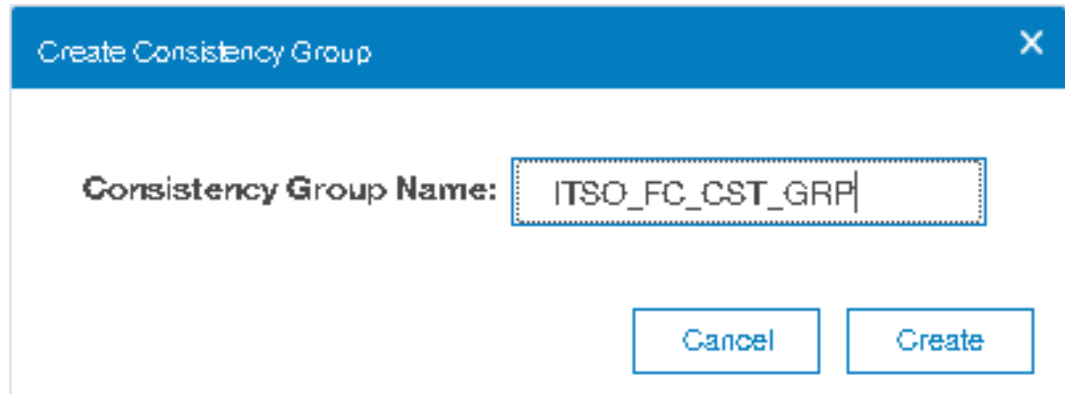


Figure 10-38 Create Consistency Group window

Consistency Group name: You can use the letters A - Z and a - z, the numbers 0 - 9, and the underscore (_) character. The volume name can be 1 - 63 characters.

Figure 10-39 shows the result.




Figure 10-39 New Consistency Group

10.4.6 Creating FlashCopy mappings in a Consistency Group

This section describes how to create FlashCopy mappings for volumes and their related targets. The source and target volumes were created before this operation.

Complete the following steps:

1. From the main pane, click **Copy Services** → **Consistency Group**. The Consistency Groups pane opens, as shown in Figure 10-40.




Figure 10-40 Consistency Group pane

2. Select in which Consistency Group you want to create the FlashCopy mapping. If you prefer not to create a FlashCopy mapping in a Consistency Group, select **Not in a Group** as shown in Figure 10-41.




Figure 10-41 Consistency Group selection

3. If you select a new Consistency Group, click **Actions** → **Create FlashCopy Mapping**, as shown in Figure 10-42.

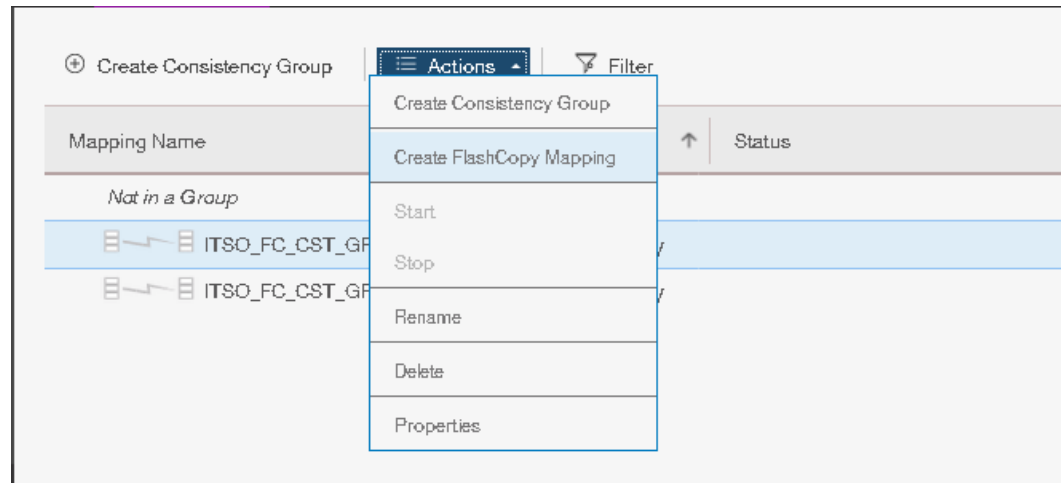


Figure 10-42 Create FlashCopy Mapping action for a Consistency Group

4. If you did not select a Consistency Group, click **Create FlashCopy Mapping**, as shown in Figure 10-43.

Consistency Groups: If no Consistency Group is defined, the mapping is a stand-alone mapping. It can be prepared and started without affecting other mappings. All mappings in the same Consistency Group must have the same status to maintain the consistency of the group.

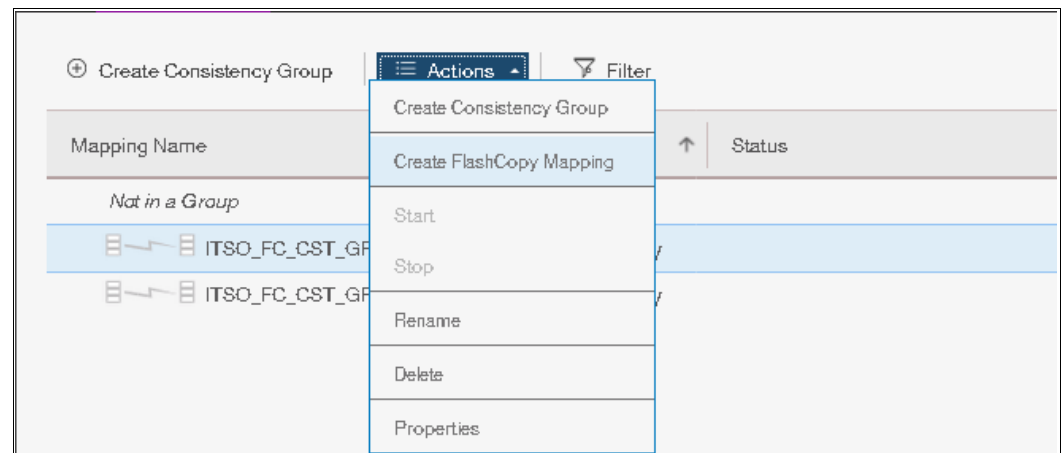


Figure 10-43 Create FlashCopy Mapping

5. The Create FlashCopy Mapping window opens, as shown in Figure 10-44 on page 501. In this window, you must create the relationships between the source volumes (the volumes that are copied) and the target volumes (the volumes that receive the copy). A mapping can be created between any two volumes in a clustered system.

Important: The source volume and the target volume must be of equal size.

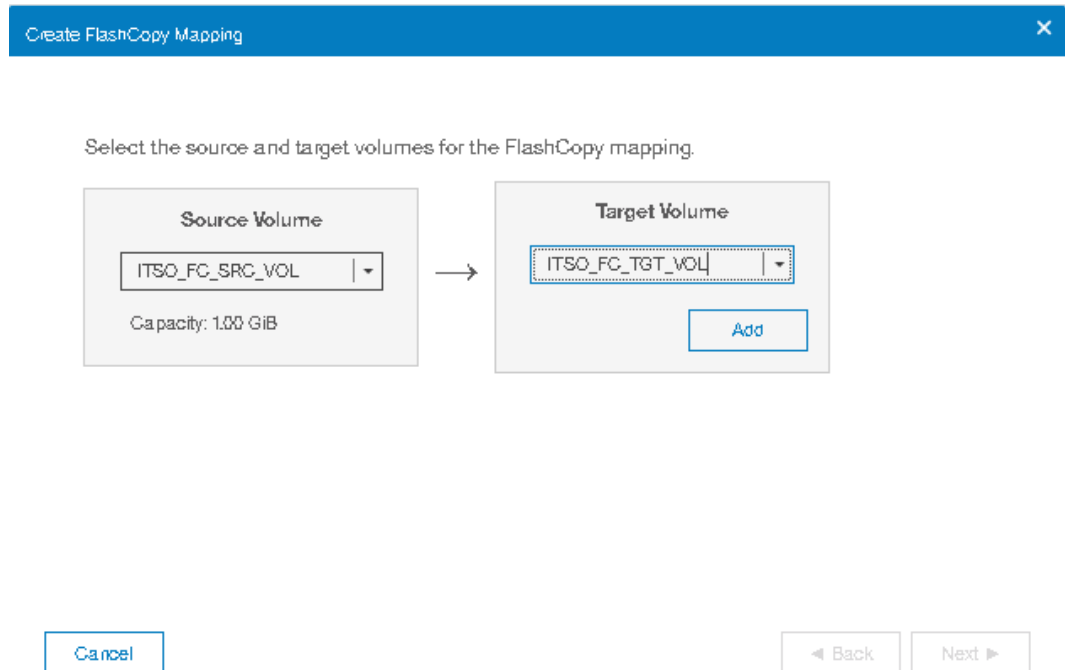


Figure 10-44 Create FlashCopy Mapping window

Tip: The volumes do not have to be in the same I/O Group or storage pool.

6. Select a volume in the Source Volume column by using the drop-down list. Then, select a volume in the Target Volume column by using the drop-down list. Click **Add**, as shown in Figure 10-45. Repeat this step to create other relationships.

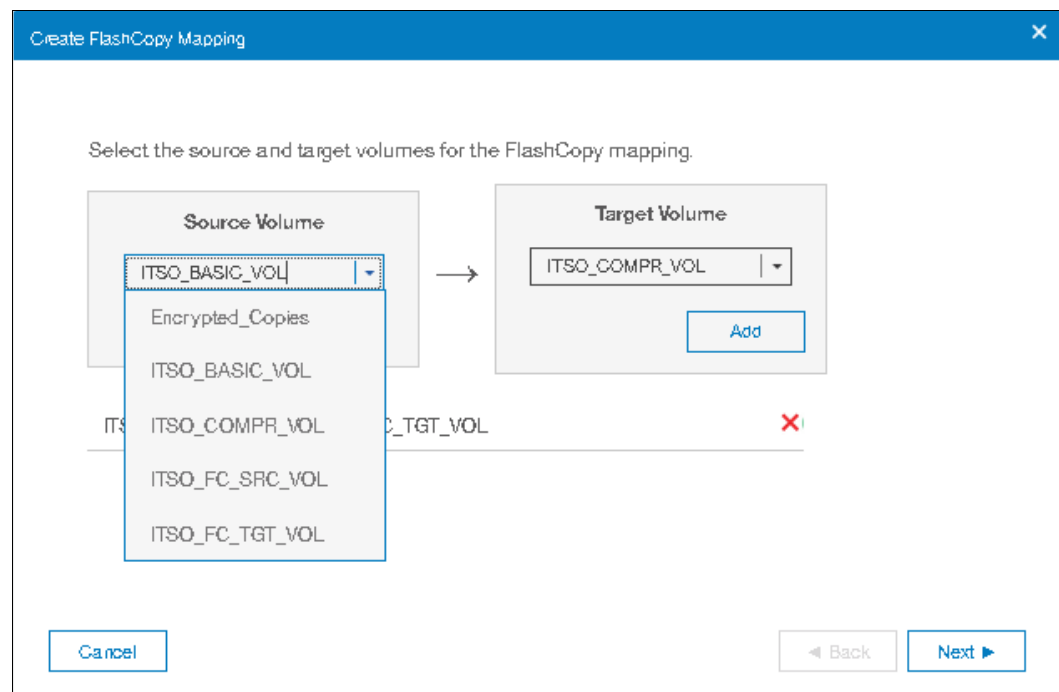


Figure 10-45 Selecting source and target volumes

To remove a relationship that was created, click **X**.

Important: The source and target volumes must be of equal size. Therefore, only the targets with the appropriate size are shown for a source volume.

7. Click **Next** after all of the relationships that you want to create are shown (Figure 10-46).

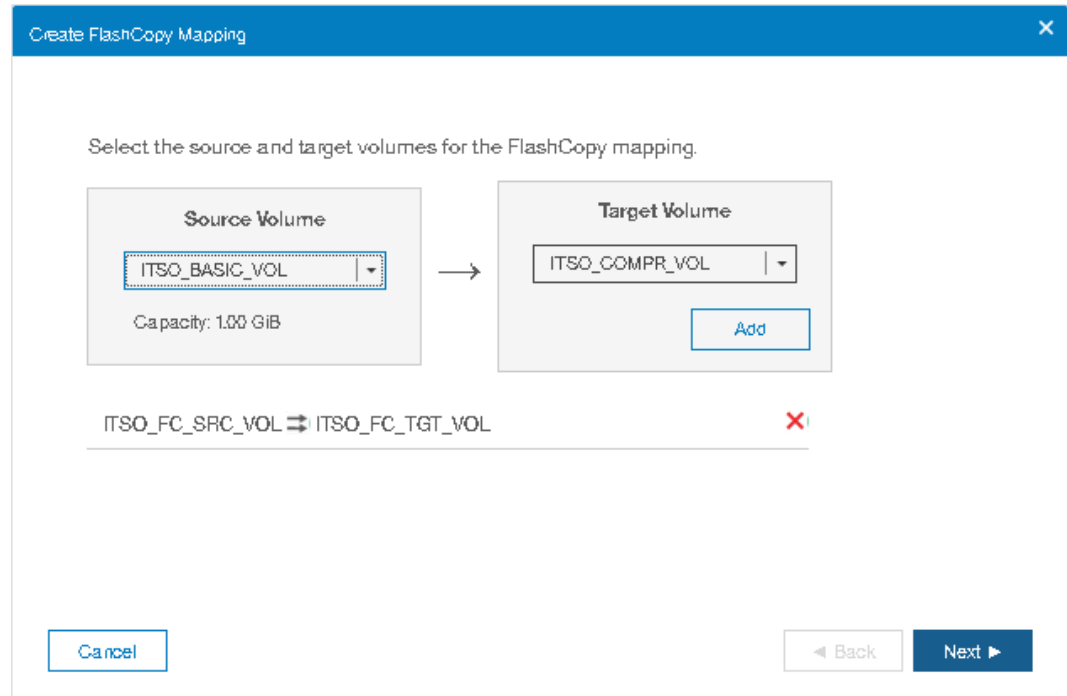


Figure 10-46 Create FlashCopy Mapping with the relationships that were created

8. In the next window, you must select one FlashCopy preset along with their customization options. The GUI provides the following presets to simplify common FlashCopy operations, as shown in Figure 10-47 on page 503:
 - Snapshot: Creates a copy-on-write point-in-time copy.
 - Clone: Creates an exact replica of the source volume on a target volume. The copy can be changed without affecting the original volume.
 - Backup: Creates a FlashCopy mapping that can be used to recover data or objects if the system experiences data loss. These backups can be copied multiple times from the source and target volumes.

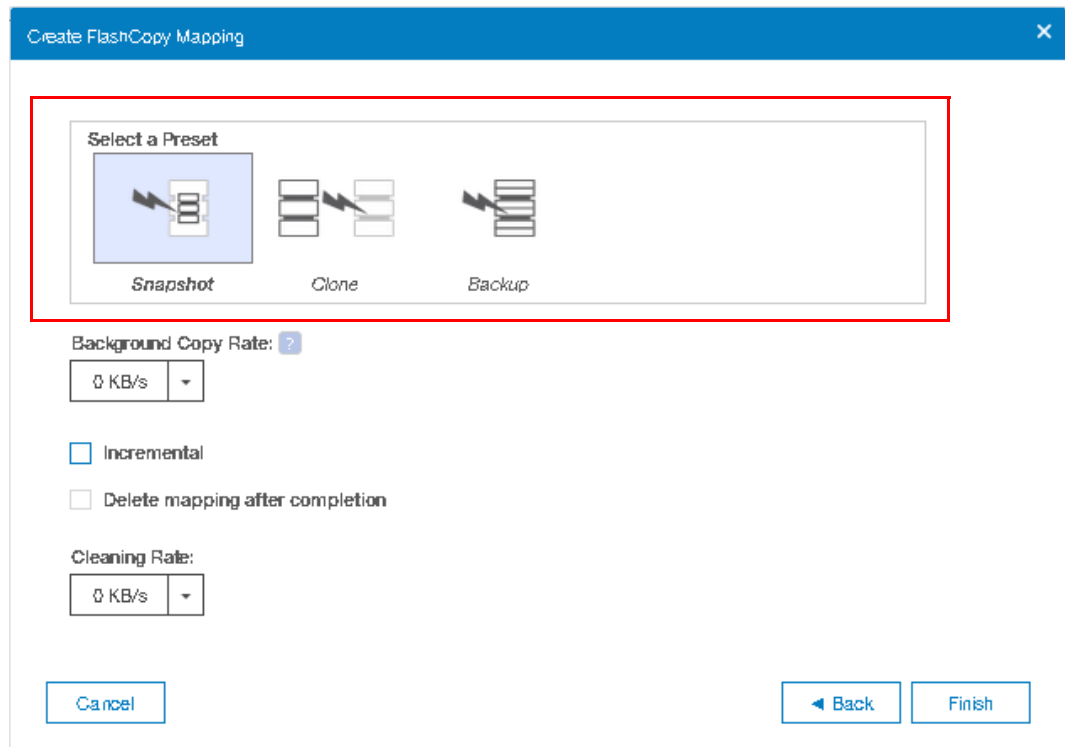


Figure 10-47 Create FlashCopy Mapping window

Whichever preset you select, based on that you can customize options based on the selected preset as shown in Figure 10-48.

If you prefer not to customize these settings, go directly to step 9.

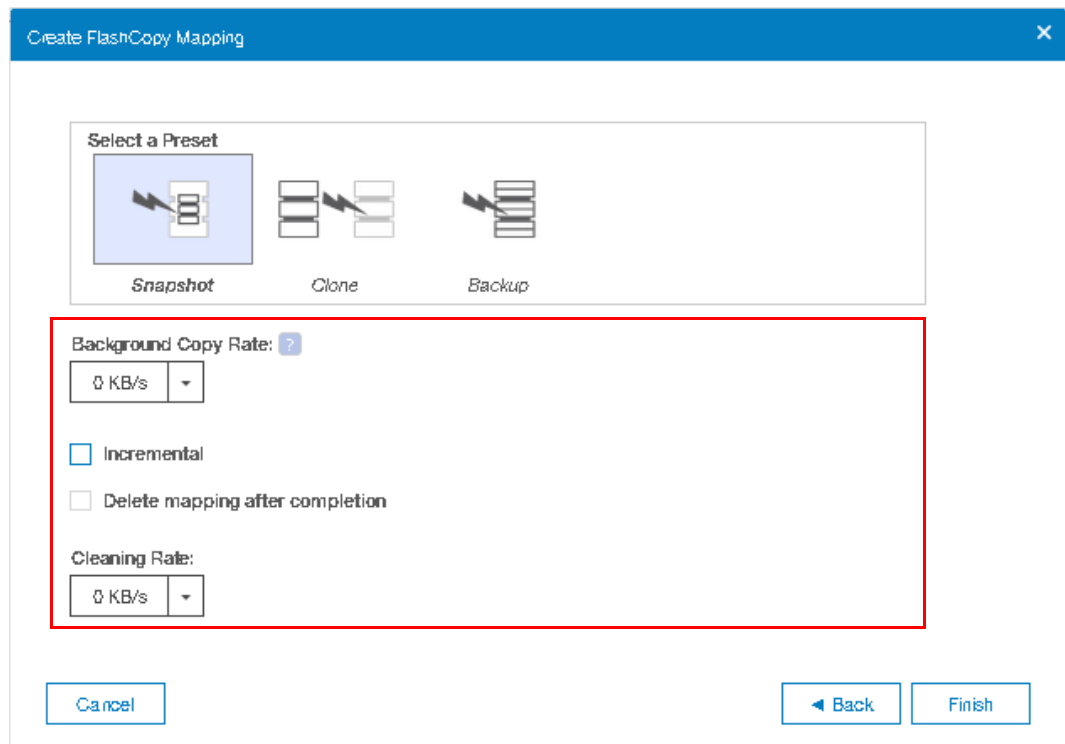


Figure 10-48 Create FlashCopy Mapping: Preset customization

9. Click **Finish**.
10. Check the result of this FlashCopy mapping in the Consistency Groups window, as shown in Figure 10-49.

For each FlashCopy mapping relationship that you created, a mapping name is automatically generated that starts with `fcmapX` where `X` is an available number. If necessary, you can rename these mappings. For more information, see 10.4.11, “Renaming FlashCopy mapping” on page 509.



Mapping Name	Status	Source Volume	Target Volume	Progress	Flash Time
Not in a Group					
fcmap0	Idle or Copied	ITSD_FC_SRC_VOL	ITSD_FC_TGT_VOL	100%	
ITSD_FC_GRP_2	Empty				

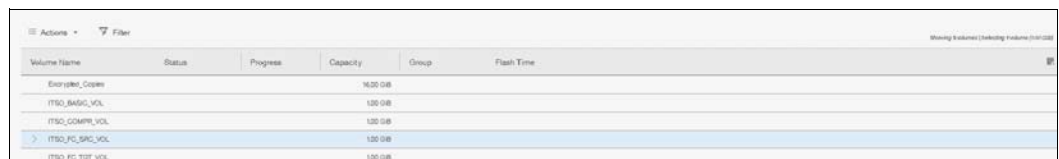
Figure 10-49 Create FlashCopy mappings result

Tip: You can start FlashCopy from the GUI. However, if you plan to handle many FlashCopy mappings or Consistency Groups periodically, or at varying times, creating a script by using the operating system shell CLI might be more convenient.

10.4.7 Showing related volumes

Complete the following steps to show related volumes for a specific FlashCopy mapping:

1. From main pane, click **Copy Services** → **FlashCopy**.
2. Select the volume (from the FlashCopy pane only) or the FlashCopy mapping that you want to view in this Consistency Group as shown in Figure 10-50



Volume Name	Status	Progress	Capacity	Group	Flash Time
Encrypted_Copies			1638 GB		
ITSD_AWDC_VOL			120 GB		
ITSD_COMPR_VOL			120 GB		
ITSD_FC_SRC_VOL			120 GB		
ITSD_FC_TGT_VOL			120 GB		

Figure 10-50 Volumes

3. Click **Actions** → **Show Related Volumes**, as shown in Figure 10-51 on page 505.

Tip: You can also right-click a FlashCopy mapping and select **Show Related Volumes**.

Actions ▾ Filter					
Volume Name	Status	Progress	Capacity	Group	Flash Time
Encrypted_Copies			16.00 GiB		
ITSO_BASIC_VOL			1.00 GiB		
ITSO_COMPR_VOL			1.00 GiB		
> ITSO_FC_SRC_VOL			1.00 GiB		
ITSO_FC_TGT_VOL			1.00 GiB		

Create Snapshot
Create Clone
Create Backup
Advanced FlashCopy... ▸
Show Related Volumes
Delete
Properties

Figure 10-51 Show Related Volumes

In the Related Volumes window (Figure 10-52), you can see the related mapping for a volume. If you click one of these volumes, you can see its properties.




Figure 10-52 Related Volumes

10.4.8 Moving a FlashCopy mapping to a Consistency Group

Complete the following steps to move a FlashCopy mapping to the Consistency Group:

1. From the main panel, click **Copy Services** → **FlashCopy** or **Consistency Groups**, or **FlashCopy Mappings**.
2. Select the FlashCopy mapping that you want to move to a Consistency Group or the FlashCopy mapping for which you want to change the Consistency Group.
3. Click **Actions** → **Move to Consistency Group**, as shown in Figure 10-53 on page 506.

Tip: You can also right-click a FlashCopy mapping and select **Move to Consistency Group**.




Figure 10-53 Move to Consistency Group action

4. In the Move FlashCopy Mapping to Consistency Group window, select the Consistency Group for this FlashCopy mapping by using the drop-down list (Figure 10-54).




Figure 10-54 Move FlashCopy mapping to Consistency Group window

5. Click **Move to Consistency Group** to confirm your changes.

10.4.9 Removing a FlashCopy mapping from a Consistency Group

Complete the following steps to remove a FlashCopy mapping from a Consistency Group:

1. From the main panel, click **Copy Services** → **FlashCopy** or **Consistency Groups**, or **FlashCopy Mappings**.
2. Select the FlashCopy mapping that you want to remove from a Consistency Group.
3. Click **Actions** → **Remove from Consistency Group**, as shown in Figure 10-55.

Tip: You can also right-click a FlashCopy mapping and select **Remove from Consistency Group**.




Figure 10-55 Remove from Consistency Group action

4. In the Remove FlashCopy Mapping from Consistency Group window, click **Remove**, as shown in Figure 10-56 on page 507.

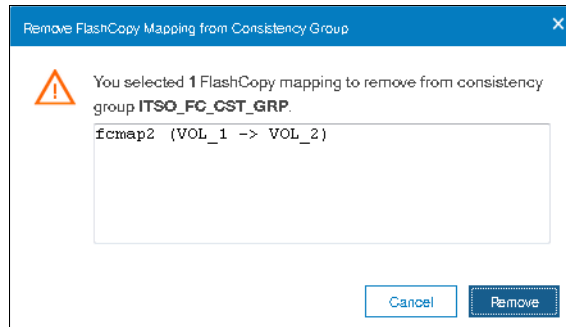


Figure 10-56 Remove FlashCopy Mapping from Consistency Group

5. Click **Remove** and the desired FlashCopy map will be removed from the consistency group.

10.4.10 Modifying a FlashCopy mapping

Complete the following steps to modify a FlashCopy mapping:

1. From the main panel, click **Copy Services** → **FlashCopy** or **Consistency Groups**, or **FlashCopy Mappings**.
2. In the table, select the FlashCopy mapping that you want to modify.
3. Click **Actions** → **Edit Properties**, as shown in Figure 10-57.




Figure 10-57 Edit Properties

Tip: You can also right-click a FlashCopy mapping and select **Edit Properties**.

4. In the Edit FlashCopy Mapping window, you can modify the Background Copy Rate from the drop-down as parameters for a selected FlashCopy mapping, as shown in Figure 10-58 on page 508.

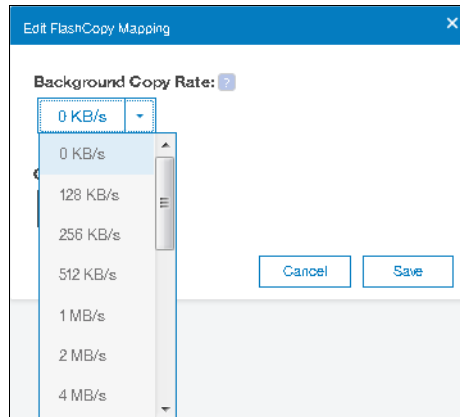


Figure 10-58 Modifying Background Copy Rate

Note: Background Copy Rate: This option determines the priority that is given to the copy process. A faster rate increases the priority of the process, which might affect the performance of other operations.

For FlashCopy background copy rates, starting from V7.8.1, the background copy rate up to 2 GB/s.

5. In the Edit Flash Copy Mapping window, you can modify the Cleaning Rate from the drop-down for the selected FlashCopy mapping as shown in Figure 10-59.

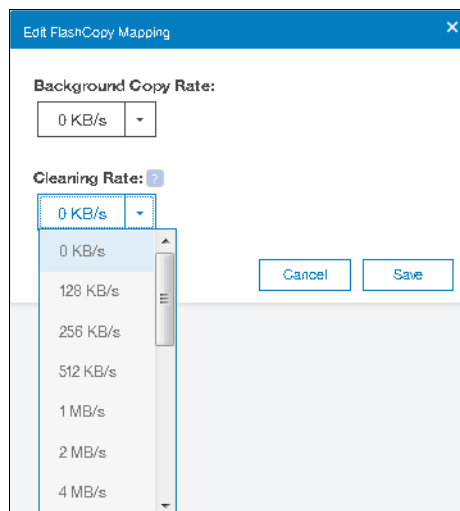


Figure 10-59 Modifying Cleaning rate

Note: Cleaning Rate: This option minimizes the amount of time that a mapping is in the stopping state. If the mapping is not complete, the target volume is offline while the mapping is stopping.

For FlashCopy background cleaning rates, starting from V7.8.1, the background cleaning rate up to 2 GB/s.

6. Click **Save** to confirm your changes.

10.4.11 Renaming FlashCopy mapping

Complete the following steps to rename a FlashCopy mapping:

1. From the main panel, click **Copy Services** → **FlashCopy** or **Consistency Groups**, or **FlashCopy Mappings**.
2. In the table, select the FlashCopy mapping that you want to rename.
3. Click **Actions** → **Rename Mapping**, as shown in Figure 10-60.

Tip: You can also right-click a FlashCopy mapping and select **Rename Mapping**.




Figure 10-60 Rename Mapping action

4. In the Rename FlashCopy Mapping window, enter the new name that you want to assign to the FlashCopy mapping and click **Rename**, as shown in Figure 10-61.




Figure 10-61 Renaming a FlashCopy mapping

FlashCopy mapping name: You can use the letters A - Z and a - z, the numbers 0 - 9, and the underscore (_) character. The FlashCopy mapping name can be 1 - 63 characters.

10.4.12 Renaming Consistency Group

To rename a Consistency Group, complete the following steps:

1. From the main panel, click **Copy Services** → **Consistency Groups**.
2. From the left pane, select the Consistency Group that you want to rename. Then, select **Actions** → **Rename**, as shown in Figure 10-62.



Figure 10-62 Renaming a Consistency Group

3. Enter the new name that you want to assign to the Consistency Group and click **Rename**, as shown in Figure 10-63.




Figure 10-63 Changing the name for a Consistency Group

Consistency Group name: The name can consist of the letters A - Z and a - z, the numbers 0 - 9, the dash (-), and the underscore (_) character. The name can be 1 - 63 characters. However, the name cannot start with a number, a dash, or an underscore.

The new Consistency Group name is displayed in the Consistency Group pane.

10.4.13 Deleting FlashCopy mapping

Complete the following steps to delete a FlashCopy mapping:

1. From the main panel, click **Copy Services** → **FlashCopy** or **Consistency Groups**, or **FlashCopy Mappings**.
2. In the table, select the FlashCopy mapping that you want to delete.

Selecting multiple FlashCopy mappings: To select multiple FlashCopy mappings, hold down Ctrl and click the other entries that you want to delete. This capability is only available in the Consistency Groups pane and the FlashCopy Mappings pane.

3. Click **Actions** → **Delete Mapping**, as shown in Figure 10-64.

Tip: You can also right-click a FlashCopy mapping and select **Delete Mapping**.



Figure 10-64 Selecting the Delete Mapping option

4. The Delete FlashCopy Mapping window opens, as shown in Figure 10-65 on page 511. In the “Verify the number of FlashCopy mappings that you are deleting” field, you must enter the number of volumes that you want to remove. This verification was added to help avoid deleting the wrong mappings.

If you still have target volumes that are inconsistent with the source volumes and you want to delete these FlashCopy mappings, select **Delete the FlashCopy mapping even when**

the data on the target volume is inconsistent, or if the target volume has other dependencies.

Click **Delete**, as shown in Figure 10-65.

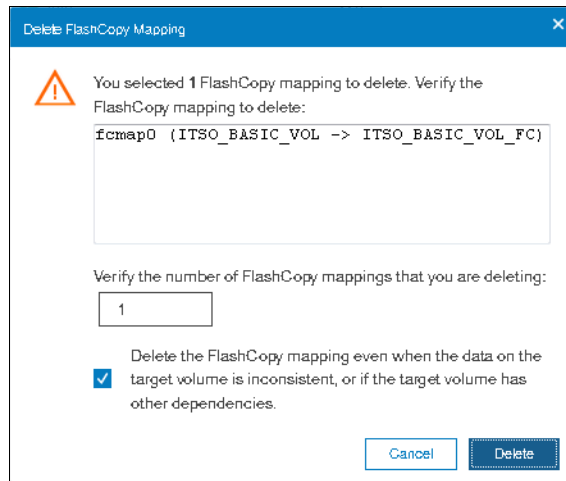


Figure 10-65 Delete FlashCopy Mapping

10.4.14 Deleting FlashCopy Consistency Group

Important: Deleting a Consistency Group does not delete the FlashCopy mappings.

Complete the following steps to delete a FlashCopy Consistency Group:

1. From the main panel, click **Copy Services** → **Consistency Groups**.
2. Select the FlashCopy Consistency Group that you want to delete.
3. Click **Actions** → **Delete**, as shown in Figure 10-66.




Figure 10-66 Delete Consistency Group action

4. The Warning window opens, as shown in Figure 10-67. Click **Yes**.




Figure 10-67 Warning window

10.4.15 Starting FlashCopy process

When the FlashCopy mapping is created, the copy process can be started. Only mappings that are not members of a Consistency Group can be started individually. Complete the following steps:

1. From the main panel, click **Copy Services** → **FlashCopy Mappings**.
2. In the table, choose the FlashCopy mapping that you want to start.
3. Click **Actions** → **Start** (as shown in Figure 10-68) to start the FlashCopy process.

Tip: You can also right-click a FlashCopy mapping and select **Start**.




Figure 10-68 Start the FlashCopy process action

4. You can check the FlashCopy progress in the Progress column of the table or in the Running Tasks status area. After the task completes, the FlashCopy mapping status is in a Copied state, as shown in Figure 10-69.



Figure 10-69 Checking the FlashCopy progress

10.4.16 Stopping FlashCopy process

When a FlashCopy copy process is stopped, the target volume becomes invalid and it is set offline by the system. The FlashCopy mapping copy must be retrIGGERED to bring the target volume online again.

Important: Stop a FlashCopy copy process only when the data on the target volume is not useful and can be discarded, or if you want to modify the FlashCopy mapping. When a FlashCopy mapping is stopped, the target volume becomes invalid and it is set offline by the system.

Complete the following steps to stop a FlashCopy copy process:

1. From the main panel, click **Copy Services** → **FlashCopy Mappings**.
2. Choose the FlashCopy mapping that you want to stop.
3. Click **Actions** → **Stop** (as shown in Figure 10-70 on page 513) to stop the FlashCopy Consistency Group copy process.




Figure 10-70 Stopping the FlashCopy copy process

The FlashCopy Mapping status changes to Stopped, as shown in Figure 10-71.



Figure 10-71 FlashCopy Mapping status

10.5 Volume mirroring and migration options

Volume mirroring is a simple RAID 1-type function that enables a volume to remain online even when the storage pool that is backing it becomes inaccessible. Volume mirroring is designed to protect the volume from storage infrastructure failures by seamless mirroring between storage pools.

Volume mirroring is provided by a specific volume mirroring function in the I/O stack, and it cannot be manipulated like a FlashCopy or other types of copy volumes. However, this feature provides migration functionality, which can be obtained by splitting the mirrored copy from the source, or by using the *migrate to* function. Volume mirroring cannot control backend storage mirroring or replication.

With volume mirroring, host I/O completes when both copies are written, and this feature is enhanced with a tunable latency tolerance. This tolerance provides an option to give preference to losing the redundancy between the two copies. This tunable timeout value is Latency or Redundancy.

The Latency tuning option, which is set with **chvdisk -mirrowritepriority latency**, is the default. It prioritizes host I/O latency, which yields a preference to host I/O over availability. However, you might need to give preference to redundancy in your environment when availability is more important than I/O response time. Use the **chvdisk -mirror writepriority redundancy** command to set the redundancy option.

Regardless of which option you choose, volume mirroring can provide extra protection for your environment.

Migration offers the following options:

- **Export to Image mode.** By using this option, you can move storage from managed mode to image mode, which is useful if you are using the Lenovo storage V series as a migration device. For example, vendor A's product cannot communicate with vendor B's product, but you must migrate existing data from vendor A to vendor B. By using Export to image mode, you can migrate data by using Copy Services functions and then return control to the native array while maintaining access to the hosts.

- **Import to Image mode.** By using this option, you can import an existing storage MDisk or logical unit number (LUN) with its existing data from an external storage system without putting metadata on it so that the existing data remains intact. After you import it, all copy services functions can be used to migrate the storage to other locations while the data remains accessible to your hosts.
- **Volume migration by using volume mirroring and then by using Split into New Volume.** By using this option, you can use the available RAID 1 functionality. You create two copies of data that initially has a set relationship (one volume with two copies, one primary and one secondary) but then break the relationship (two volumes, both primary and no relationship between them) to make them independent copies of data.

You can use this to migrate data between storage pools and devices. You might use this option if you want to move volumes to multiple storage pools. Each volume can have two copies at a time, so you can add only one copy to the original volume, and then you must split those copies to create another copy of the volume.

- **Volume migration by using move to another pool.** By using this option, you can move any volume between storage pools without any interruption to the host access. This option is a quicker version of the “Volume Mirroring and Split into New Volume” option. You might use this option if you want to move volumes in a single step, or you do not have a volume mirror copy already.

Migration: While these migration methods do not disrupt access, you must take a brief outage to install the host drivers for your Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems if you do not already have them installed.

With volume mirroring, you can move data to different MDisk within the same storage pool or move data between different storage pools. Using volume mirroring over volume migration is beneficial because with volume mirroring, storage pools do not need to have the same extent size as is the case with volume migration.

Note: Volume mirroring does not create a second volume before you split copies. Volume mirroring adds a second copy of the data under the same volume, so you end up having one volume presented to the host with two copies of data connected to this volume. Only splitting copies creates another volume, and then both volumes have only one copy of the data.

Starting with V7.3 and the introduction of the new cache architecture, mirrored volume performance has been significantly improved. Now, lower cache is beneath the volume mirroring layer, which means that both copies have their own cache.

This approach helps in cases of having copies of different types, for example generic and compressed, because now both copies use its independent cache and performs its own read prefetch. Destaging of the cache can now be done independently for each copy, so one copy does not affect performance of a second copy.

Also, because the IBM Storwize for Lenovo and Lenovo Storage V series destage algorithm is MDisk aware, it can tune or adapt the destaging process, depending on MDisk type and usage, for each copy independently.

10.6 Native IP replication

Before we describe Remote Copy features that benefit from the use of multiple Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems, it is important to describe the partnership option introduced with V7.2 native IP replication.

10.6.1 Native IP replication technology

Remote Mirroring over IP communication is supported on The Lenovo storage V series systems by using Ethernet communication links. The IP replication uses innovative *Bridgeworks SANSlide* technology to optimize network bandwidth and utilization. This new function enables the use of a lower-speed and lower-cost networking infrastructure for data replication.

Bridgeworks' SANSlide technology, which is integrated into the controller firmware, uses artificial intelligence to help optimize network bandwidth use and adapt to changing workload and network conditions.

This technology can improve remote mirroring network bandwidth usage up to three times, which can enable clients to deploy a less costly network infrastructure, or speed up remote replication cycles to enhance disaster recovery effectiveness.

With an Ethernet network data flow, the data transfer can slow down over time. This condition occurs because of the latency that is caused by waiting for the acknowledgment of each set of packets that are sent. The next packet set cannot be sent until the previous packet is acknowledged, as shown in Figure 10-72.

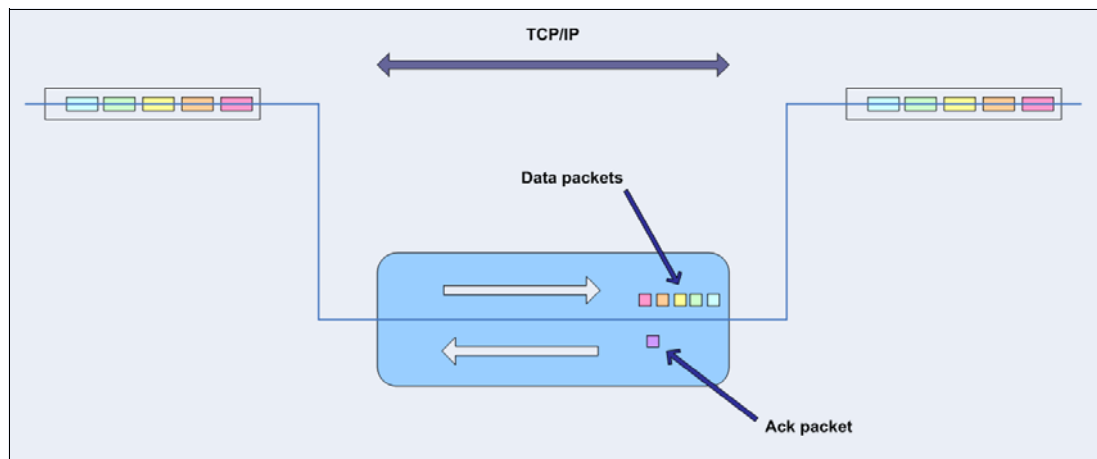


Figure 10-72 Typical Ethernet network data flow

However, by using the embedded IP replication, this behavior can be eliminated with the enhanced parallelism of the data flow by using multiple virtual connections (VC) that share IP links and addresses. The artificial intelligence engine can dynamically adjust the number of VCs, receive window size, and packet size as appropriate to maintain optimum performance. While the engine is waiting for one VC's ACK, it sends more packets across other VCs. If packets are lost from any VC, data is automatically retransmitted, as shown in Figure 10-73.

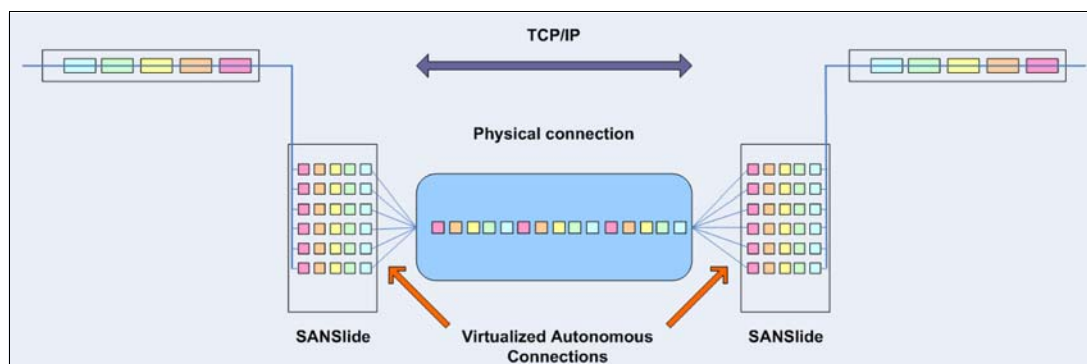


Figure 10-73 Optimized network data flow by using Bridgeworks SANSlide technology

With native IP partnership, the following Copy Services features are supported:

- Metro Mirror (MM)

Referred to as *synchronous replication*, MM provides a consistent copy of a source virtual disk on a target virtual disk. Data is written to the target virtual disk synchronously after it is written to the source virtual disk so that the copy is continuously updated.

- Global Mirror (GM) and GM with Change Volumes

Referred to as *asynchronous replication*, GM provides a consistent copy of a source virtual disk on a target virtual disk. Data is written to the target virtual disk asynchronously so that the copy is continuously updated. However, the copy might not contain the last few updates if a disaster recovery (DR) operation is performed. An added extension to GM is GM with Change Volumes. GM with Change Volumes is the preferred method for use with native IP replication.

10.6.2 Lenovo Storage V series System Layers

A Lenovo storage V-series family system can be in one of the two layers: the *replication* layer or the *storage* layer. The system layer affects how the system interacts with Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems. The IBM SAN Volume Controller is always set to replication layer and this parameter is unchangeable.

In the storage layer, an IBM Storwize for Lenovo family system has the following characteristics and requirements:

- The system can perform MM and GM replication with other storage-layer systems.
- The system can provide external storage for replication-layer systems or IBM SAN Volume Controller.
- The system cannot use a storage-layer system as external storage.

In the replication layer, an IBM SAN Volume Controller or an IBM Storwize for Lenovo family system has the following characteristics and requirements:

- ▶ The system can perform MM and GM replication with other replication-layer systems or IBM SAN Volume Controller.
- ▶ The system cannot provide external storage for a replication-layer system or an IBM SAN Volume Controller.
- ▶ The system can use a storage-layer system as external storage.

An IBM Storwize for Lenovo family system is in the storage layer by default, but the layer can be changed. For example, you might want to change an IBM Storwize V5000 for Lenovo to a replication layer to complete Global Mirror or Metro Mirror replication with a SAN Volume Controller system.

To change the storage layer of an existing IBM Storwize for Lenovo system with internal storage before you add a second system into the SAN zone, you do not need to stop I/O operations. However, if the system has Fibre Channel (FC) connections to another IBM Storwize for Lenovo family or SAN Volume Controller system in the SAN fabric, I/O operations must be stopped temporarily. In this scenario, the FC ports must be disabled (for example, by unplugging all the FC ports, changing zoning, disabling switch ports) before you change the system layer. Then, you must re-enable the FC ports.

Note: Before you change the layer of an IBM Storwize for Lenovo family system, the following conditions must be met:

- ▶ No host object can be configured with worldwide port names (WWPNs) from an IBM Storwize for Lenovo family system.
- ▶ No system partnerships can be defined.
- ▶ No IBM Storwize for Lenovo family system can be visible on the SAN fabric.

In your IBM Storwize for Lenovo system, use the **lssystem** command to check the current system layer, as shown in Example 10-2.

Example 10-2 Output from lssystem command showing the system layer

```
IBM_Storwize:ITS0_5K:superuser>lssystem
id 000001002140020E
name ITS0_V5K
...
lines omitted for brevity
...
easy_tier_acceleration off
has_nas_key no
layer replication
...
```

Note: Consider the following rules for creating remote partnerships between the IBM SAN Volume Controller and IBM Storwize for Lenovo Family systems:

- ▶ An IBM SAN Volume Controller is always in the replication layer.
- ▶ By default, the IBM Storwize for Lenovo systems are in the storage layer but can be changed to the replication layer.
- ▶ A system can form partnerships only with systems in the same layer.
- ▶ Starting in software V6.4, an IBM SAN Volume Controller or IBM Storwize for Lenovo system in the replication layer can virtualize an IBM Storwize for Lenovo in the storage layer.

10.6.3 IP partnership limitations

The following prerequisites and assumptions must be considered before IP partnership between two Lenovo Storage V-series systems can be established:

- ▶ The IBM Storwize for Lenovo systems are successfully installed with V7.2 or later code levels.
- ▶ The systems must have the necessary licenses that enable remote copy partnerships to be configured between two systems. No separate license is required to enable IP partnership.
- ▶ The storage SANs are configured correctly and the correct infrastructure to support the Spectrum Virtualize systems in remote copy partnerships over IP links is in place.
- ▶ The two systems must be able to ping each other and perform the discovery.
- ▶ The maximum number of partnerships between the local and remote systems, including both IP and Fibre Channel (FC) partnerships, is limited to the current maximum that is supported, which is three partnerships (four systems total).
- ▶ Only a single partnership over IP is supported.
- ▶ A system can have simultaneous partnerships over FC and IP, but with separate systems. The FC zones between two systems must be removed before an IP partnership is configured.
- ▶ IP partnerships are supported on both 10 gigabits per second (Gbps) links and 1 Gbps links. However, the intermix of both on a single link is not supported.
- ▶ The maximum supported round-trip time is 80 milliseconds (ms) for 1 Gbps links.
- ▶ The maximum supported round-trip time is 10 ms for 10 Gbps links.
- ▶ The minimum supported link bandwidth is 10 Mbps.
- ▶ The inter-cluster heartbeat traffic uses 1 Mbps per link.
- ▶ Only nodes from two I/O Groups can have ports that are configured for an IP partnership.
- ▶ Migrations of remote copy relationships directly from FC-based partnerships to IP partnerships are not supported.
- ▶ IP partnerships between the two systems can be over IPv4 or IPv6 only, but not both.
- ▶ Virtual LAN (VLAN) tagging of the IP addresses that are configured for remote copy is supported starting with V7.4.0.
- ▶ Management IP and Internet SCSI (iSCSI) IP on the same port can be in a different network starting with V7.4.0.
- ▶ An added layer of security is provided by using Challenge Handshake Authentication Protocol (CHAP) authentication.
- ▶ Transmission Control Protocol (TCP) ports 3260 and 3265 are used for IP partnership communications. Therefore, these ports must be open in firewalls between the systems.
- ▶ Only a single Remote Copy (RC) data session per physical link can be established. It is intended that only one connection (for sending/receiving Remote Copy data) is made for each independent physical link between the systems.

Note: A physical link is the physical IP link between the two sites, A (local) and B (remote). Multiple IP addresses on local system A could be connected (by Ethernet switches) to this physical link. Similarly, multiple IP addresses on remote system B could be connected (by Ethernet switches) to the same physical link. At any point in time, only a single IP address on cluster A can form an RC data session with an IP address on cluster B.

- ▶ The maximum throughput is restricted based on the use of 1 Gbps or 10 Gbps Ethernet ports, and varies based on distance (for example, round-trip latency) and quality of communication link (for example, packet loss):
 - One 1 Gbps port might transfer up to 110 megabytes per second (MBps) unidirectional, 190 MBps bidirectional
 - Two 1 Gbps ports might transfer up to 220 MBps unidirectional, 325 MBps bidirectional
 - One 10 Gbps port might transfer up to 240 MBps unidirectional, 350 MBps bidirectional
 - Two 10 Gbps port might transfer up to 440 MBps unidirectional, 600 MBps bidirectional

Note: The Bandwidth setting definition when the IP partnerships are created changed. Previously, the bandwidth setting defaulted to 50 MB, and was the maximum transfer rate from the primary site to the secondary site for initial sync/resyncs of volumes.

The Link Bandwidth setting is now configured by using megabits (Mb) not MB. You set the Link Bandwidth setting to a value that the communication link can sustain, or to what is allocated for replication. The Background Copy Rate setting is now a percentage of the Link Bandwidth. The Background Copy Rate setting determines the available bandwidth for the initial sync and resyncs or for GM with Change Volumes.

For further information on IP replication requirements and limitations, and supported configurations, please see Information Center for Lenovo Storage V3700 V2, V3700 V2 XP and V5030 at:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_ippartnershipreqs.html

10.6.4 VLAN support

Starting with V7.4.0, VLAN tagging is supported for both iSCSI host attachment and IP replication. Hosts and remote-copy operations can connect to the system through Ethernet ports. Each traffic type has different bandwidth requirements, which can interfere with each other if they share the same IP connections. VLAN tagging creates two separate connections on the same IP network for different types of traffic. The system supports VLAN configuration on both IPv4 and IPv6 connections.

When the VLAN ID is configured for the IP addresses that are used for either iSCSI host attach or IP replication, the appropriate VLAN settings on the Ethernet network and servers must be configured correctly in order not to experience connectivity issues. After the VLANs are configured, changes to the VLAN settings will disrupt iSCSI and IP replication traffic to and from the partnerships.

During the VLAN configuration for each IP address, the VLAN settings for the local and failover ports on two nodes of an I/O Group can differ. To avoid any service disruption, switches must be configured so the failover VLANs are configured on the local switch ports and the failover of IP addresses from a failing node to a surviving node succeeds. If failover

VLANs are not configured on the local switch ports, there are no paths to Lenovo Storage V3700 V2, V3700 V2 XP and V5030 system node canisters during a node canister failure and the replication fails.

Consider the following requirements and procedures when implementing VLAN tagging:

- ▶ VLAN tagging is supported for IP partnership traffic between two systems.
- ▶ VLAN provides network traffic separation at the layer 2 level for Ethernet transport.
- ▶ VLAN tagging by default is disabled for any IP address of a node port. You can use the CLI or GUI to optionally set the VLAN ID for port IPs on both systems in the IP partnership.
- ▶ When a VLAN ID is configured for the port IP addresses that are used in remote copy port groups, appropriate VLAN settings on the Ethernet network must also be properly configured to prevent connectivity issues.

Setting VLAN tags for a port is disruptive. Therefore, VLAN tagging requires that you stop the partnership first before you configure VLAN tags. Then, restart again when the configuration is complete.

For further information on configuring VLAN for IP replication, please see Information Center for Lenovo Storage V3700 V2, V3700 V2 XP and V5030 at:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_vlanconfigipreplication.html

10.6.5 IP partnership and terminology

The IP partnership terminology and abbreviations that are used are listed in Table 10-7.

Table 10-7 Terminology for IP partnership

IP partnership terminology	Description
Remote copy group or Remote copy port group	The following numbers group a set of IP addresses that are connected to the same physical link. Therefore, only IP addresses that are part of the same remote copy group can form remote copy connections with the partner system: <ul style="list-style-type: none"> ▶ 0 – Ports that are not configured for remote copy ▶ 1 – Ports that belong to remote copy port group 1 ▶ 2 – Ports that belong to remote copy port group 2 Each IP address can be shared for iSCSI host attach and remote copy functionality. Therefore, appropriate settings must be applied to each IP address.
IP partnership	Two systems that are partnered to perform remote copy over native IP links.
FC partnership	Two systems that are partnered to perform remote copy over native Fibre Channel links.
Failover	Failure of a node within an I/O group causes the volume access to go through the surviving node. The IP addresses fail over to the surviving node in the I/O group. When the configuration node of the system fails, management IPs also fail over to an alternative node.
Failback	When the failed node rejoins the system, all failed over IP addresses are failed back from the surviving node to the rejoined node, and virtual disk access is restored through this node.
linkbandwidthmbits	Aggregate bandwidth of all physical links between two sites in Mbps.

IP partnership terminology	Description
IP partnership or partnership over native IP links	These terms are used to describe the IP partnership feature.
Discovery	Process by which two Lenovo Storage V-series systems exchange information about their IP address configuration. For IP-based partnerships, only IP addresses configured for Remote Copy are discovered. For example, the first Discovery takes place when the user is running the mkippartnership CLI command. Subsequent Discoveries can take place as a result of user activities (configuration changes) or as a result of hardware failures (for example, node failure, ports failure, and so on).

10.6.6 States of IP partnership

The different partnership states in IP partnership are listed in Table 10-8.

Table 10-8 States of IP partnership

State	Systems connected	Support for active remote copy I/O	Comments
Partially_Configured_Local	No	No	This state indicates that the initial discovery is complete.
Fully_Configured	Yes	Yes	Discovery successfully completed between two systems, and the two systems can establish remote copy relationships.
Fully_Configured_Stopped	Yes	Yes	The partnership is stopped on the system.
Fully_Configured_Remote_Stopped	Yes	No	The partnership is stopped on the remote system.
Not_Present	Yes	No	The two systems cannot communicate with each other. This state is also seen when data paths between the two systems are not established.
Fully_Configured_Exceeded	Yes	No	There are too many systems in the network, and the partnership from the local system to remote system is disabled.
Fully_Configured_Excluded	No	No	The connection is excluded because of too many problems, or either system cannot support the I/O work load for the Metro Mirror and Global Mirror relationships.

The following steps must be completed to establish two systems in the IP partnerships:

1. The administrator configures the CHAP secret on both the systems. This step is not mandatory, and users can choose to not configure the CHAP secret.
2. The administrator configures the system IP addresses on both local and remote systems so that they can discover each other over the network.

3. If you want to use VLANs, configure your LAN switches and Ethernet ports to use VLAN tagging (for more information about VLAN tagging, see 10.6.4, “VLAN support” on page 519).
4. The administrator configures the systems ports on each node in both of the systems by using the GUI (or the **cfgportip** CLI command), and completes the following steps:
 - a. Configure the IP addresses for remote copy data.
 - b. Add the IP addresses in the respective remote copy port group.
 - c. Define whether the host access on these ports over iSCSI is allowed.
5. The administrator establishes the partnership with the remote system from the local system where the partnership state then changes to the `Partially_Configured_Local` state.
6. The administrator establishes the partnership from the remote system with the local system, and if successful, the partnership state then changes to the `Fully_Configured` state, which implies that the partnerships over the IP network were successfully established. The partnership state momentarily remains in the `Not_Present` state before moving to the `Fully_Configured` state.
7. The administrator creates MM, GM, and GM with Change Volume relationships.

Partnership consideration: When the partnership is created, no master or auxiliary status is defined or implied. The partnership is equal. The concepts of *master or auxiliary* and *primary or secondary* apply to volume relationships only, not to system partnerships.

10.6.7 Remote copy groups

This section describes remote copy groups (or remote copy port groups) and different ways to configure the links between the two remote systems. The two Lenovo Storage V-series systems can be connected to each other over one link or, at most, two links. To address the requirement to enable the systems to know about the physical links between the two sites, the concept of remote copy port groups was introduced.

Remote copy port group ID is a numerical tag associated with an IP port of Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems to indicate which physical IP link it is connected to. Multiple nodes could be connected to the same physical long-distance link, and must therefore share the same remote copy port group id.

In scenarios where there are two physical links between the local and remote clusters, two remote copy port group IDs must be used to designate which IP addresses are connected to which physical link. This configuration must be done by the system administrator using the GUI or the **cfgportip** CLI command.

Remember: IP ports on both partners must have been configured with identical remote copy port group IDs for the partnership to be established correctly.

The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems IP addresses that are connected to the same physical link are designated with identical remote copy port groups. The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems support three remote copy groups: 0, 1, and 2.

The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems IP addresses are, by default, in remote copy port group 0. Ports in port group 0 are not considered for creating remote copy data paths between two systems. For partnerships to be established over IP

links directly, IP ports must be configured in remote copy group 1 if a single inter-site link exists, or in remote copy groups 1 and 2 if two inter-site links exist.

You can assign one IPv4 address and one IPv6 address to each Ethernet port on the system platforms. Each of these IP addresses can be shared between iSCSI host attach and the IP partnership. The user must configure the required IP address (IPv4 or IPv6) on an Ethernet port with a remote copy port group.

The administrator might want to use IPv6 addresses for remote copy operations and use IPv4 addresses on that same port for iSCSI host attach. This configuration also implies that for two systems to establish an IP partnership, both systems must have IPv6 addresses that are configured.

Administrators can choose to dedicate an Ethernet port for IP partnership only. In that case, host access must be explicitly disabled for that IP address and any other IP address that is configured on that Ethernet port.

Note: To establish an IP partnership, each Lenovo Storage V3700 V2, V3700 V2 XP and V5030 nodes must have only a single remote copy port group that is configured, 1 or 2. The remaining IP addresses must be in remote copy port group 0.

10.7 Remote Copy

This section describes the Remote Copy services, which are a synchronous remote copy called *Metro Mirror (MM)*, asynchronous remote copy called *Global Mirror (GM)*, and Global Mirror with Change Volumes. The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems provide a single point of control when remote copy is enabled in your network (regardless of the disk subsystems that are used) if those disk subsystems are supported by the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems.

The general application of remote copy services is to maintain two real-time synchronized copies of a volume. Often, two copies are geographically dispersed between two Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems, although it is possible to use MM or GM within a single system (within an I/O Group). If the master copy fails, you can enable an auxiliary copy for I/O operation.

Tips: Intracluster MM/GM uses more resources within the system when compared to an intercluster MM/GM relationship, where resource allocation is shared between the systems. Use intercluster MM/GM when possible. For mirroring volumes in the same system, it is better to use Volume Mirroring or the FlashCopy feature.

A typical application of this function is to set up a dual-site solution that uses two Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems. The first site is considered the *primary site* or *production site*, and the second site is considered the *backup site* or *failover site*, which is activated when a failure at the first site is detected.

10.7.1 Multiple Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems mirroring

Each Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems can maintain up to three partner system relationships, which enables as many as four systems to be directly

associated with each other. This system partnership capability enables the implementation of disaster recovery (DR) solutions.

Note: For more information about restrictions and limitations of native IP replication, see 10.6.3, “IP partnership limitations” on page 518.

Figure 10-74 shows an example of a multiple system mirroring configuration.

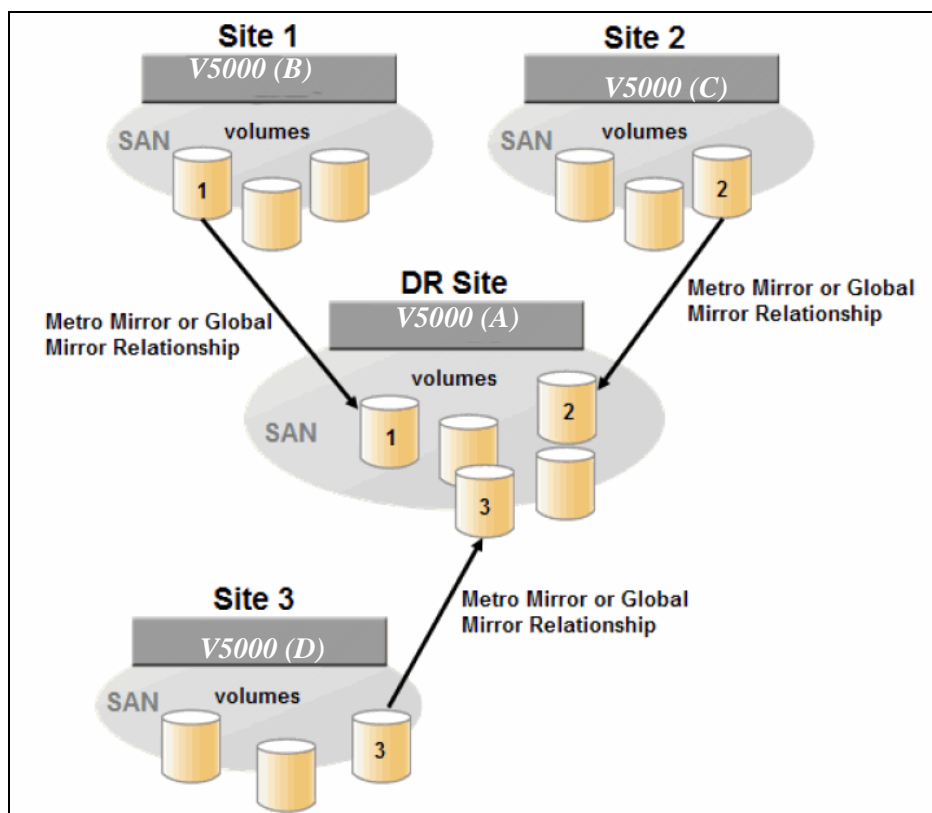


Figure 10-74 Multiple system mirroring configuration example

Supported multiple system mirroring topologies

Multiple system mirroring supports various partnership topologies, as shown in the example in Figure 10-75. This example is a star topology ($A \rightarrow B$, $A \rightarrow C$, and $A \rightarrow D$).




Figure 10-75 Star topology

Figure 10-75 on page 524 shows four systems in a star topology, with System A at the center. System A can be a central DR site for the three other locations.

By using a star topology, you can migrate applications by using a process, such as the one described in the following example:

1. Suspend application at A.
2. Remove the $A \rightarrow B$ relationship.
3. Create the $A \rightarrow C$ relationship (or the $B \rightarrow C$ relationship).
4. Synchronize to system C, and ensure that $A \rightarrow C$ is established:
 - $A \rightarrow B$, $A \rightarrow C$, $A \rightarrow D$, $B \rightarrow C$, $B \rightarrow D$, and $C \rightarrow D$
 - $A \rightarrow B$, $A \rightarrow C$, and $B \rightarrow C$

Figure 10-76 shows an example of a triangle topology ($A \rightarrow B$, $A \rightarrow C$, and $B \rightarrow C$).




Figure 10-76 Triangle topology

Figure 10-77 shows an example of Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems fully connected topology ($A \rightarrow B$, $A \rightarrow C$, $A \rightarrow D$, $B \rightarrow D$, and $C \rightarrow D$).




Figure 10-77 Fully connected topology

Figure 10-77 is a fully connected mesh in which every system has a partnership to each of the three other systems. This topology enables volumes to be replicated between any pair of systems, for example $A \rightarrow B$, $A \rightarrow C$, and $B \rightarrow C$.

Figure 10-78 shows a daisy-chain topology.




Figure 10-78 Daisy-chain topology

Although systems can have up to three partnerships, volumes can be part of only one remote copy relationship, for example A → B.

System partnership intermix: All of the preceding topologies are valid for the intermix of the IBM SAN Volume Controller with the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems if the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems are set to the replication layer and running controller firmware code 6.3.0 or later.

10.7.2 Importance of write ordering

Many applications that use block storage have a requirement to survive failures, such as loss of power or a software crash, and to not lose data that existed before the failure. Because many applications must perform large numbers of update operations in parallel, maintaining write ordering is key to ensuring the correct operation of applications after a disruption.

An application that performs a high volume of database updates is designed with the concept of dependent writes. With dependent writes, it is important to ensure that an earlier write completed before a later write is started. Reversing or performing the order of writes differently than the application intended can undermine the application's algorithms and can lead to problems, such as detected or undetected data corruption.

The Metro Mirror and Global Mirror implementation operates in a manner that is designed to always keep a consistent image at the secondary site. The Global Mirror implementation uses complex algorithms that operate to identify sets of data and number those sets of data in sequence. The data is then applied at the secondary site in the defined sequence.

Operating in this manner ensures that if the relationship is in a `Consistent_Synchronized` state, the Global Mirror target data is at least crash consistent, and supports quick recovery through application crash recovery facilities.

Remote Copy Consistency Groups

A Remote Copy Consistency Group can contain an arbitrary number of relationships up to the maximum number of MM/GM relationships that is supported by the system. MM/GM commands can be issued to a Remote Copy Consistency Group.

Therefore, these commands can be issued simultaneously for all MM/GM relationships that are defined within that Consistency Group, or to a single MM/GM relationship that is not part of a Remote Copy Consistency Group. For example, when a **starttrcconsistgrp** command is issued to the Consistency Group, all of the MM/GM relationships in the Consistency Group are started at the same time.

Figure 10-79 on page 527 shows the concept of Metro Mirror Consistency Groups. The same applies to Global Mirror Consistency Groups.

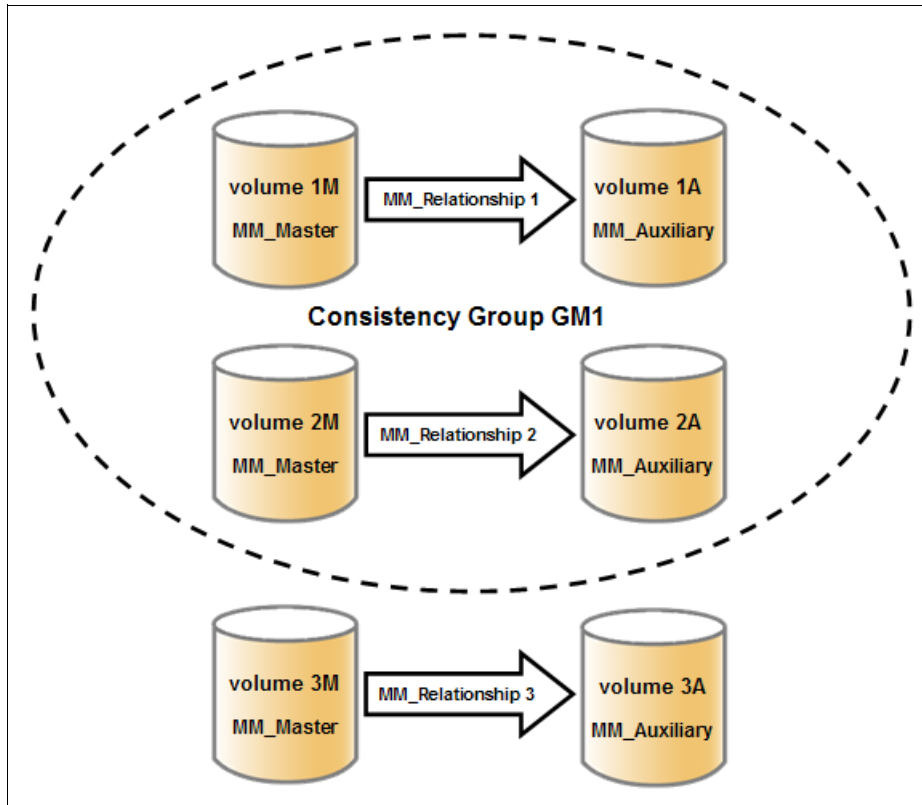


Figure 10-79 Metro Mirror Consistency Group

Because the MM_Relationship 1 and 2 are part of the Consistency Group, they can be handled as one entity. The stand-alone MM_Relationship 3 is handled separately.

Certain uses of MM/GM require the manipulation of more than one relationship. Remote Copy Consistency Groups can group relationships so that they are manipulated in unison.

Consider the following points:

- ▶ MM/GM relationships can be part of a Consistency Group, or they can be stand-alone and, therefore, are handled as single instances.
- ▶ A Consistency Group can contain zero or more relationships. An empty Consistency Group with zero relationships in it has little purpose until it is assigned its first relationship, except that it has a name.
- ▶ All relationships in a Consistency Group must have corresponding master and auxiliary volumes.
- ▶ All relationships in one Consistency Group must be the same type, for example only Metro Mirror or only Global Mirror.

Although Consistency Groups can be used to manipulate sets of relationships that do not need to satisfy these strict rules, this manipulation can lead to undesired side effects. The rules behind a Consistency Group mean that certain configuration commands are prohibited. These configuration commands are not prohibited if the relationship is not part of a Consistency Group.

For example, consider the case of two applications that are independent, yet they are placed into a single Consistency Group. If an error occurs, synchronization is lost and a background

copy process is required to recover synchronization. While this process is progressing, MM/GM rejects attempts to enable access to the auxiliary volumes of either application.

If one application finishes its background copy more quickly than the other application, MM/GM still refuses to grant access to its auxiliary volumes even though it is safe in this case. The MM/GM policy is to refuse access to the entire Consistency Group if any part of it is inconsistent.

Stand-alone relationships and Consistency Groups share a common configuration and state model. All of the relationships in a non-empty Consistency Group have the same state as the Consistency Group.

10.7.3 Remote copy intercluster communication

In the traditional Fibre Channel, the intercluster communication between systems in a Metro Mirror and Global Mirror partnership is performed over the SAN. This section describes this communication path.

For more information about intercluster communication between systems in an IP partnership, see 10.6.6, “States of IP partnership” on page 521.

Zoning

The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems FC ports on each system must communicate with each other to create the partnership. Switch zoning is critical to facilitating intercluster communication.

Intercluster communication channels

When a system partnership is defined on a pair of systems, the following intercluster communication channels are established:

- ▶ A single control channel, which is used to exchange and coordinate configuration information
- ▶ I/O channels between each of these nodes in the systems

These channels are maintained and updated as nodes and links appear and disappear from the fabric, and are repaired to maintain operation where possible. If communication between the systems is interrupted or lost, an event is logged (and the Metro Mirror and Global Mirror relationships stop).

Alerts: You can configure the system to raise Simple Network Management Protocol (SNMP) traps to the enterprise monitoring system to alert on events that indicate an interruption in internode communication occurred.

Intercluster links

All Lenovo Storage V3700 V2, V3700 V2 XP and V5030 node canisters maintain a database of other devices that are visible on the fabric. This database is updated as devices appear and disappear.

Devices that advertise themselves as Lenovo storage V-series family product nodes are categorized according to the system to which they belong. Nodes that belong to the same system establish communication channels between themselves and begin to exchange messages to implement clustering and the functional protocols of controller firmware.

Nodes that are in separate systems do not exchange messages after initial discovery is complete, unless they are configured together to perform a remote copy relationship.

The intercluster link carries control traffic to coordinate activity between two systems. The link is formed between one node in each system. The traffic between the designated nodes is distributed among logins that exist between those nodes.

If the designated node fails (or all of its logins to the remote system fail), a new node is chosen to carry control traffic. This node change causes the I/O to pause, but it does not put the relationships in a `ConsistentStopped` state.

Note: It is advised to use `chsystem` with `-partnerfcportmask` to dedicate several FC ports only to system-to-system traffic to ensure that remote copy is not affected by other traffic, such as host-to-node traffic or node-to-node traffic within the same system.

10.7.4 Metro Mirror overview

Metro Mirror establishes a synchronous relationship between two volumes of equal size. The volumes in a Metro Mirror relationship are referred to as the master (primary) volume and the auxiliary (secondary) volume. Traditional FC Metro Mirror is primarily used in a metropolitan area or geographical area, up to a maximum distance of 300 km (186.4 miles) to provide synchronous replication of data.

With synchronous copies, host applications write to the master volume, but they do not receive confirmation that the write operation completed until the data is written to the auxiliary volume. This action ensures that both the volumes have identical data when the copy completes. After the initial copy completes, the Metro Mirror function always maintains a fully synchronized copy of the source data at the target site.

Metro Mirror has the following characteristics:

- ▶ Zero recovery point objective (RPO)
- ▶ Synchronous
- ▶ Production application performance that is affected by round-trip latency

Increased distance directly affects host I/O performance because the writes are synchronous. Use the requirements for application performance when you are selecting your Metro Mirror auxiliary location.

Consistency Groups can be used to maintain data integrity for dependent writes, which is similar to FlashCopy Consistency Groups (FlashCopy Consistency Groups are described in 10.3, “Implementing FlashCopy” on page 457).

The Lenovo Storage V series system provides intracluster and intercluster Metro Mirror.

Intracluster Metro Mirror

Intracluster Metro Mirror performs the intracluster copying of a volume, in which both volumes belong to the same system and I/O Group within the system. Because it is within the same I/O Group, there must be sufficient bitmap space within the I/O Group for both sets of volumes and licensing on the system.

Important: Performing Metro Mirror across I/O Groups within a system is not supported.

Intercluster Metro Mirror

Intercluster Metro Mirror performs intercluster copying of a volume, in which one volume belongs to a system and the other volume belongs to a separate system.

Two Lenovo Storage V series systems must be defined in a partnership, which must be performed on both systems to establish a fully functional Metro Mirror partnership.

By using standard single-mode connections, the supported distance between two systems in a Metro Mirror partnership is 10 km (6.2 miles), although greater distances can be achieved by using extenders. For extended distance solutions, contact your Lenovo representative.

Limit: When a local fabric and a remote fabric are connected for Metro Mirror purposes, the inter-switch link (ISL) hop count between a local node and a remote node cannot exceed seven.

10.7.5 Synchronous remote copy

Metro Mirror is a fully synchronous remote copy technique that ensures that writes are committed at both the master and auxiliary volumes before write completion is acknowledged to the host, but only if writes to the auxiliary volumes are possible.

Events, such as a loss of connectivity between systems, can cause mirrored writes from the master volume and the auxiliary volume to fail. In that case, Metro Mirror suspends writes to the auxiliary volume and enables I/O to the master volume to continue to avoid affecting the operation of the master volumes.

Figure 10-80 shows how a write to the master volume is mirrored to the cache of the auxiliary volume before an acknowledgment of the write is sent back to the host that issued the write. This process ensures that the auxiliary is synchronized in real time if it is needed in a failover situation.

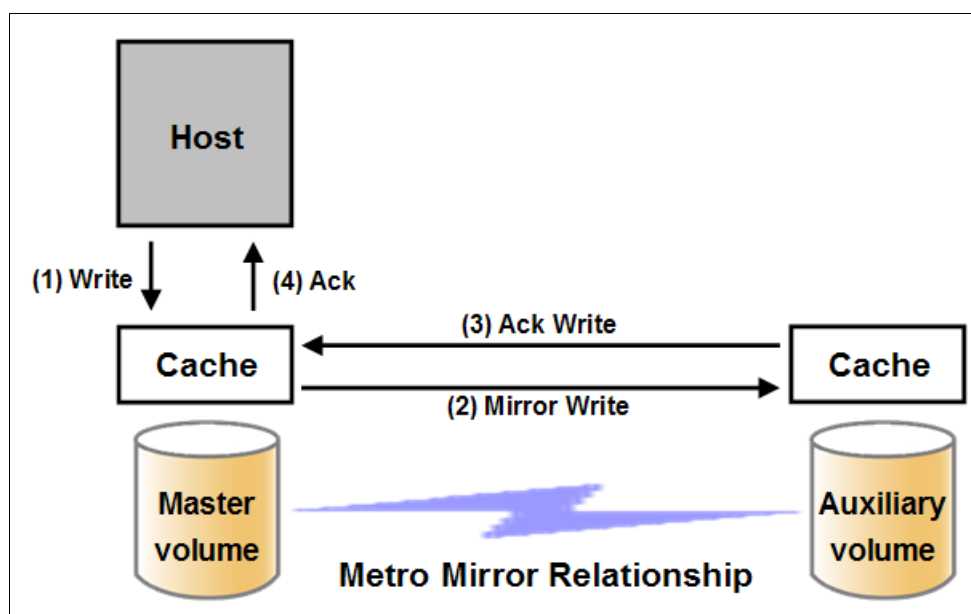


Figure 10-80 Write on volume in Metro Mirror relationship

However, this process also means that the application is exposed to the latency and bandwidth limitations (if any) of the communication link between the master and auxiliary

volumes. This process might lead to unacceptable application performance, particularly when placed under peak load. Therefore, the use of traditional Fibre Channel Metro Mirror has distance limitations that are based on your performance requirements. It does not support more than 300 km (186.4 miles).

10.7.6 Metro Mirror features

The Metro Mirror function supports the following features:

- ▶ Synchronous remote copy of volumes that are dispersed over metropolitan distances.
- ▶ The Metro Mirror relationships between volume pairs, with each volume in a pair that is managed by a Lenovo storage V-series system (requires V6.3.0 or later).
- ▶ Supports intracluster Metro Mirror where both volumes belong to the same system (and I/O Group).
- ▶ The Lenovo Storage V-series system supports intercluster Metro Mirror where each volume belongs to a separate system. You can configure a specific system for partnership with another system. All intercluster Metro Mirror processing occurs between two Lenovo storage V-series systems that are configured in a partnership.
- ▶ Intercluster and intracluster Metro Mirror can be used concurrently.
- ▶ The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems do not require that a control network or fabric is installed to manage Metro Mirror. For intercluster Metro Mirror, the system maintains a control link between two systems. This control link is used to control the state and coordinate updates at either end. The control link is implemented on top of the same FC fabric connection that the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems use for Metro Mirror I/O.
- ▶ The controller firmware implements a configuration model that maintains the Metro Mirror configuration and state through major events, such as failover, recovery, and resynchronization, to minimize user configuration action through these events.

The controller firmware supports the resynchronization of changed data so that write failures that occur on the master or auxiliary volumes do not require a complete resynchronization of the relationship.

10.7.7 Metro Mirror attributes

The Metro Mirror function possesses the following attributes:

- ▶ A partnership is created between two Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems or an IBM SAN Volume Controller system and Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems operating in the replication layer (for intercluster Metro Mirror).
- ▶ A Metro Mirror relationship is created between two volumes of the same size.
- ▶ To manage multiple Metro Mirror relationships as one entity, relationships can be made part of a Metro Mirror Consistency Group, which ensures data consistency across multiple Metro Mirror relationships and provides ease of management.
- ▶ When a Metro Mirror relationship is started and when the background copy completes, the relationship becomes consistent and synchronized.
- ▶ After the relationship is synchronized, the auxiliary volume holds a copy of the production data at the primary, which can be used for DR.
- ▶ The auxiliary volume is in read-only mode when relationship is active.

- ▶ To access the auxiliary volume, the Metro Mirror relationship must be stopped with the access option enabled, before write I/O is allowed to the auxiliary.
- ▶ The remote host server is mapped to the auxiliary volume, and the disk is available for I/O.

10.7.8 Practical use of Metro Mirror

The master volume is the production volume, and updates to this copy are mirrored in real time to the auxiliary volume. The contents of the auxiliary volume that existed when the relationship was created are deleted.

Switching copy direction: The copy direction for a Metro Mirror relationship can be switched so that the auxiliary volume becomes the master, and the master volume becomes the auxiliary, which is similar to the FlashCopy restore option. However, although the FlashCopy target volume can operate in read/write mode, the target volume of the started remote copy is always in read-only mode.

While the Metro Mirror relationship is active, the auxiliary volume is not accessible for host application write I/O at any time. The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems allow read-only access to the auxiliary volume when it contains a consistent image. IBM Storwize for Lenovo allows boot time operating system discovery to complete without an error, so that any hosts at the secondary site can be ready to start the applications with minimum delay, if required.

For example, many operating systems must read LBA zero to configure a logical unit. Although read access is allowed at the auxiliary in practice, the data on the auxiliary volumes cannot be read by a host because most operating systems write a “dirty bit” to the file system when it is mounted. Because this write operation is not allowed on the auxiliary volume, the volume cannot be mounted.

This access is provided only where consistency can be ensured. However, coherency cannot be maintained between reads that are performed at the auxiliary and later write I/Os that are performed at the master.

To enable access to the auxiliary volume for host operations, you must stop the Metro Mirror relationship by specifying the `-access` parameter. While access to the auxiliary volume for host operations is enabled, the host must be instructed to mount the volume before the application can be started, or instructed to perform a recovery process.

For example, the Metro Mirror requirement to enable the auxiliary copy for access differentiates it from third-party mirroring software on the host, which aims to emulate a single, reliable disk regardless of what system is accessing it. Metro Mirror retains the property that there are two volumes in existence but it suppresses one volume while the copy is being maintained.

The use of an auxiliary copy demands a conscious policy decision by the administrator that a failover is required, and that the tasks to be performed on the host that is involved in establishing the operation on the auxiliary copy are substantial. The goal is to make this copy rapid (much faster when compared to recovering from a backup copy) but not seamless.

The failover process can be automated through failover management software. The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems provide SNMP traps and programming (or scripting) for the CLI to enable this automation.

10.7.9 Global Mirror Overview

This section describes the Global Mirror copy service, which is an asynchronous remote copy service. This service provides and maintains a consistent mirrored copy of a source volume to a target volume.

Global Mirror establishes a Global Mirror relationship between two volumes of equal size. The volumes in a Global Mirror relationship are referred to as the *master* (source) volume and the *auxiliary* (target) volume, which is the same as Metro Mirror. Consistency Groups can be used to maintain data integrity for dependent writes, which is similar to FlashCopy Consistency Groups.

Global Mirror writes data to the auxiliary volume asynchronously, which means that host writes to the master volume provide the host with confirmation that the write is complete before the I/O completes on the auxiliary volume.

Global Mirror has the following characteristics:

- ▶ Near-zero RPO
- ▶ Asynchronous
- ▶ Production application performance that is affected by I/O sequencing preparation time

Intracuster Global Mirror

Although Global Mirror is available for intracuster, it has no functional value for production use. Intracuster Metro Mirror provides the same capability with less processor use. However, leaving this functionality in place simplifies testing and supports client experimentation and testing (for example, to validate server failover on a single test system). As with Intracuster Metro Mirror, you must consider the increase in the license requirement, because source and target exist on the same system.

Intercluster Global Mirror

Intercluster Global Mirror operations require a pair of Lenovo storage V-series systems that are connected by several intercluster links. The two systems must be defined in a partnership to establish a fully functional Global Mirror relationship.

Limit: When a local fabric and a remote fabric are connected for Global Mirror purposes, the ISL hop count between a local node and a remote node must not exceed seven hops.

10.7.10 Asynchronous remote copy

Global Mirror is an asynchronous remote copy technique. In asynchronous remote copy, the write operations are completed on the primary site and the write acknowledgment is sent to the host before it is received at the secondary site. An update of this write operation is sent to the secondary site at a later stage, which provides the capability to perform remote copy over distances that exceed the limitations of synchronous remote copy.

The Global Mirror function provides the same function as Metro Mirror remote copy, but over long-distance links with higher latency without requiring the hosts to wait for the full round-trip delay of the long-distance link.

Figure 10-81 on page 534 shows that a write operation to the master volume is acknowledged back to the host that is issuing the write before the write operation is mirrored to the cache for the auxiliary volume.

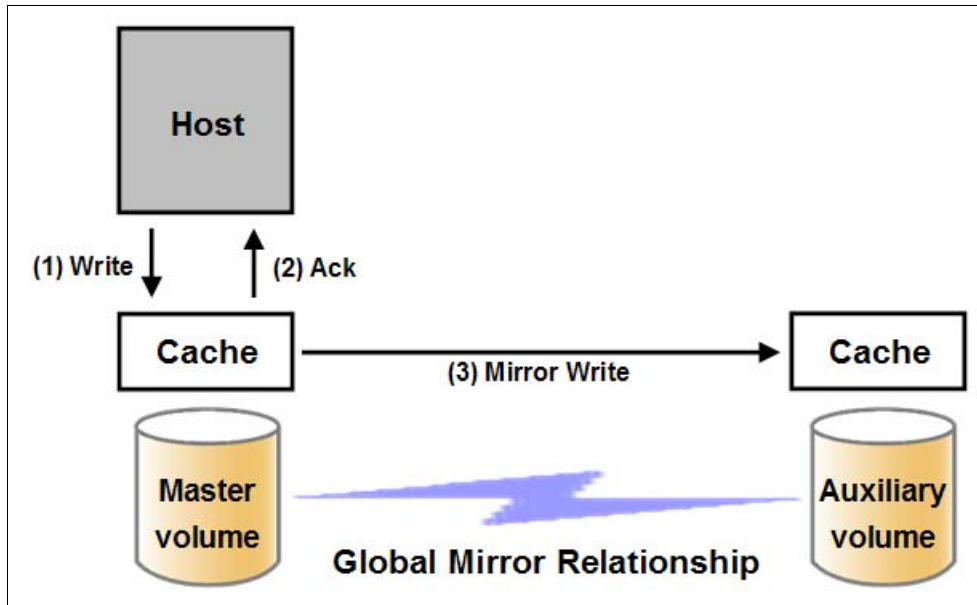


Figure 10-81 Global Mirror write sequence

The Global Mirror algorithms maintain a consistent image on the auxiliary always. They achieve this consistent image by identifying sets of I/Os that are active concurrently at the master, assigning an order to those sets, and applying those sets of I/Os in the assigned order at the secondary. As a result, Global Mirror maintains the features of Write Ordering and Read Stability.

The multiple I/Os within a single set are applied concurrently. The process that marshals the sequential sets of I/Os operates at the secondary system. Therefore, the process is not subject to the latency of the long-distance link. These two elements of the protocol ensure that the throughput of the total system can be grown by increasing system size while maintaining consistency across a growing data set.

Global Mirror write I/O from production system to a secondary system requires serialization and sequence-tagging before being sent across the network to a remote site (to maintain a write-order consistent copy of data).

To avoid affecting the production site, Lenovo storage V series system supports more parallelism in processing and managing Global Mirror writes on the secondary system by using the following methods:

- ▶ Secondary system nodes store replication writes in new redundant non-volatile cache
- ▶ Cache content details are shared between nodes
- ▶ Cache content details are batched together to make node-to-node latency less of an issue
- ▶ Nodes intelligently apply these batches in parallel as soon as possible
- ▶ Nodes internally manage and optimize Global Mirror secondary write I/O processing

In a failover scenario where the secondary site must become the master source of data, certain updates might be missing at the secondary site. Therefore, any applications that use this data must have an external mechanism for recovering the missing updates and reapplying them; for example, a transaction log replay.

Global Mirror is supported over FC, FC over IP (FCIP), FC over Ethernet (FCoE), and native IP connections. The maximum supported round-trip latency between sites depends on the type of partnership between systems, the version of software, and the system hardware that is used.

Figure 10-82 lists the maximum round-trip latency. This restriction applies to all variant of remote mirroring. More configuration requirements and guidelines apply to systems that perform remote mirroring over extended distances, where the round-trip time is greater than 80 ms.

Software version	System node hardware	Partnership		
		FC	1 Gbps IP	10 Gbps IP
7.3.0 and earlier	All	80 ms	80 ms	10 ms
7.4.0 and later	◦ Storwize® V5000 Gen2	250 ms		
	All other models	80 ms		

Figure 10-82 Supported Remote mirroring latency

10.7.11 Global Mirror features

Global Mirror supports the following features:

- ▶ Asynchronous remote copy of volumes that are dispersed over metropolitan-scale distances.
- ▶ The Global Mirror relationship between a volume pair, with each volume in the pair being managed by Lenovo storage V-series system is implemented.
- ▶ The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 support intracluster Global Mirror where both volumes belong to the same system (and I/O Group).
- ▶ The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 intercluster Global Mirror in which each volume belongs to its separate Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems. A Lenovo Storage V3700 V2, V3700 V2 XP and V5030 system can be configured for partnership with 1 - 3 other systems. For more information about IP partnership restrictions, see 10.6.3, “IP partnership limitations” on page 518.
- ▶ Intercluster and intracluster Global Mirror can be used concurrently, but not for the same volume.
- ▶ The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems do not require a control network or fabric to be installed to manage Global Mirror. For intercluster Global Mirror, the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems maintain a control link between the two systems. This control link is used to control the state and to coordinate the updates at either end. The control link is implemented on top of the same FC fabric connection that the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems use for Global Mirror I/O.
- ▶ The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems implement a configuration model that maintains the Global Mirror configuration and state through major events, such as failover, recovery, and resynchronization, to minimize user configuration action through these events.
- ▶ The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems implement flexible resynchronization support, enabling it to resynchronize volume pairs that experienced write I/Os to both disks, and to resynchronize only those regions that changed.
- ▶ An optional feature for Global Mirror is a delay simulation to be applied on writes that are sent to auxiliary volumes. It is useful in intracluster scenarios for testing purposes.

Colliding writes

Before V4.3.1, the Global Mirror algorithm required that only a single write is active on any 512-byte logical block address (LBA) of a volume. If a further write is received from a host while the auxiliary write is still active (even though the master write might complete), the new

host write is delayed until the auxiliary write is complete. This restriction is needed if a series of writes to the auxiliary must be tried again (which is called *reconstruction*). Conceptually, the data for reconstruction comes from the master volume.

If multiple writes are allowed to be applied to the master for a sector, only the most recent write gets the correct data during reconstruction. If reconstruction is interrupted for any reason, the intermediate state of the auxiliary is inconsistent. Applications that deliver such write activity do not achieve the performance that Global Mirror is intended to support. A volume statistic is maintained about the frequency of these collisions.

An attempt is made to allow multiple writes to a single location to be outstanding in the Global Mirror algorithm. There is still a need for master writes to be serialized, and the intermediate states of the master data must be kept in a non-volatile journal while the writes are outstanding to maintain the correct write ordering during reconstruction. Reconstruction must never overwrite data on the auxiliary with an earlier version. The volume statistic that is monitoring colliding writes is now limited to those writes that are not affected by this change.

Figure 10-83 shows a colliding write sequence example.

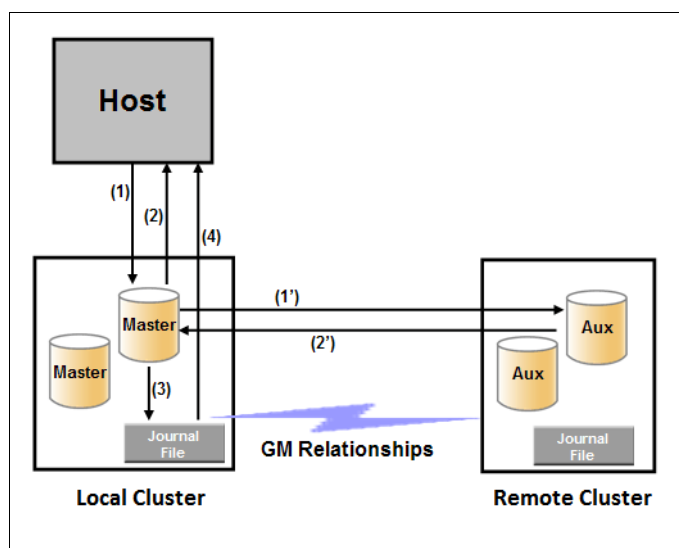


Figure 10-83 Colliding writes example

The following numbers correspond to the numbers that are shown in Figure 10-83:

- (1) The first write is performed from the host to LBA X.
- (2) The host is provided acknowledgment that the write completed even though the mirrored write to the auxiliary volume is not yet complete.
- (1') and (2') occur asynchronously with the first write.
- (3) The second write is performed from the host also to LBA X. If this write occurs before (2'), the write is written to the journal file.
- (4) The host is provided acknowledgment that the second write is complete.

Delay simulation

An optional feature for Global Mirror enables a delay simulation to be applied on writes that are sent to auxiliary volumes. This feature enables you to perform testing that detects colliding writes. Therefore, you can use this feature to test an application before the full deployment of the feature. The feature can be enabled separately for each of the intracluster or intercluster Global Mirrors.

You specify the delay setting by using the **chsystem** command and view the delay by using the **lssystem** command. The `gm_intra_cluster_delay_simulation` field expresses the amount of time that intracluster auxiliary I/Os are delayed. The `gm_inter_cluster_delay_simulation` field expresses the amount of time that intercluster auxiliary I/Os are delayed. A value of zero disables the feature.

Tip: If you are experiencing repeated problems with the delay on your link, make sure that the delay simulator was properly disabled.

10.7.12 Using Change Volumes with Global Mirror

Global Mirror is designed to achieve an RPO as low as possible so that data is as up-to-date as possible. This design places several strict requirements on your infrastructure. In certain situations with low network link quality, congested hosts, or overloaded hosts, you might be affected by multiple 1920 congestion errors.

Congestion errors happen in the following primary situations:

- ▶ Congestion at the source site through the host or network
- ▶ Congestion in the network link or network path
- ▶ Congestion at the target site through the host or network

Global Mirror has functionality that is designed to address the following conditions, which might negatively affect certain Global Mirror implementations:

- ▶ The estimation of the bandwidth requirements tends to be complex.
- ▶ Ensuring the latency and bandwidth requirements can be met is often difficult.
- ▶ Congested hosts on the source or target site can cause disruption.
- ▶ Congested network links can cause disruption with only intermittent peaks.

To address these issues, *Change Volumes* were added as an option for Global Mirror relationships. Change Volumes use the FlashCopy functionality, but they cannot be manipulated as FlashCopy volumes because they are for a special purpose only. Depending on the cycling mode defined, Change Volumes replicate point-in-time images on a cycling period.

Note: The cycling mode can be either **none** or **multi**. When cycling mode is set to **none**, the Global Mirror will behave identically to Global Mirror without Change Volumes. When cycling mode is set to **multi**, the Global Mirror will behave as described in this section.

The cycling mode can be changed only when the relationship is stopped and in `consistent_stopped` or `inconsistent_stopped` status.

The default cycling period is 300 seconds.

Your change rate needs to include only the condition of the data at the point-in-time that the image was taken, rather than all the updates during the period. The use of this function can provide significant reductions in replication volume.

Global Mirror with Change Volumes has the following characteristics:

- ▶ Larger RPO
- ▶ Point-in-time copies
- ▶ Asynchronous
- ▶ Possible system performance resource requirements because point-in-time copies are created locally

Figure 10-84 shows a simple Global Mirror relationship without Change Volumes.




Figure 10-84 Global Mirror without Change Volumes

With Change Volumes, this environment looks as it is shown in Figure 10-85.




Figure 10-85 Global Mirror with Change Volumes

With Change Volumes, a FlashCopy mapping exists between the primary volume and the primary Change Volume. The mapping is updated on the cycling period (60 seconds to one day). The primary Change Volume is then replicated to the secondary Global Mirror volume at the target site, which is then captured in another Change Volume on the target site. This approach provides an always consistent image at the target site and protects your data from being inconsistent during resynchronization.

Figure 10-86 shows how Change Volumes might save you replication traffic.

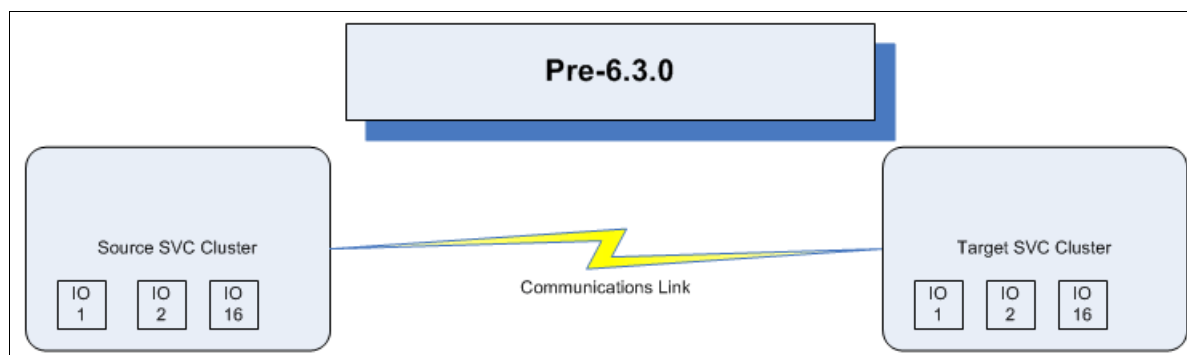


Figure 10-86 Global Mirror I/O replication without Change Volumes

In Figure 10-86 on page 538, you can see several I/Os on the source and the same number on the target, and in the same order. Assuming that this data is the same set of data being updated repeatedly, this approach results in wasted network traffic. The I/O can be completed much more efficiently, as shown in Figure 10-87.

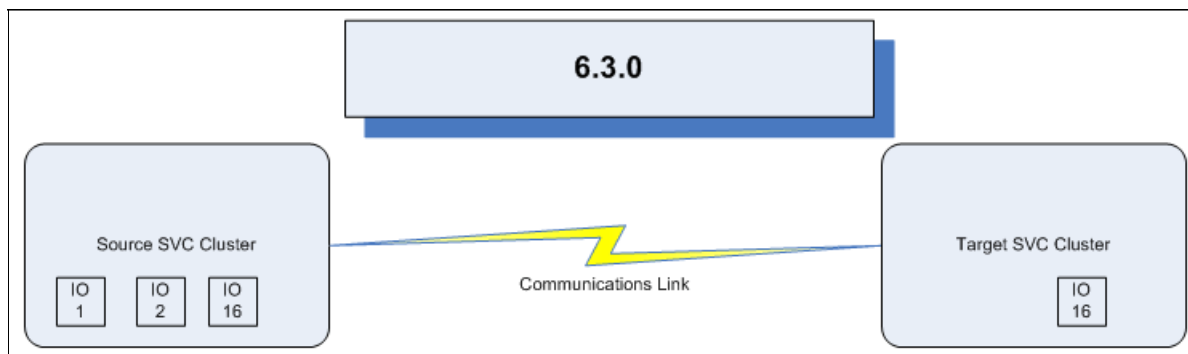


Figure 10-87 Global Mirror I/O with Change Volumes V6.3.0 and beyond

In Figure 10-87, the same data is being updated repeatedly. Therefore, Change Volumes demonstrate significant I/O transmission savings by needing to send I/O number 16 only, which was the last I/O before the cycling period.

You can adjust the cycling period by using the **chrcrelationship -cycleperiodseconds <60 - 86400>** command from the CLI. If a copy does not complete in the cycle period, the next cycle does not start until the prior cycle completes. For this reason, the use of Change Volumes gives you the following possibilities for RPO:

- ▶ If your replication completes in the cycling period, your RPO is twice the cycling period.
- ▶ If your replication does not complete within the cycling period, RPO is twice the completion time. The next cycling period starts immediately after the prior cycling period is finished.

Carefully consider your business requirements versus the performance of Global Mirror with Change Volumes. Global Mirror with Change Volumes increases the intercluster traffic for more frequent cycling periods. Therefore, selecting the shortest cycle periods possible is not always the answer. In most cases, the default must meet requirements and perform well.

Important: When you create your Global Mirror volumes with Change Volumes, make sure that you remember to select the Change Volume on the auxiliary (target) site. Failure to do so leaves you exposed during a resynchronization operation.

10.7.13 Distribution of work among nodes

For the best performance, MM/GM volumes must have their preferred nodes evenly distributed among the nodes of the systems. Each volume within an I/O Group has a preferred node property that can be used to balance the I/O load between nodes in that group. MM/GM also uses this property to route I/O between systems.

If this preferred practice is not maintained, for example, source volumes are assigned to only one node in the I/O group, you can change the preferred node for each volume to distribute volumes evenly between the nodes. You can also change the preferred node for volumes that are in a remote copy relationship without affecting the host I/O to a particular volume.

The remote copy relationship type does not matter. (The remote copy relationship type can be MM, GM, or GM with Change Volumes.) You can change the preferred node both to the source and target volumes that are participating in the remote copy relationship.

10.7.14 Background copy performance

The background copy performance is subject to sufficient Redundant Array of Independent Disks (RAID) controller bandwidth. Performance is also subject to other potential bottlenecks, such as the intercluster fabric, and possible contention from host I/O for the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems bandwidth resources.

Background copy I/O is scheduled to avoid bursts of activity that might have an adverse effect on system behavior. An entire grain of tracks on one volume is processed at around the same time but not as a single I/O. Double buffering is used to try to use sequential performance within a grain. However, the next grain within the volume might not be scheduled for some time. Multiple grains might be copied simultaneously, and might be enough to satisfy the requested rate, unless the available resources cannot sustain the requested rate.

Global Mirror paces the rate at which background copy is performed by the appropriate relationships. Background copy occurs on relationships that are in the `InconsistentCopying` state with a status of `Online`.

The quota of background copy (configured on the intercluster link) is divided evenly between all nodes that are performing background copy for one of the eligible relationships. This allocation is made irrespective of the number of disks for which the node is responsible. Each node in turn divides its allocation evenly between the multiple relationships that are performing a background copy.

The default value of the background copy is 25 megabytes per second (MBps), per volume.

Important: The background copy value is a system-wide parameter that can be changed dynamically but only on a per-system basis and not on a per-relationship basis. Therefore, the copy rate of all relationships changes when this value is increased or decreased. In systems with many remote copy relationships, increasing this value might affect overall system or intercluster link performance. The background copy rate can be changed from 1 - 1000 MBps.

10.7.15 Thin-provisioned background copy

Metro Mirror and Global Mirror relationships preserve the space-efficiency of the master. Conceptually, the background copy process detects a deallocated region of the master and sends a special *zero buffer* to the auxiliary.

If the auxiliary volume is thin-provisioned and the region is deallocated, the special buffer prevents a write and, therefore, an allocation. If the auxiliary volume is not thin-provisioned or the region in question is an allocated region of a thin-provisioned volume, a buffer of “real” zeros is synthesized on the auxiliary and written as normal.

10.7.16 Methods of synchronization

This section describes two methods that can be used to establish a synchronized relationship.

Full synchronization after creation

The full synchronization after creation method is the default method. It is the simplest method in that it requires no administrative activity apart from issuing the necessary commands. However, in certain environments, the available bandwidth can make this method unsuitable.

Use the following command sequence for a single relationship:

- ▶ Run **mkrcrelationship** without specifying the **-sync** option.
- ▶ Run **starttrcrelationship** without specifying the **-clean** option.

Synchronized before creation

In this method, the administrator must ensure that the master and auxiliary volumes contain identical data before creating the relationship by using the following technique:

- ▶ Both disks are created with the security delete feature to make all data zero.
- ▶ A complete tape image (or other method of moving data) is copied from one disk to the other disk.

With this technique, do not allow I/O on the master or auxiliary before the relationship is established. Then, the administrator must run the following commands:

- ▶ Run **mkrcrelationship** with the **-sync** flag.
- ▶ Run **starttrcrelationship** without the **-clean** flag.

Important: Failure to perform these steps correctly can cause MM/GM to report the relationship as consistent when it is not, therefore creating a data loss or data integrity exposure for hosts accessing data on the auxiliary volume.

10.7.17 Practical use of Global Mirror

The practical use of Global Mirror is similar to the Metro Mirror described in 10.7.8, “Practical use of Metro Mirror” on page 532. The main difference between the two remote copy modes is that Global Mirror and Global Mirror with Change Volumes are mostly used on much larger distances than Metro Mirror. Weak link quality or insufficient bandwidth between the primary and secondary sites can also be a reason to prefer asynchronous Global Mirror over synchronous Metro Mirror. Otherwise, the use cases for Metro Mirror and Global Mirror are the same.

10.7.18 Valid combinations of FlashCopy, Metro Mirror, and Global Mirror

Table 10-9 lists the combinations of FlashCopy and Metro Mirror or Global Mirror functions that are valid for a single volume.

Table 10-9 Valid combination for a single volume

FlashCopy	Metro Mirror or Global Mirror source	Metro Mirror or Global Mirror target
FlashCopy Source	Supported	Supported
FlashCopy Target	Supported	Not supported

10.7.19 Remote Copy configuration limits

Table 10-10 on page 542 lists the Metro Mirror and Global Mirror configuration limits.

Table 10-10 Metro Mirror configuration limits

Parameter	Value
Number of Metro Mirror or Global Mirror Consistency Groups per system	256
Number of Metro Mirror or Global Mirror relationships per system	4096
Number of Metro Mirror or Global Mirror relationships per Consistency Group	No limit is imposed beyond the Remote Copy relationships per system limit
Total volume size per I/O Group	There is a per-I/O Group limit of 1024 terabytes (TB) on the quantity of master and auxiliary volume address spaces that can participate in Metro Mirror and Global Mirror relationships. This maximum configuration uses all 512 MiB of bitmap space for the I/O Group and allows 10 MiB of space for all remaining copy services features.
Total number of Global Mirror with Change Volumes relationships per system	256

For further details on the configuration limits, please see Information Center Lenovo Storage > Lenovo Storage V3700 V2/V5030 Series > Version 8.1.0 > Configuring > Configuring > Known issues and limitations with Virtual Volumes at:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_vmwareknownvvoliss.html

10.7.20 Remote Copy states and events

This section describes the various states of a MM/GM relationship and the conditions that cause them to change. In Figure 10-88 on page 543, the MM/GM relationship diagram shows an overview of the status that can apply to a MM/GM relationship in a connected state.

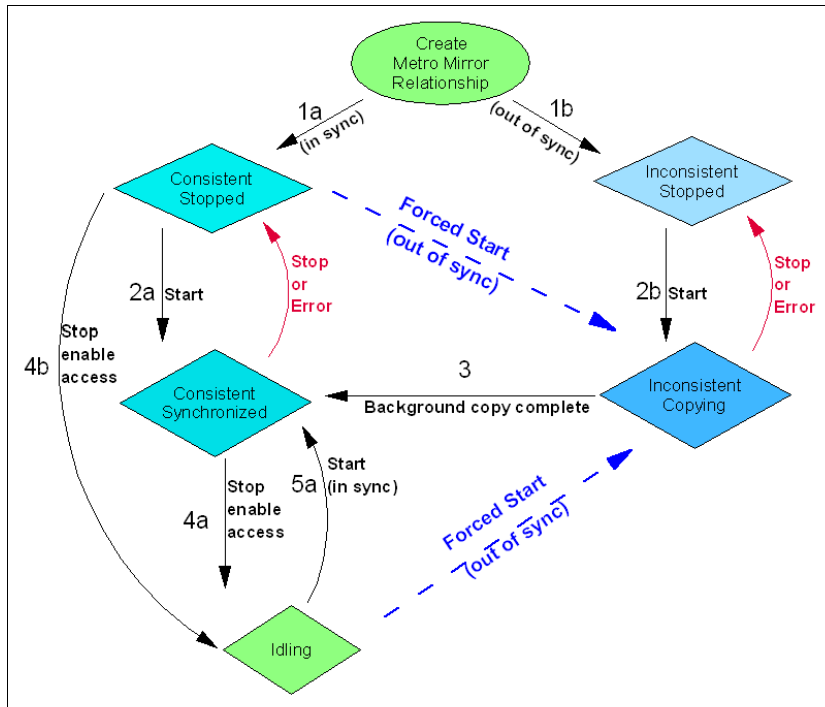


Figure 10-88 Metro Mirror or Global Mirror mapping state diagram

When the MM/GM relationship is created, you can specify whether the auxiliary volume is already in sync with the master volume, and the background copy process is then skipped. This capability is useful when MM/GM relationships are established for volumes that were created with the format option.

The following step identifiers are shown in Figure 10-88:

- Step 1:
 - a. The MM/GM relationship is created with the **-sync** option, and the MM/GM relationship enters the ConsistentStopped state.
 - b. The MM/GM relationship is created without specifying that the master and auxiliary volumes are in sync, and the MM/GM relationship enters the InconsistentStopped state.
- Step 2:
 - a. When a MM/GM relationship is started in the ConsistentStopped state, the MM/GM relationship enters the ConsistentSynchronized state. Therefore, no updates (write I/O) were performed on the master volume while in the ConsistentStopped state. Otherwise, the **-force** option must be specified, and the MM/GM relationship then enters the InconsistentCopying state while the background copy is started.
 - b. When a MM/GM relationship is started in the InconsistentStopped state, the MM/GM relationship enters the InconsistentCopying state while the background copy is started.
- Step 3:

When the background copy completes, the MM/GM relationship transitions from the InconsistentCopying state to the ConsistentSynchronized state.

- Step 4:
 - a. When a MM/GM relationship is stopped in the `ConsistentSynchronized` state, the MM/GM relationship enters the `Idling` state when you specify the `-access` option, which enables write I/O on the auxiliary volume.
 - b. When a MM/GM relationship is stopped in the `ConsistentSynchronized` state without an `-access` parameter, the auxiliary volumes remain read-only and the state of the relationship changes to `ConsistentStopped`.
 - c. To enable write I/O on the auxiliary volume, when the MM/GM relationship is in the `ConsistentStopped` state, issue the `svctask stopprcrelationship` command, which specifies the `-access` option, and the MM/GM relationship enters the `Idling` state.
- Step 5:
 - a. When a MM/GM relationship is started from the `Idling` state, you must specify the `-primary` argument to set the copy direction. If no write I/O was performed (to the master or auxiliary volume) while in the `Idling` state, the MM/GM relationship enters the `ConsistentSynchronized` state.
 - b. If write I/O was performed to the master or auxiliary volume, the `-force` option must be specified and the MM/GM relationship then enters the `InconsistentCopying` state while the background copy is started. The background process copies only the data that changed on the primary volume while the relationship was stopped.

Stop on Error

When a MM/GM relationship is stopped (intentionally, or because of an error), the state changes. For example, the MM/GM relationships in the `ConsistentSynchronized` state enter the `ConsistentStopped` state, and the MM/GM relationships in the `InconsistentCopying` state enter the `InconsistentStopped` state.

If the connection is broken between the two systems that are in a partnership, all (intercluster) MM/GM relationships enter a `Disconnected` state. For more information, see “Connected versus disconnected” on page 544.

Common states: Stand-alone relationships and Consistency Groups share a common configuration and state model. All MM/GM relationships in a Consistency Group have the same state as the Consistency Group.

State overview

In the following sections, we provide an overview of the various MM/GM states.

Connected versus disconnected

Under certain error scenarios (for example, a power failure at one site that causes one complete system to disappear), communications between two systems in an MM/GM relationship can be lost. Alternatively, the fabric connection between the two systems might fail, which leaves the two systems running but they cannot communicate with each other.

When the two systems can communicate, the systems and the relationships that spans them are described as *connected*. When they cannot communicate, the systems and the relationships spanning them are described as *disconnected*.

In this state, both systems are left with fragmented relationships and are limited regarding the configuration commands that can be performed. The disconnected relationships are portrayed as having a changed state. The new states describe what is known about the relationship and the configuration commands that are permitted.

When the systems can communicate again, the relationships are reconnected. MM/GM automatically reconciles the two state fragments, considering any configuration or other event that occurred while the relationship was disconnected. As a result, the relationship can return to the state that it was in when it became disconnected, or it can enter a new state.

Relationships that are configured between volumes in the same Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems (intracluster) are never described as being in a disconnected state.

Consistent versus inconsistent

Relationships that contain volumes that are operating as secondaries can be described as being consistent or inconsistent. Consistency Groups that contain relationships can also be described as being consistent or inconsistent. The consistent or inconsistent property describes the relationship of the data on the auxiliary to the data on the master volume. It can be considered a property of the auxiliary volume.

An auxiliary volume is described as *consistent* if it contains data that might be read by a host system from the master if power failed at an imaginary point while I/O was in progress, and power was later restored. This imaginary point is defined as the *recovery point*.

The requirements for consistency are expressed regarding activity at the master up to the recovery point. The auxiliary volume contains the data from all of the writes to the master for which the host received successful completion and that data was not overwritten by a subsequent write (before the recovery point).

Consider writes for which the host did not receive a successful completion (that is, it received bad completion or no completion at all). If the host then performed a read from the master of that data that returned successful completion and no later write was sent (before the recovery point), the auxiliary contains the same data as the data that was returned by the read from the master.

From the point of view of an application, consistency means that an auxiliary volume contains the same data as the master volume at the recovery point (the time at which the imaginary power failure occurred). If an application is designed to cope with an unexpected power failure, this assurance of consistency means that the application can use the auxiliary and begin operation as though it was restarted after the hypothetical power failure. Again, maintaining the application write ordering is the key property of consistency.

If a relationship (or set of relationships) is inconsistent and an attempt is made to start an application by using the data in the secondaries, the following outcomes are possible:

- ▶ The application might decide that the data is corrupted and crash or exit with an event code.
- ▶ The application might fail to detect that the data is corrupted and return erroneous data.
- ▶ The application might work without a problem.

Because of the risk of data corruption, and in particular undetected data corruption, MM/GM strongly enforces the concept of consistency and prohibits access to inconsistent data.

Consistency as a concept can be applied to a single relationship or a set of relationships in a Consistency Group. Write ordering is a concept that an application can maintain across

several disks that are accessed through multiple systems. Therefore, consistency must operate across all of those disks.

When you are deciding how to use Consistency Groups, the administrator must consider the scope of an application's data and consider all of the interdependent systems that communicate and exchange information.

If two programs or systems communicate and store details as a result of the information exchanged, either of the following actions might occur:

- ▶ All of the data that is accessed by the group of systems must be placed into a single Consistency Group.
- ▶ The systems must be recovered independently (each within its own Consistency Group). Then, each system must perform recovery with the other applications to become consistent with them.

Consistent versus synchronized

A copy that is consistent and up-to-date is described as *synchronized*. In a synchronized relationship, the master and auxiliary volumes differ only in regions where writes are outstanding from the host.

Consistency does not mean that the data is up-to-date. A copy can be consistent and yet contain data that was frozen at a point in the past. Write I/O might continue to a master but not be copied to the auxiliary. This state arises when it becomes impossible to keep data up-to-date and maintain consistency. An example is a loss of communication between systems when you are writing to the auxiliary.

When communication is lost for an extended period, MM/GM tracks the changes that occurred on the master, but not the order or the details of such changes (write data). When communication is restored, it is impossible to synchronize the auxiliary without sending write data to the auxiliary out of order. Therefore, consistency is lost.

The following policies can be used to cope with this situation:

- ▶ Make a point-in-time copy of the consistent auxiliary before you allow the auxiliary to become inconsistent. If there is a disaster before consistency is achieved again, the point-in-time copy target provides a consistent (although out-of-date) image.
- ▶ Accept the loss of consistency and the loss of a useful auxiliary while synchronizing the auxiliary.

Detailed states

In the following sections, we describe the states that are portrayed to the user for either Consistency Groups or relationships. We also describe information that is available in each state. The major states are designed to provide guidance about the available configuration commands.

InconsistentStopped

InconsistentStopped is a connected state. In this state, the master is accessible for read and write I/O, but the auxiliary is not accessible for read or write I/O. A copy process must be started to make the auxiliary consistent.

This state is entered when the relationship or Consistency Group was *InconsistentCopying* and suffered a persistent error or received a **stop** command that caused the copy process to stop.

A **start** command causes the relationship or Consistency Group to move to the InconsistentCopying state. A **stop** command is accepted, but has no effect.

If the relationship or Consistency Group becomes disconnected, the auxiliary side transitions to InconsistentDisconnected. The master side transitions to IdlingDisconnected.

InconsistentCopying

InconsistentCopying is a connected state. In this state, the master is accessible for read and write I/O, but the auxiliary is not accessible for read or write I/O. This state is entered after a **start** command is issued to an InconsistentStopped relationship or a Consistency Group.

It is also entered when a forced start is issued to an Idling or ConsistentStopped relationship or Consistency Group. In this state, a background copy process runs that copies data from the master to the auxiliary volume.

In the absence of errors, an InconsistentCopying relationship is active, and the copy progress increases until the copy process completes. In certain error situations, the copy progress might halt or even regress.

A persistent error or **stop** command places the relationship or Consistency Group into an InconsistentStopped state. A **start** command is accepted but has no effect.

If the background copy process completes on a stand-alone relationship or on all relationships for a Consistency Group, the relationship or Consistency Group transitions to the ConsistentSynchronized state.

If the relationship or Consistency Group becomes disconnected, the auxiliary side transitions to InconsistentDisconnected. The master side transitions to IdlingDisconnected.

ConsistentStopped

ConsistentStopped is a connected state. In this state, the auxiliary contains a consistent image, but it might be out-of-date in relation to the master. This state can arise when a relationship was in a ConsistentSynchronized state and experienced an error that forces a Consistency Freeze. It can also arise when a relationship is created with a CreateConsistentFlag set to TRUE.

Normally, write activity that follows an I/O error causes updates to the master, and the auxiliary is no longer synchronized. In this case, consistency must be given up for a period to reestablish synchronization. You must use a **start** command with the **-force** option to acknowledge this condition, and the relationship or Consistency Group transitions to InconsistentCopying. Enter this command only after all outstanding events are repaired.

In the unusual case where the master and the auxiliary are still synchronized (perhaps following a user stop, and no further write I/O was received), a **start** command takes the relationship to ConsistentSynchronized. No **-force** option is required. Also, in this case, you can enter a **switch** command that moves the relationship or Consistency Group to ConsistentSynchronized and reverses the roles of the master and the auxiliary.

If the relationship or Consistency Group becomes disconnected, the auxiliary transitions to ConsistentDisconnected. The master transitions to IdlingDisconnected.

An informational status log is generated whenever a relationship or Consistency Group enters the ConsistentStopped state with a status of OnLine. You can configure this event to generate an SNMP trap that can be used to trigger automation or manual intervention to issue a **start** command following a loss of synchronization.

ConsistentSynchronized

ConsistentSynchronized is a connected state. In this state, the master volume is accessible for read and write I/O, and the auxiliary volume is accessible for read-only I/O. Writes that are sent to the master volume are also sent to the auxiliary volume. Either successful completion must be received for both writes, the write must be failed to the host, or a state must transition out of the ConsistentSynchronized state before a write is completed to the host.

A **stop** command takes the relationship to the ConsistentStopped state. A **stop** command with the **-access** parameter takes the relationship to the Idling state.

A **switch** command leaves the relationship in the ConsistentSynchronized state, but it reverses the master and auxiliary roles (it switches the direction of replicating data). A **start** command is accepted, but has no effect.

If the relationship or Consistency Group becomes disconnected, the same transitions are made as for ConsistentStopped.

Idling

Idling is a connected state. Both master and auxiliary volumes operate in the master role. Therefore, both master and auxiliary volumes are accessible for write I/O.

In this state, the relationship or Consistency Group accepts a **start** command. MM/GM maintains a record of regions on each disk that received write I/O while they were idling. This record is used to determine what areas must be copied following a **start** command.

The **start** command must specify the new copy direction. A **start** command can cause a loss of consistency if either volume in any relationship received write I/O, which is indicated by the Synchronized status. If the **start** command leads to loss of consistency, you must specify the **-force** parameter.

Following a **start** command, the relationship or Consistency Group transitions to ConsistentSynchronized if there is no loss of consistency, or to InconsistentCopying if there is a loss of consistency.

Also, the relationship or Consistency Group accepts a **-clean** option on the **start** command while in this state. If the relationship or Consistency Group becomes disconnected, both sides change their state to IdlingDisconnected.

IdlingDisconnected

IdlingDisconnected is a disconnected state. The target volumes in this half of the relationship or Consistency Group are all in the master role and accept read or write I/O.

The priority in this state is to recover the link to restore the relationship or consistency.

No configuration activity is possible (except for deletes or stops) until the relationship becomes connected again. At that point, the relationship transitions to a connected state. The exact connected state that is entered depends on the state of the other half of the relationship or Consistency Group, which depends on the following factors:

- ▶ The state when it became disconnected
- ▶ The write activity since it was disconnected
- ▶ The configuration activity since it was disconnected

If both halves are IdlingDisconnected, the relationship becomes Idling when it is reconnected.

While `IdlingDisconnected`, if a write I/O is received that causes the loss of synchronization (synchronized attribute transitions from true to false) and the relationship was not already stopped (either through a user stop or a persistent error), an event is raised to notify you of the condition. This same event also is raised when this condition occurs for the `ConsistentSynchronized` state.

When the relationship or Consistency Group becomes connected again, the relationship becomes `InconsistentCopying` automatically unless either of the following conditions are true:

- ▶ The relationship was `InconsistentStopped` when it became disconnected.
- ▶ The user issued a **stop** command while disconnected.

In either case, the relationship or Consistency Group becomes `InconsistentStopped`.

ConsistentDisconnected

`ConsistentDisconnected` is a disconnected state. The target volumes in this half of the relationship or Consistency Group are all in the auxiliary role, and accept read I/O but *not* write I/O.

This state is entered from `ConsistentSynchronized` or `ConsistentStopped` when the auxiliary side of a relationship becomes disconnected.

In this state, the relationship or Consistency Group displays an attribute of `FreezeTime`, which is the point when Consistency was frozen. When it is entered from `ConsistentStopped`, it retains the time that it had in that state. When it is entered from `ConsistentSynchronized`, the `FreezeTime` shows the last time at which the relationship or Consistency Group was known to be consistent. This time corresponds to the time of the last successful heartbeat to the other system.

A **stop** command with the `-access` flag set to true transitions the relationship or Consistency Group to the `IdlingDisconnected` state. This state allows write I/O to be performed to the auxiliary volume and is used as part of a DR scenario.

When the relationship or Consistency Group becomes connected again, the relationship or Consistency Group becomes `ConsistentSynchronized` only if this action does not lead to a loss of consistency. The following conditions must be true:

- ▶ The relationship was `ConsistentSynchronized` when it became disconnected.
- ▶ No writes received successful completion at the master while disconnected.

Otherwise, the relationship becomes `ConsistentStopped`. The `FreezeTime` setting is retained.

Empty

This state applies only to Consistency Groups. It is the state of a Consistency Group that has no relationships and no other state information to show. It is entered when a Consistency Group is first created. It is exited when the first relationship is added to the Consistency Group, at which point the state of the relationship becomes the state of the Consistency Group.

10.8 Consistency protection for Remote and Global mirror

Prior to V7.8.1, Metro Mirror and regular Global Mirror relationships and consistency groups stop when

- ▶ the link between systems goes down; or

- a secondary volume goes offline

Consistency of the secondary volume would be lost during resynchronization, so the relationship is automatically stopped and a 1720 error is raised, which would require the user to restart the relationship manually.

Global Mirror with Change Volumes (GMCV) relationships use a secondary change volume to retain a consistent copy during resync and automatically restart when they can.

From V7.8.1, Metro Mirror and regular Global Mirror also behave like GMCV relationship if a secondary change volume is configured which will do the following:

- Makes Metro Mirror and Global Mirror more suited to links with intermittent connectivity and IP replication
- Stop as before if no secondary change volume configured

The consistency protection mechanism for metro mirror and regular global mirror uses change volumes and has following characteristics:

- It is a tweak to existing Metro Mirror and Global Mirror copy types using technology already in Global Mirror with Change Volumes
- Does not need FlashCopy license
- Uses two FlashCopy maps per relationship per system (so maximum of 2500 relationships on a 10k volume-capable system)
- Supported on all systems that can have remote mirroring license
- Requires both participating systems to be at V7.8.1 or later

Consistency protection for metro or regular Global mirror can be enabled by configuring a secondary change volume and no further configuration needed to enable this behavior. All relationships in a consistency group have to be so configured for this behavior to work on any relationship in the consistency group.

Table 10-11 describes the events and the expected behavior when consistency protection mechanism has been enabled for metro mirror and regular global mirror.

Table 10-11 Events and expected behavior

Event	Expected behavior for the relationship
link down or secondary volume offline relationship is consistent_synchronized (started and in sync)	<ul style="list-style-type: none"> ► retain the secondary consistent copy ► prepare for resynchronization ► go to the consistent_copying state ► automatically resume replication as and when possible ► go back to consistent_synchronized when complete

Event	Expected behavior for the relationship
relationship is restarted, relationship was stopped (consistent_stopped or idling) and the two copies are different	<ul style="list-style-type: none"> ▶ retain the secondary consistent copy ▶ prepare for resynchronization ▶ go to the consistent_copying state ▶ replicate differences as and when possible ▶ go back to consistent_synchronized when complete

Read/write access can be enabled as normal during resynchronization, rewinding the secondary to the last consistent image. The hosts reading the secondary volume during resynchronization will see data from the last consistent image. Also at the end of the cycle, data may need to be cleaned from the change volume to the next FlashCopy map in the cascade.

Note: Change volume should be created as thin provisioned, but in theory, can grow to 100%

A change volume must be:

- ▶ Used by the relationship that owns it.
- ▶ In the same I/O group as the associated master or auxiliary volume.
- ▶ The same size as the associated master or auxiliary volume.

A change volume is owned and used by the associated Remote Copy relationship. Therefore, it cannot be:

- ▶ Mapped to a host.
- ▶ Used as source or target of any FlashCopy maps.
- ▶ Part of any other relationship.
- ▶ A filesystem disk

Assigning a change volume to a relationship requires new FlashCopy mappings to be created between the master or auxiliary volume and the associated change volume. Therefore, there must be sufficient unallocated FlashCopy memory in the target I/O group or the command fails.

10.9 Remote Copy commands

This section presents commands that need to be issued to create and operate remote copy services.

10.9.1 Remote Copy process

The MM/GM process includes the following steps:

1. A system partnership is created between two Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems or IBM SAN Volume Controller (for intercluster MM/GM).
2. A MM/GM relationship is created between two volumes of the same size.
3. To manage multiple MM/GM relationships as one entity, the relationships can be made part of a MM/GM Consistency Group to ensure data consistency across multiple MM/GM relationships, or for ease of management.
4. The MM/GM relationship is started and when the background copy completes, the relationship is consistent and synchronized.
5. When synchronized, the auxiliary volume holds a copy of the production data at the master that can be used for disaster recovery.
6. To access the auxiliary volume, the MM/GM relationship must be stopped with the access option enabled before write I/O is submitted to the auxiliary.

Following these commands, the remote host server is mapped to the auxiliary volume and the disk is available for I/O.

The command set for MM/GM contains the following broad groups:

- ▶ Commands to create, delete, and manipulate relationships and Consistency Groups
- ▶ Commands to cause state changes

If a configuration command affects more than one system, MM/GM performs the work to coordinate configuration activity between the systems. Certain configuration commands can be performed only when the systems are connected, and fail with no effect when they are disconnected.

Other configuration commands are permitted even though the systems are disconnected. The state is reconciled automatically by MM/GM when the systems become connected again.

For any command (with one exception) a single system receives the command from the administrator. This design is significant for defining the context for a CreateRelationship **mkrcrelationship** or CreateConsistencyGroup **mkrcconsistgrp** command, in which case the system that is receiving the command is called the *local system*.

The exception is a command that sets systems into a MM/GM partnership. The **mkfcppartnership** and **mkippartnership** commands must be issued on both, the local and remote systems.

The commands in this section are described as an abstract command set, and are implemented by either of the following methods:

- ▶ CLI can be used for scripting and automation.
- ▶ GUI can be used for one-off tasks.

10.9.2 Listing available system partners

Use the **lspartnershipcandidate** command to list the systems that are available for setting up a two-system partnership. This command is a prerequisite for creating MM/GM relationships.

Note: This command is not supported on IP partnerships. Use **mkippartnership** for IP connections.

10.9.3 Changing the system parameters

When you want to change system parameters specific to any remote copy or Global Mirror only, use the **chsystem** command. The **chsystem** command features the following parameters for MM/GM:

► **-relationshipbandwidthlimit** *cluster_relationship_bandwidth_limit*

This parameter controls the maximum rate at which any one remote copy relationship can synchronize. The default value for the relationship bandwidth limit is 25 MBps, but this value can now be specified 1 - 100,000 MBps. The partnership overall limit is controlled by the **chpartnership -linkbandwidthmbits** command, and must be set on each involved system.

Important: Do not set this value higher than the default without first establishing that the higher bandwidth can be sustained without affecting the host's performance. The limit must never be higher than the maximum that is supported by the infrastructure connecting the remote sites, regardless of the compression rates that you might achieve.

► **-gmlinktolerance** *link_tolerance*

This parameter specifies the maximum period that the system tolerates delay before stopping Global Mirror relationships. Specify values 60 - 86,400 seconds in increments of 10 seconds. The default value is 300. Do not change this value except under the direction of IBM Support.

► **-gmmaxhostdelay** *max_host_delay*

This parameter specifies the maximum time delay, in milliseconds, at which the Global Mirror link tolerance timer starts counting down. This threshold value determines the additional effect that Global Mirror operations can add to the response times of the Global Mirror source volumes. You can use this parameter to increase the threshold from the default value of 5 milliseconds.

► **-gminterdelaysimulation** *link_tolerance*

This parameter specifies the number of milliseconds that I/O activity (intercluster copying to an auxiliary volume) is delayed. This parameter enables you to test performance implications before Global Mirror is deployed and a long-distance link is obtained. Specify a value of 0 - 100 milliseconds in 1-millisecond increments. The default value is 0. Use this argument to test each intercluster Global Mirror relationship separately.

► **-gmintradelaysimulation** *link_tolerance*

This parameter specifies the number of milliseconds that I/O activity (intracluster copying to an auxiliary volume) is delayed. By using this parameter, you can test performance implications before Global Mirror is deployed and a long-distance link is obtained. Specify a value of 0 - 100 milliseconds in 1-millisecond increments. The default value is 0. Use this argument to test each intracluster Global Mirror relationship separately.

► **-maxreplicationdelay** *max_replication_delay*

This parameter sets a maximum replication delay in seconds. The value must be a number 1 - 360. This feature sets the maximum number of seconds to be tolerated to complete a single I/O. If I/O can't complete within the *max_replication_delay* the 1920

event is reported. This is the system-wide setting. When set to 0, the feature is disabled. This applies to Metro Mirror and Global Mirror relationships.

Use the **chsystem** command to adjust these values, as shown in the following example:

```
chsystem -gmlinktolerance 300
```

You can view all of these parameter values by using the **lssystem <system_name>** command.

We focus on the **gmlinktolerance** parameter in particular. If poor response extends past the specified tolerance, a 1920 event is logged and one or more GM relationships automatically stop to protect the application hosts at the primary site. During normal operations, application hosts experience a minimal effect from the response times because the GM feature uses asynchronous replication.

However, if GM operations experience degraded response times from the secondary system for an extended period, I/O operations begin to queue at the primary system. This queue results in an extended response time to application hosts. In this situation, the **gmlinktolerance** feature stops GM relationships, and the application host's response time returns to normal.

After a 1920 event occurs, the GM auxiliary volumes are no longer in the `consistent_synchronized` state until you fix the cause of the event and restart your GM relationships. For this reason, ensure that you monitor the system to track when these 1920 events occur.

You can disable the **gmlinktolerance** feature by setting the **gmlinktolerance** value to 0 (zero). However, the **gmlinktolerance** feature cannot protect applications from extended response times if it is disabled. It might be appropriate to disable the **gmlinktolerance** feature under the following circumstances:

- ▶ During SAN maintenance windows in which degraded performance is expected from SAN components, and application hosts can withstand extended response times from GM volumes.
- ▶ During periods when application hosts can tolerate extended response times and it is expected that the **gmlinktolerance** feature might stop the GM relationships. For example, if you test by using an I/O generator that is configured to stress the back-end storage, the **gmlinktolerance** feature might detect the high latency and stop the GM relationships. Disabling the **gmlinktolerance** feature prevents this result at the risk of exposing the test host to extended response times.

A 1920 event indicates that one or more of the SAN components cannot provide the performance that is required by the application hosts. This situation can be temporary (for example, a result of a maintenance activity) or permanent (for example, a result of a hardware failure or an unexpected host I/O workload).

If 1920 events are occurring, it can be necessary to use a performance monitoring and analysis tool, such as the IBM Virtual Storage Center, to help identify and resolve the problem.

10.9.4 System partnership

To create a Lenovo storage V series system partnership, use the **mkfcpartnership** command for traditional Fibre Channel (FC or FCoE) connections or **mkippartnership** for IP-based connections.

The **svctask mkfcpartnership** command

Use the **mkfcpartnership** command to establish a one-way MM/GM partnership between the local system and a remote system. Alternatively, use **mkippartnership** to create IP-based partnership.

To establish a fully functional MM/GM partnership, you must issue this command on both systems. This step is a prerequisite for creating MM/GM relationships between volumes on the Lenovo Storage V series systems.

When the partnership is created, you can specify the bandwidth to be used by the background copy process between the local and remote system. If it is not specified, the bandwidth defaults to 50 MBps. The bandwidth must be set to a value that is less than or equal to the bandwidth that can be sustained by the intercluster link.

Background copy bandwidth effect on foreground I/O latency

The background copy bandwidth determines the rate at which the background copy is attempted for MM/GM. The background copy bandwidth can affect foreground I/O latency in one of the following ways:

- ▶ The following result can occur if the background copy bandwidth is set too high compared to the MM/GM intercluster link capacity:
 - The background copy I/Os can back up on the MM/GM intercluster link.
 - There is a delay in the synchronous auxiliary writes of foreground I/Os.
 - The foreground I/O latency increases as perceived by applications.
- ▶ If the background copy bandwidth is set too high for the storage at the primary site, background copy read I/Os overload the primary storage and delay foreground I/Os.
- ▶ If the background copy bandwidth is set too high for the storage at the secondary site, background copy writes at the secondary site overload the auxiliary storage and again delay the synchronous secondary writes of foreground I/Os.

To set the background copy bandwidth optimally, ensure that you consider all three resources: primary storage, intercluster link bandwidth, and auxiliary storage. Provision the most restrictive of these three resources between the background copy bandwidth and the peak foreground I/O workload.

Perform this provisioning by calculation or by determining experimentally how much background copy can be allowed before the foreground I/O latency becomes unacceptable. Then, reduce the background copy to accommodate peaks in workload.

The **chpartnership** command

To change the bandwidth that is available for background copy in the system partnership, use the **chpartnership -backgroundcopyrate <percentage_of_link_bandwidth>** command to specify the percentage of whole link capacity to be used by background copy process.

10.9.5 Creating a Metro Mirror/Global Mirror consistency group

Use the **mkrcconsistgrp** command to create an empty MM/GM Consistency Group.

The MM/GM consistency group name must be unique across all consistency groups that are known to the systems owning this consistency group. If the consistency group involves two systems, the systems must be in communication throughout the creation process.

The new consistency group does not contain any relationships and is in the Empty state. You can add MM/GM relationships to the group (upon creation or afterward) by using the **chrelationship** command.

10.9.6 Creating a Metro Mirror/Global Mirror relationship

Use the **mkrcrelationship** command to create a new MM/GM relationship. This relationship persists until it is deleted.

Optional parameter: If you do not use the **-global** optional parameter, a Metro Mirror relationship is created rather than a Global Mirror relationship.

The auxiliary volume must be equal in size to the master volume or the command fails. If both volumes are in the same system, they must be in the same I/O Group. The master and auxiliary volume cannot be in an existing relationship, and they cannot be the target of a FlashCopy mapping. This command returns the new relationship (**relationship_id**) when successful.

When the MM/GM relationship is created, you can add it to an existing Consistency Group, or it can be a stand-alone MM/GM relationship if no Consistency Group is specified.

The **lsrcrelationshipcandidate** command

Use the **lsrcrelationshipcandidate** command to list the volumes that are eligible to form an MM/GM relationship.

When the command is issued, you can specify the master volume name and auxiliary system to list the candidates that comply with the prerequisites to create a MM/GM relationship. If the command is issued with no parameters, all of the volumes that are not disallowed by another configuration state, such as being a FlashCopy target, are listed.

10.9.7 Changing Metro Mirror/Global Mirror relationship

Use the **chrcrelationship** command to modify the following properties of an MM/GM relationship:

- ▶ Change the name of an MM/GM relationship.
- ▶ Add a relationship to a group.
- ▶ Remove a relationship from a group using the **-force** flag.

Adding an MM/GM relationship: When an MM/GM relationship is added to a Consistency Group that is not empty, the relationship must have the same state and copy direction as the group to be added to it.

10.9.8 Changing Metro Mirror/Global Mirror consistency group

Use the **chrcconsistgrp** command to change the name of an MM/GM Consistency Group.

10.9.9 Starting Metro Mirror/Global Mirror relationship

Use the **starttrcrelationship** command to start the copy process of an MM/GM relationship.

When the command is issued, you can set the copy direction if it is undefined, and, optionally, you can mark the auxiliary volume of the relationship as clean. The command fails if it is used as an attempt to start a relationship that is already a part of a consistency group.

You can issue this command only to a relationship that is connected. For a relationship that is idling, this command assigns a copy direction (master and auxiliary roles) and begins the copy process. Otherwise, this command restarts a previous copy process that was stopped by a **stop** command or by an I/O error.

If the resumption of the copy process leads to a period when the relationship is inconsistent, you must specify the **-force** parameter when the relationship is restarted. This situation can arise if, for example, the relationship was stopped and then further writes were performed on the original master of the relationship.

The use of the **-force** parameter here is a reminder that the data on the auxiliary becomes inconsistent while resynchronization (background copying) takes place. Therefore, this data is unusable for DR purposes before the background copy completes.

In the `Idling` state, you must specify the master volume to indicate the copy direction. In other connected states, you can provide the **-primary** argument, but it must match the existing setting.

10.9.10 Stopping Metro Mirror/Global Mirror relationship

Use the **stopprcrelationship** command to stop the copy process for a relationship. You can also use this command to enable write access to a consistent auxiliary volume by specifying the **-access** parameter.

This command applies to a stand-alone relationship. It is rejected if it is addressed to a relationship that is part of a Consistency Group. You can issue this command to stop a relationship that is copying from master to auxiliary.

If the relationship is in an inconsistent state, any copy operation stops and does not resume until you issue a **startprcrelationship** command. Write activity is no longer copied from the master to the auxiliary volume. For a relationship in the `ConsistentSynchronized` state, this command causes a Consistency Freeze.

When a relationship is in a consistent state (that is, in the `ConsistentStopped`, `ConsistentSynchronized`, or `ConsistentDisconnected` state), you can use the **-access** parameter with the **stopprcrelationship** command to enable write access to the auxiliary volume.

10.9.11 Starting Metro Mirror/Global Mirror consistency group

Use the **startprcconsistgrp** command to start an MM/GM consistency group. You can issue this command only to a consistency group that is connected.

For a consistency group that is idling, this command assigns a copy direction (master and auxiliary roles) and begins the copy process. Otherwise, this command restarts a previous copy process that was stopped by a **stop** command or by an I/O error.

10.9.12 Stopping Metro Mirror/Global Mirror consistency group

Use the **startrcconsistgrp** command to stop the copy process for an MM/GM consistency group. You can also use this command to enable write access to the auxiliary volumes in the group if the group is in a consistent state.

If the consistency group is in an inconsistent state, any copy operation stops and does not resume until you issue the **startrcconsistgrp** command. Write activity is no longer copied from the master to the auxiliary volumes that belong to the relationships in the group. For a consistency group in the ConsistentSynchronized state, this command causes a Consistency Freeze.

When a consistency group is in a consistent state (for example, in the ConsistentStopped, ConsistentSynchronized, or ConsistentDisconnected state), you can use the **-access** parameter with the **stoprcconsistgrp** command to enable write access to the auxiliary volumes within that group.

10.9.13 Deleting Metro Mirror/Global Mirror relationship

Use the **rmrcrelationship** command to delete the relationship that is specified. Deleting a relationship deletes only the logical relationship between the two volumes. It does not affect the volumes themselves.

If the relationship is disconnected at the time that the command is issued, the relationship is deleted only on the system on which the command is being run. When the systems reconnect, the relationship is automatically deleted on the other system.

Alternatively, if the systems are disconnected and you still want to remove the relationship on both systems, you can issue the **rmrcrelationship** command independently on both of the systems.

A relationship cannot be deleted if it is part of a consistency group. You must first remove the relationship from the consistency group.

If you delete an inconsistent relationship, the auxiliary volume becomes accessible even though it is still inconsistent. This situation is the one case in which MM/GM does not inhibit access to inconsistent data.

10.9.14 Deleting Metro Mirror/Global Mirror consistency group

Use the **rmrcconsistgrp** command to delete an MM/GM consistency group. This command deletes the specified consistency group. You can issue this command for any existing consistency group.

If the consistency group is disconnected at the time that the command is issued, the consistency group is deleted only on the system on which the command is being run. When the systems reconnect, the consistency group is automatically deleted on the other system.

Alternatively, if the systems are disconnected and you still want to remove the consistency group on both systems, you can issue the **rmrcconsistgrp** command separately on both of the systems.

If the consistency group is not empty, the relationships within it are removed from the consistency group before the group is deleted. These relationships then become stand-alone

relationships. The state of these relationships is not changed by the action of removing them from the consistency group.

10.9.15 Reversing Metro Mirror/Global Mirror relationship

Use the **switchrcrelationship** command to reverse the roles of the master volume and the auxiliary volume when a stand-alone relationship is in a consistent state. When the command is issued, the wanted master must be specified.

10.9.16 Reversing Metro Mirror/Global Mirror consistency group

Use the **switchrcconsistgrp** command to reverse the roles of the master volume and the auxiliary volume when a consistency group is in a consistent state. This change is applied to all of the relationships in the consistency group. When the command is issued, the wanted master must be specified.

Important: Remember that by reversing the roles, your current source volumes become targets, and target volumes become source volumes. Therefore, you lose write access to your current primary volumes.

10.10 Managing Remote Copy using the GUI

It is often easier to control working with Metro Mirror or Global Mirror by using the GUI, if you have few mappings. When many mappings are used, run your commands by using the CLI. In this section, we describe the tasks that you can perform at a remote copy level.

The following panes are used to visualize and manage your remote copies:

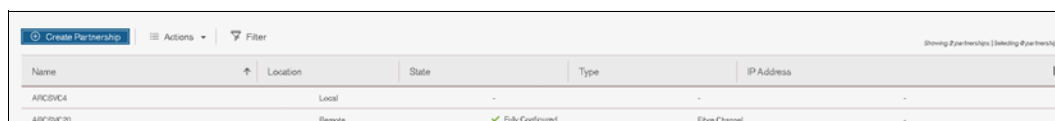
- The Remote Copy pane, as shown in Figure 10-89.
To access the Remote Copy pane, move the mouse pointer over the Copy Services selection and click **Remote Copy**.



Name	State	Master Volume	Auxiliary Volume
rcrcl1	Consistent Synchronized	ld_gm_4	ld_gm_4
rcrcl0	Consistent Synchronized	ld_gm_6	ld_gm_6

Figure 10-89 Remote Copy pane

- The Partnerships pane, as shown in Figure 10-90.
To access the Partnerships pane, move the mouse pointer over the Copy Services selection and click **Partnerships**.



Name	Location	State	Type	IP Address
APC8VCA	Local	-	-	-
APC8V20	Remote	Fully Configured	Fibre Channel	-

Figure 10-90 Partnerships pane

10.10.1 Creating Fibre Channel partnership

To create an FC partnership between the systems running, use the GUI and complete the following steps:

1. From the main navigation pane, click **Copy Services** → **Partnerships** as shown in Figure 10-91.

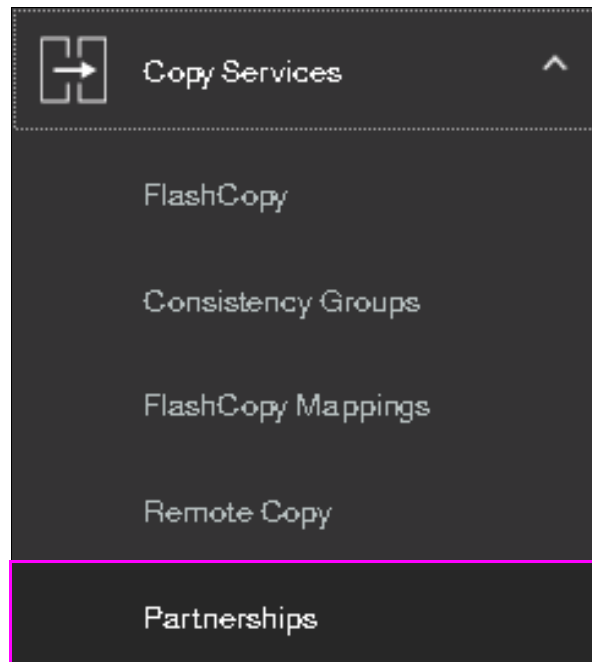


Figure 10-91 Selecting Partnerships

2. The Partnership pane opens as shown in Figure 10-92.




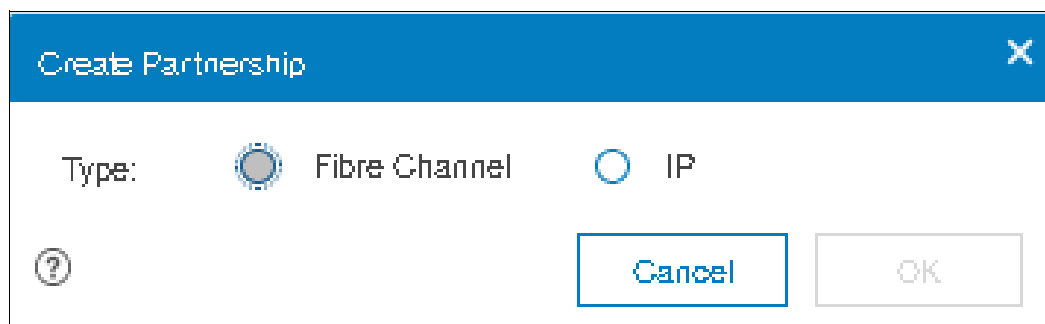
Figure 10-92 Partnership pane

3. Click **Create Partnership** to create a partnership with another Lenovo storage V series system, as shown in Figure 10-93.



Figure 10-93 Create a partnership

4. In the **Create Partnership** window, indicate the partnership type, either Fibre Channel or IP as shown in Figure 10-94.



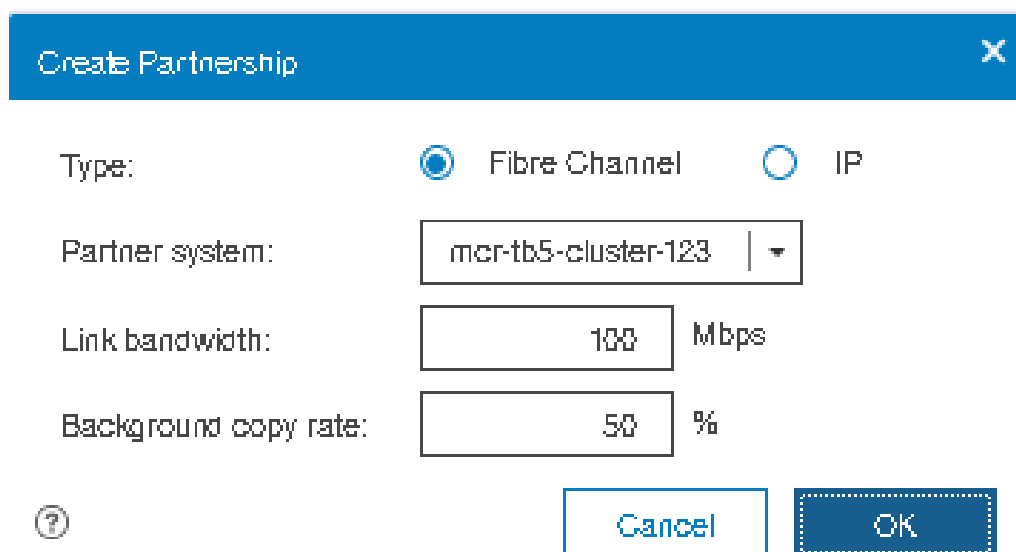
The image shows a 'Create Partnership' dialog box. At the top is a blue header bar with the title 'Create Partnership' and a close button (X). Below the header, the 'Type:' label is followed by two radio buttons. The first radio button is selected and is labeled 'Fibre Channel'. The second radio button is unselected and is labeled 'IP'. At the bottom left is a help icon (a question mark inside a circle). At the bottom right are two buttons: 'Cancel' and 'OK'.

Figure 10-94 Partnership type

5. For Fibre Channel partnership, select an available partner system from the drop-down list. If no candidate is available, the following error message is displayed:

This system does not have any candidates.

- Enter a link bandwidth in megabits per second (Mbps) that is used by the background copy process between the systems in the partnership.
- Enter the background copy rate.
- Click **OK** to confirm the partnership relationship as shown in Figure 10-95.



The image shows the 'Create Partnership' dialog box with 'Fibre Channel' selected. Below the 'Type:' section, there is a 'Partner system:' label followed by a drop-down menu showing 'mcr-tb5-cluster-123'. Below that is a 'Link bandwidth:' label followed by a text box containing '100' and the unit 'Mbps'. Below that is a 'Background copy rate:' label followed by a text box containing '50' and the unit '%'. At the bottom left is a help icon (a question mark inside a circle). At the bottom right are two buttons: 'Cancel' and 'OK'.

Figure 10-95 Create fibre channel partnership

Note: If you choose IP partnership, you must provide the IP address of the partner system and the partner system's CHAP key.

6. You will get a confirmation window as shown in Figure 10-96 on page 562.

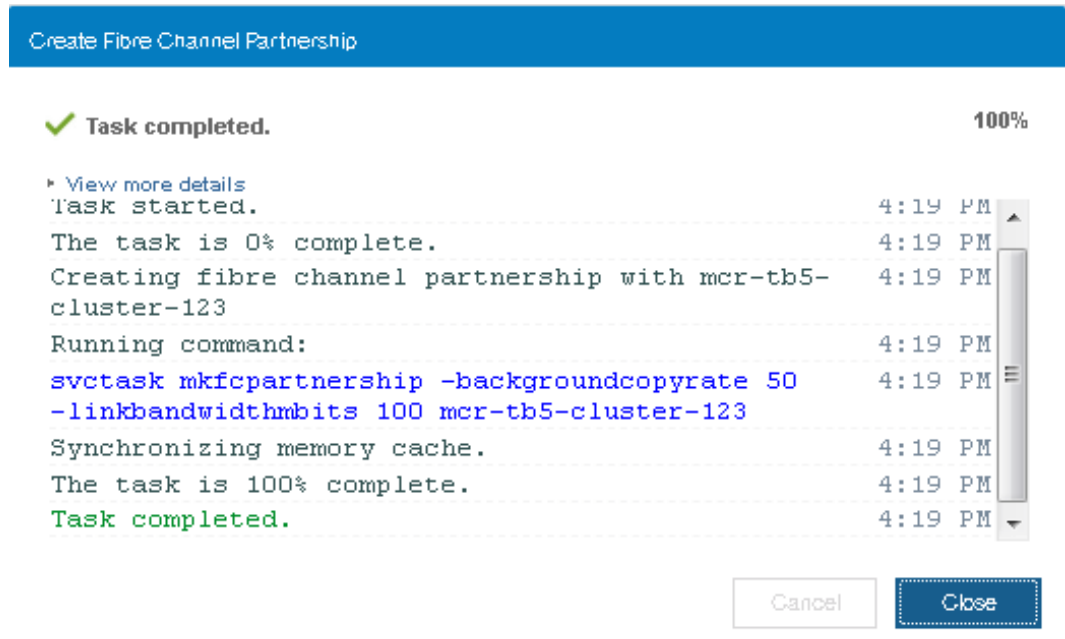


Figure 10-96 Fibre channel partnership created

7. As shown in Figure 10-97, our partnership is in the Partially Configured state because this work was performed only on one side of the partnership so far.

Create Partnership Actions Filter						
Showing 2 partnerships (selecting 0 partnerships)						
Name	Location	State	Type	IP Address		
mcr-95-cluster-29	Local					
mcr-95-cluster-123	Remote	Partially Configured	Local	Fibre Channel		

Figure 10-97 Viewing system partnerships

To fully configure the partnership between both systems, perform the same steps on the other system in the partnership. For simplicity and brevity, we show only the two most significant windows when the partnership is fully configured.

8. Starting the GUI at the partner system, select **mcr-tb5-cluster-29** for the system partnership. We specify the available bandwidth for the background copy (100 Mbps) and then click **OK**.

Now that both sides of the system partnership are defined, the resulting windows are similar at both of the systems, as shown in Figure 10-98.

Create Partnership Actions Filter						
Showing 2 partnerships (selecting 0 partnerships)						
Name	Location	State	Type	IP Address		
mcr-95-cluster-29	Local					
mcr-95-cluster-123	Remote	Fully Configured	Fibre Channel			

Figure 10-98 Fully configured remote partnership

10.10.2 Creating stand-alone remote copy relationships

In this section, we create remote copy mappings for volumes with their respective remote targets. The source and target volumes were created before this operation was done on both systems. The target volume must have the same size as the source volume.

Complete the following steps to create stand-alone copy relationships:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. Select **Not in a Group** and then click **Action** as shown in Figure 10-99.

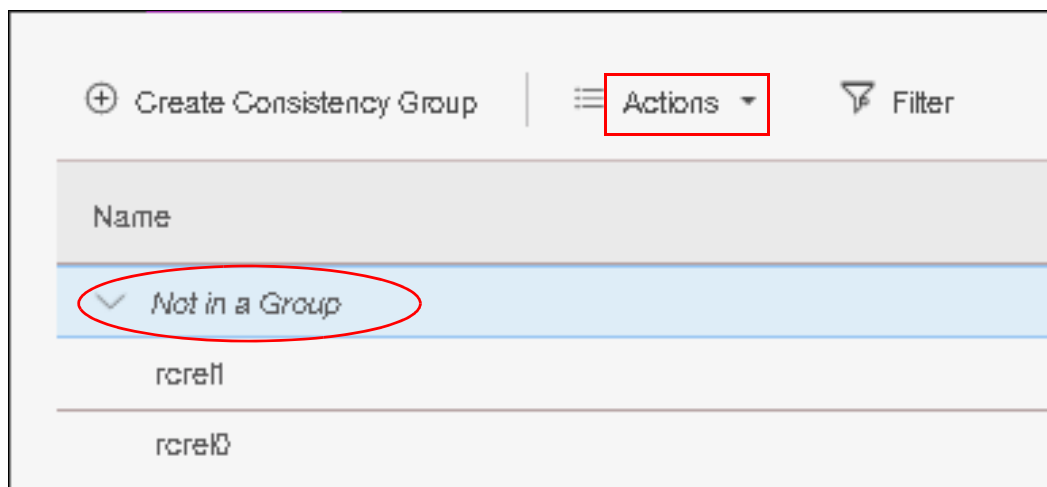


Figure 10-99 Creating a new Remote Copy relationship without consistency group

3. Click **Create Relationship**, as shown in Figure 10-100.

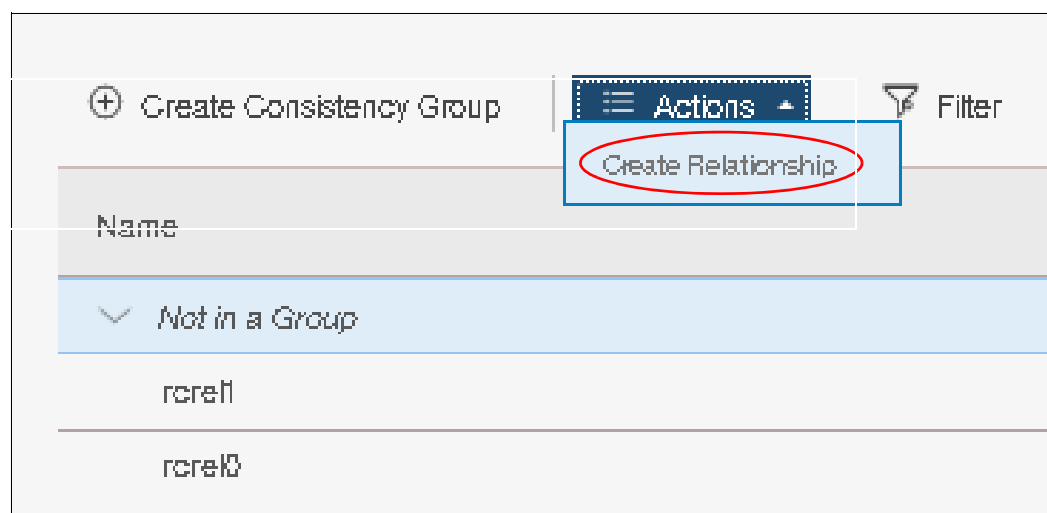


Figure 10-100 Create Relationship action

4. In the Create Relationship window, select one of the following types of relationships that you want to create (as shown in Figure 10-101 on page 564):
 - Metro Mirror
 - Global Mirror
 - Global Mirror with Change Volumes

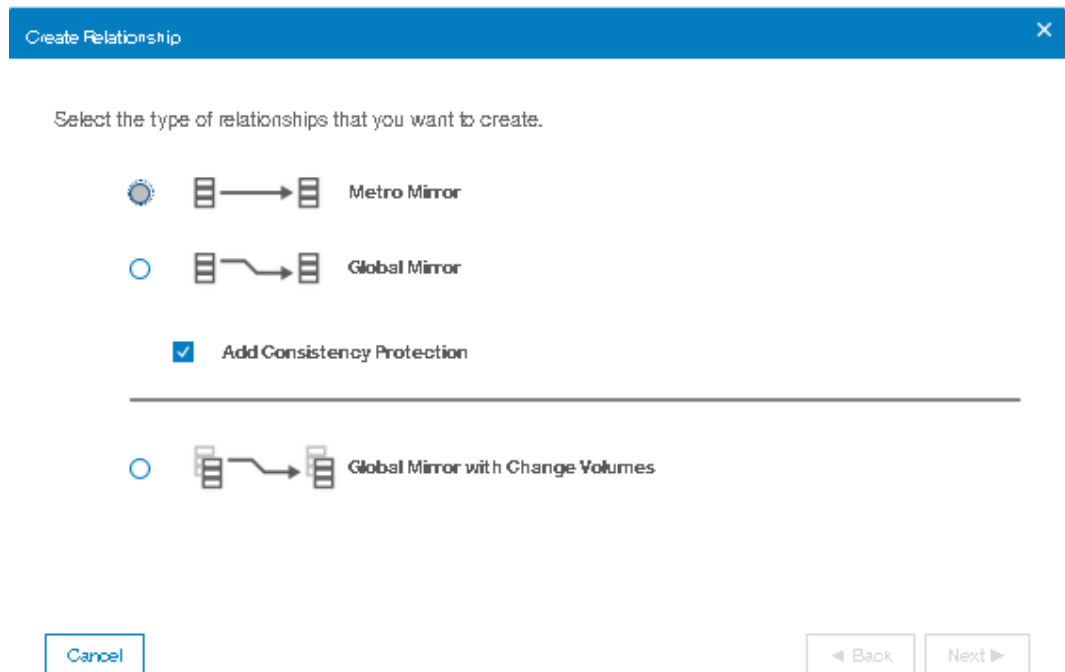


Figure 10-101 Select the type of relationship that you want to create

5. We want to create a Metro Mirror relationship. See Figure 10-102. Click **Next**.

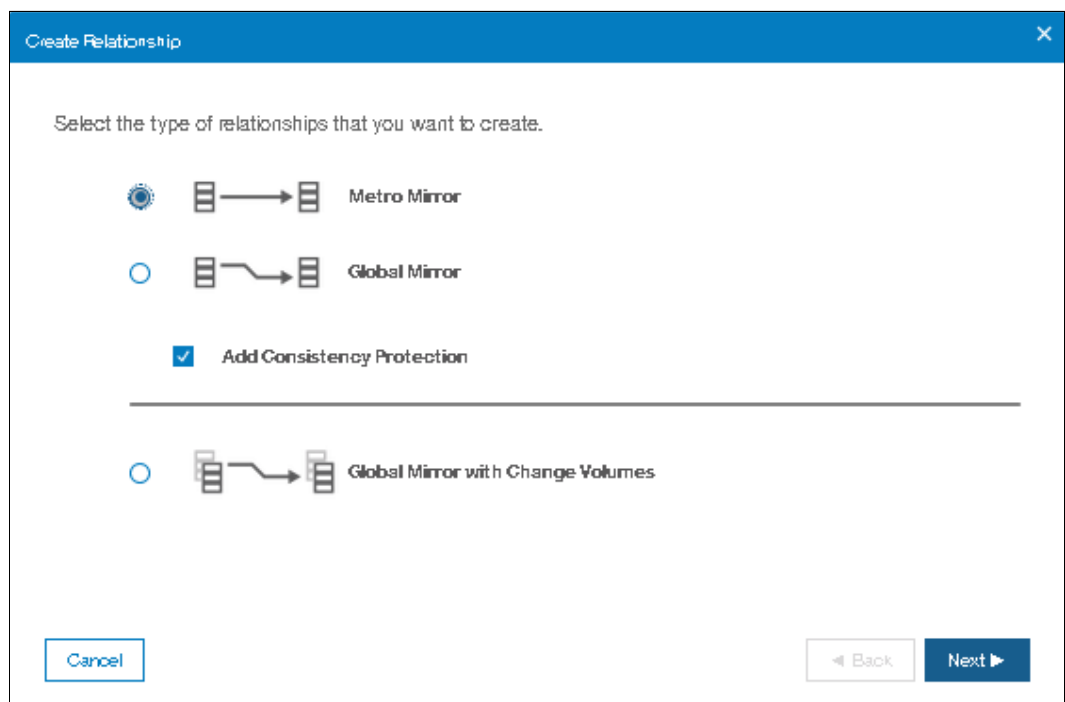
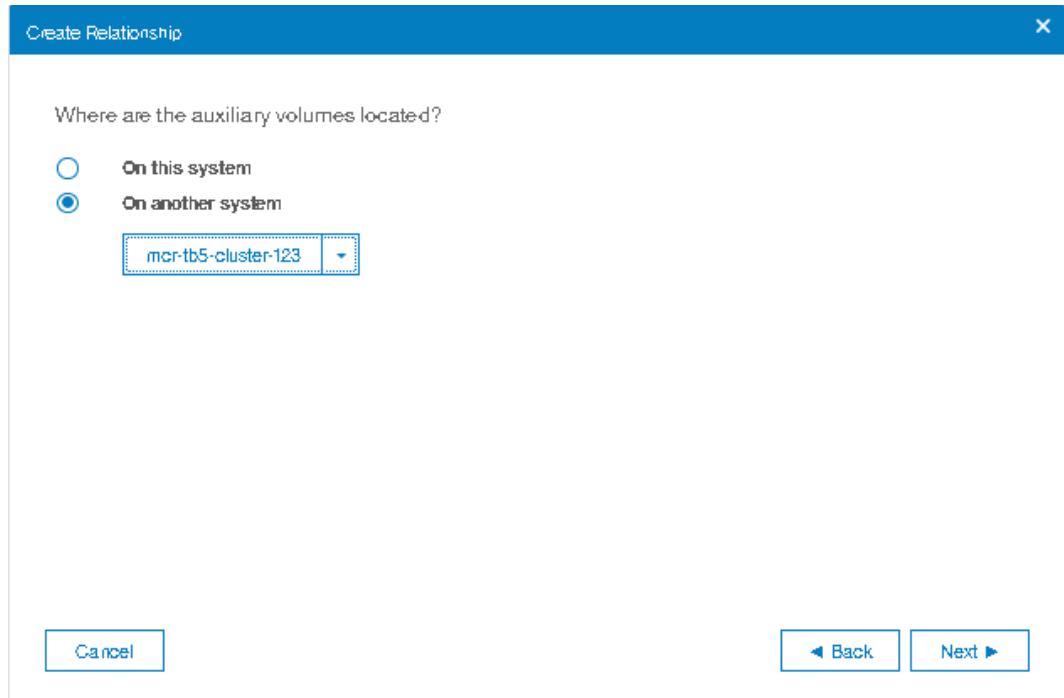


Figure 10-102 Selecting Metro Mirror as the type of relationship

Note: Starting from V7.8.1, consistency protection via Change Volume has been enabled by default. Refer to 10.8, “Consistency protection for Remote and Global mirror” on page 549 for more information on consistency protection.

6. In the next window, select the location of the auxiliary volumes, as shown in Figure 10-103:
- **On this system**, which means that the volumes are local.
 - **On another system**, which means that you select the remote system from the drop-down list.

After you make a selection, click **Next**.



Create Relationship

Where are the auxiliary volumes located?

☐ On this system

☒ On another system

mcr-tb5-cluster-123

Cancel

Back Next

Figure 10-103 Specifying the location of the auxiliary volumes

7. In the New Relationship window that is shown in Figure 10-104, you can create relationships. Select a master volume in the Master drop-down list. Then, select an auxiliary volume in the Auxiliary drop-down list for this master and click **Add**. If needed, repeat this step to create other relationships.

The screenshot shows a window titled "Create Relationship" with a close button (X) in the top right corner. Below the title bar, there is a instruction: "Select the master and auxiliary volumes to use in the relationship." The main area contains two panels. The left panel, labeled "Master", has a drop-down menu showing "ITSO_RC_SRC_VOL_1" and a capacity of "3.00 GiB". An arrow points from this panel to the right panel, labeled "Auxiliary". The right panel has a drop-down menu showing "ITSO_RC_TGT_VOL_1" and a blue "Add" button with a hand cursor icon. At the bottom of the window, there are three buttons: "Cancel", "Back", and "Next".

Figure 10-104 Select a volume for mirroring

Important: The master and auxiliary volumes must be of equal size. Therefore, only the targets with the appropriate size are shown in the list for a specific source volume.

8. Since **Add Consistency Protection** was checked as shown in 5 on page 564, you will get a dialog window asking whether you want to add change volume as shown in Figure 10-105.

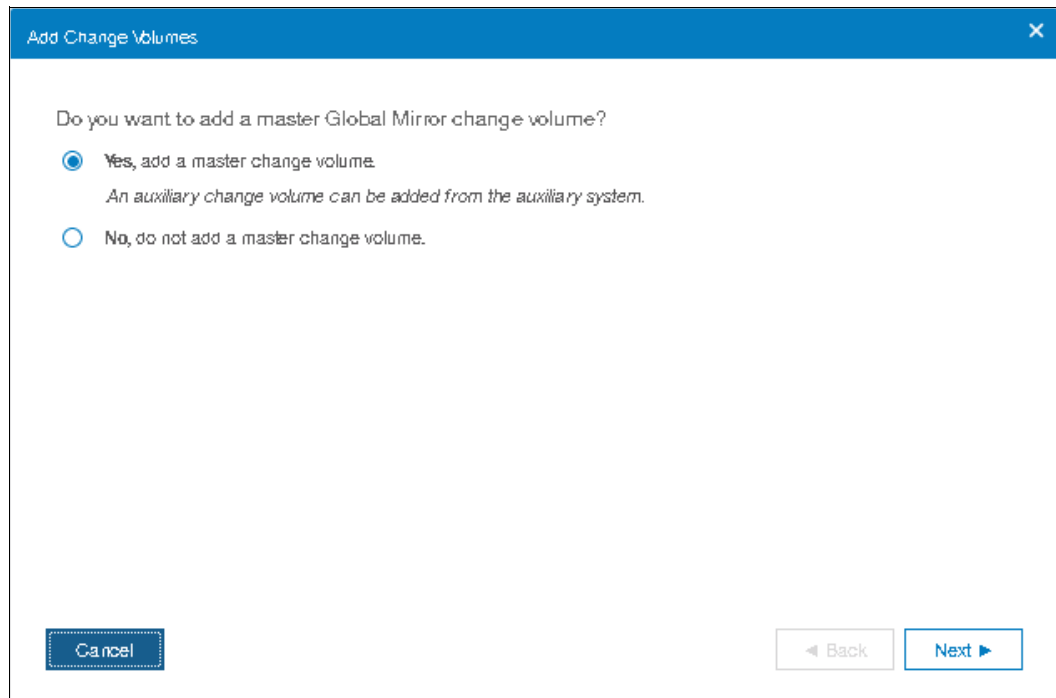


Figure 10-105 Add Change Volume

Click **Next**. You will get a dialog window asking whether you want to add a new change volume or use an existing one. In our example, we chose to create a new master change volume.as shown in Figure 10-106.

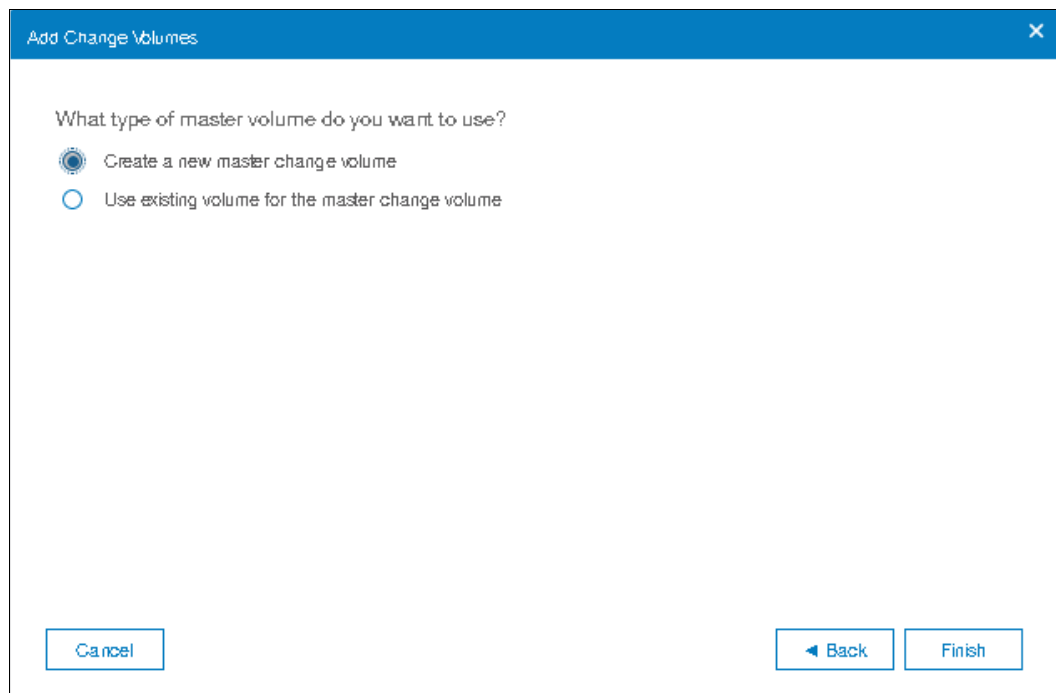
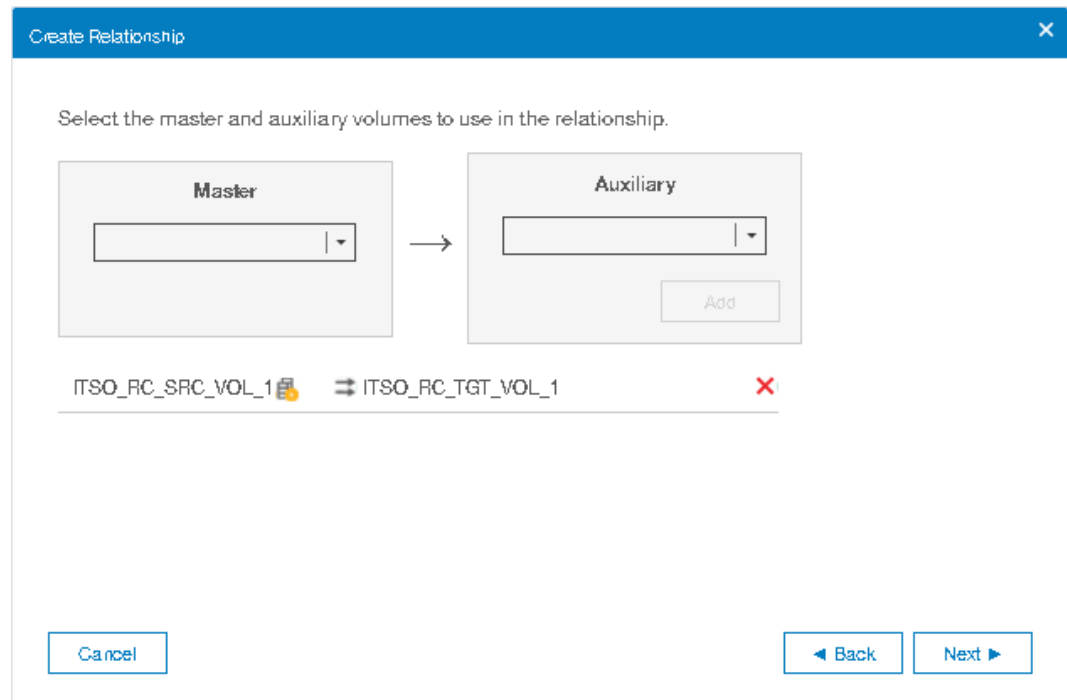


Figure 10-106 Create a new master change volume or use an existing one

9. Click **Finish**. A dialog window will be presented where you can remove a relationship that was created, click **X**, as shown in Figure 10-107.

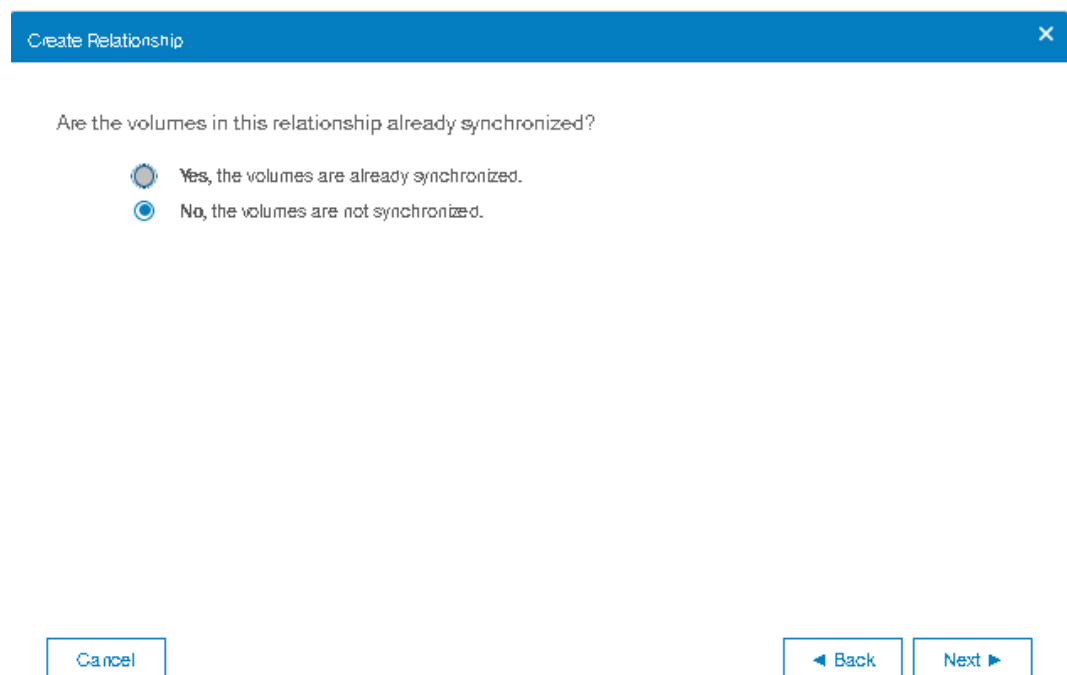


The 'Create Relationship' dialog box has a blue title bar with the text 'Create Relationship' and a close button (X). The main area contains the instruction 'Select the master and auxiliary volumes to use in the relationship.' Below this, there are two sections: 'Master' and 'Auxiliary'. Each section has a dropdown menu. An arrow points from the Master dropdown to the Auxiliary dropdown. Below these sections, there is a list of volume pairs. The first pair is 'ITSO_RC_SRC_VOL_1' (with a yellow icon) and 'ITSO_RC_TGT_VOL_1' (with a red X icon). At the bottom, there are three buttons: 'Cancel', 'Back' (with a left arrow), and 'Next' (with a right arrow).

Figure 10-107 Create the relationships between the master and auxiliary volumes

After all of the relationships that you want to create are shown, click **Next**.

10. Specify whether the volumes are synchronized, as shown in Figure 10-108. Then, click **Next**.



The 'Create Relationship' dialog box has a blue title bar with the text 'Create Relationship' and a close button (X). The main area contains the question 'Are the volumes in this relationship already synchronized?'. Below this, there are two radio button options: 'Yes, the volumes are already synchronized.' (which is selected) and 'No, the volumes are not synchronized.' At the bottom, there are three buttons: 'Cancel', 'Back' (with a left arrow), and 'Next' (with a right arrow).

Figure 10-108 Volumes are already synchronized

11. In the next window, select whether you want to start to copy the data and click **Finish**, as shown in Figure 10-109.

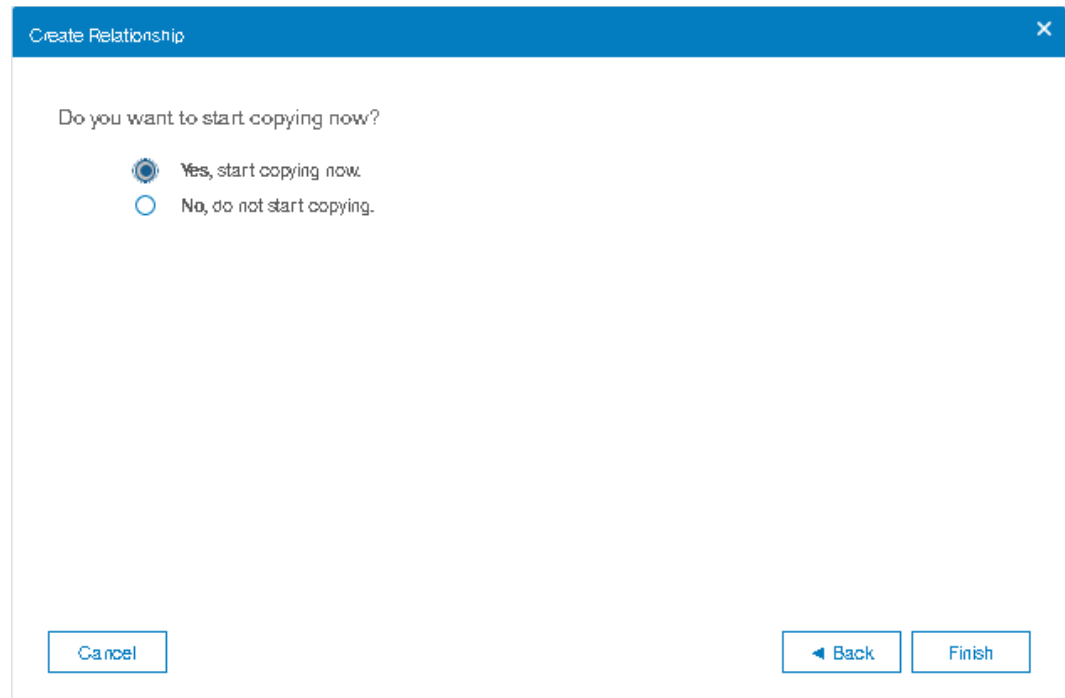


Figure 10-109 Synchronize now

12. Figure 10-110 shows that the task to create the relationship is complete.

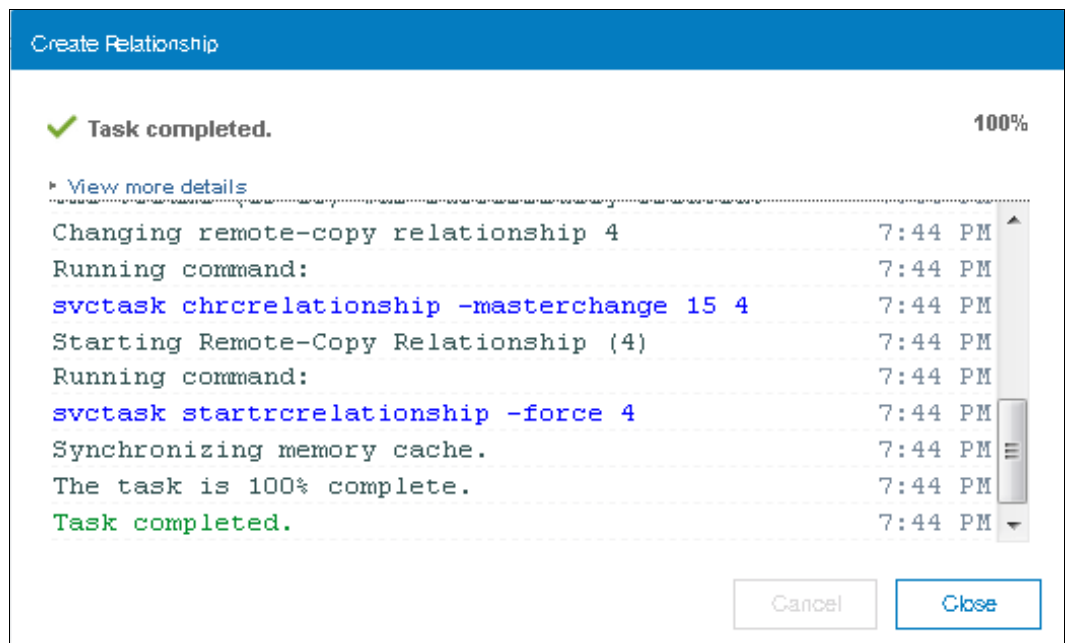
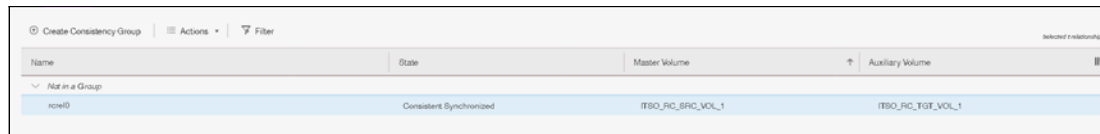


Figure 10-110 Creation of Remote Copy relationship complete

The relationships are visible in the Remote Copy pane. If you selected to copy the data, you can see that the status is Consistent Copying. You can check the copying progress in the Running Tasks status area.

After the copy is finished, the relationship status changes to Consistent synchronized. Figure 10-111 shows the Consistent Synchronized status.



Name	State	Master Volume	Auxiliary Volume
rcel0	Consistent Synchronized	ITSO_RC_SRC_VOL_1	ITSO_RC_TGT_VOL_1

Figure 10-111 Consistent copy of the mirrored volumes

10.10.3 Creating Consistency Group

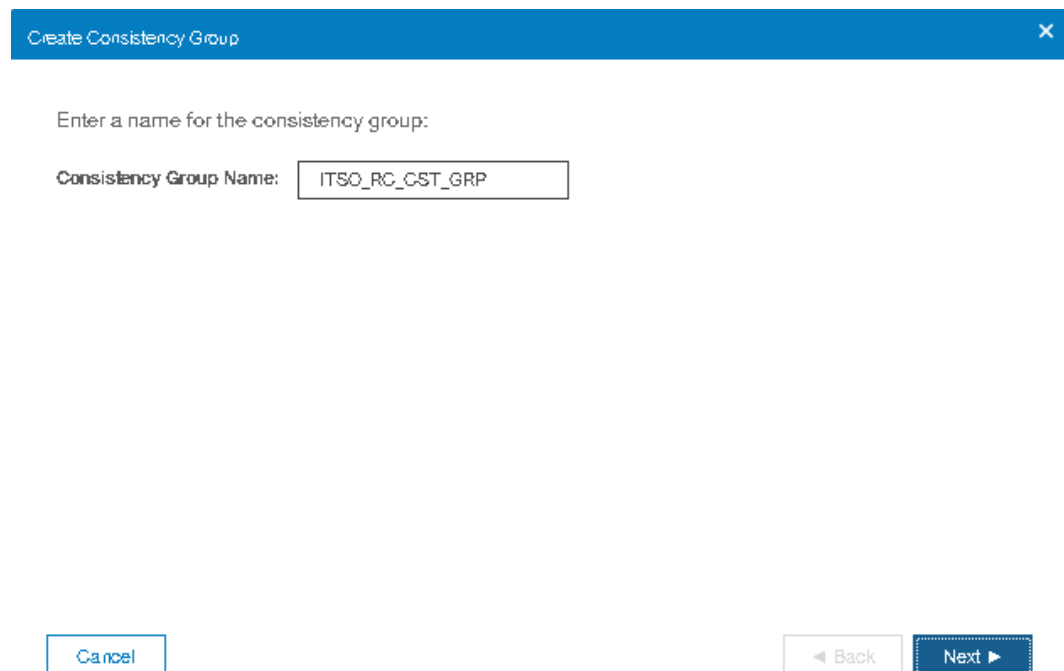
To create a Consistency Group, complete the following steps:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. Click **Create Consistency Group**, as shown in Figure 10-112.



Figure 10-112 Selecting the Create Consistency Group option

3. Enter a name for the Consistency Group, and then, click **Next**, as shown in Figure 10-113.



Create Consistency Group

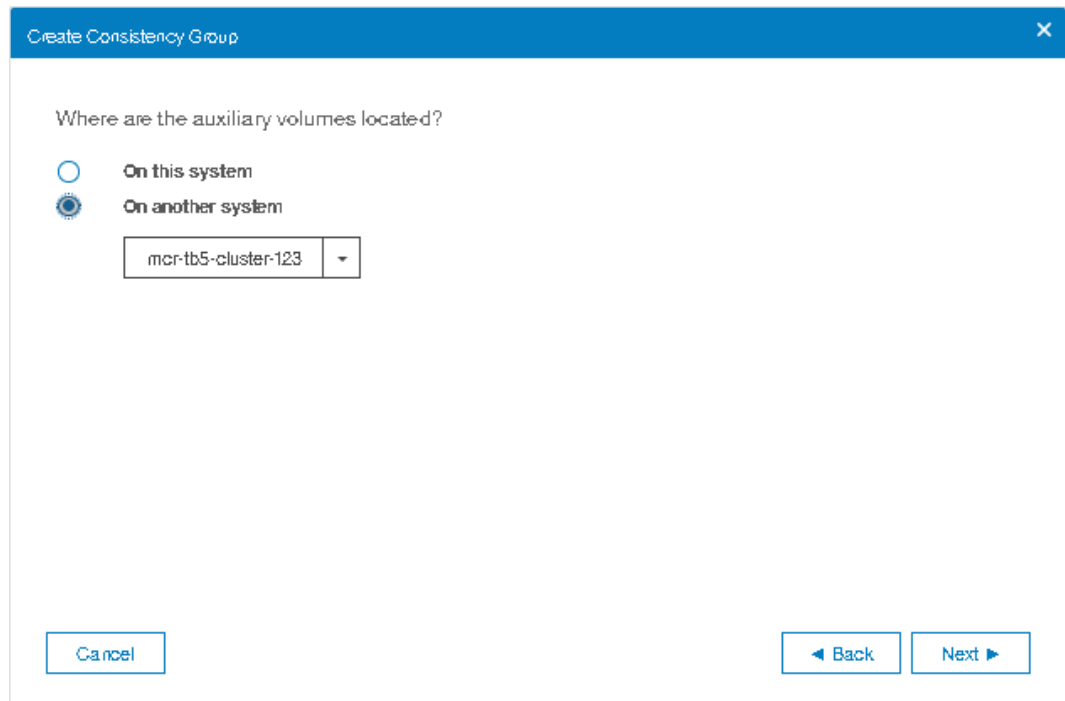
Enter a name for the consistency group:

Consistency Group Name:

Figure 10-113 Enter a Consistency Group name

4. In the next window, select where the auxiliary volumes are located, as shown in Figure 10-114 on page 571:
 - On this system, which means that the volumes are local
 - On another system, which means that you select the remote system in the drop-down list

After you make a selection, click **Next**.



Create Consistency Group

Where are the auxiliary volumes located?

☐ On this system

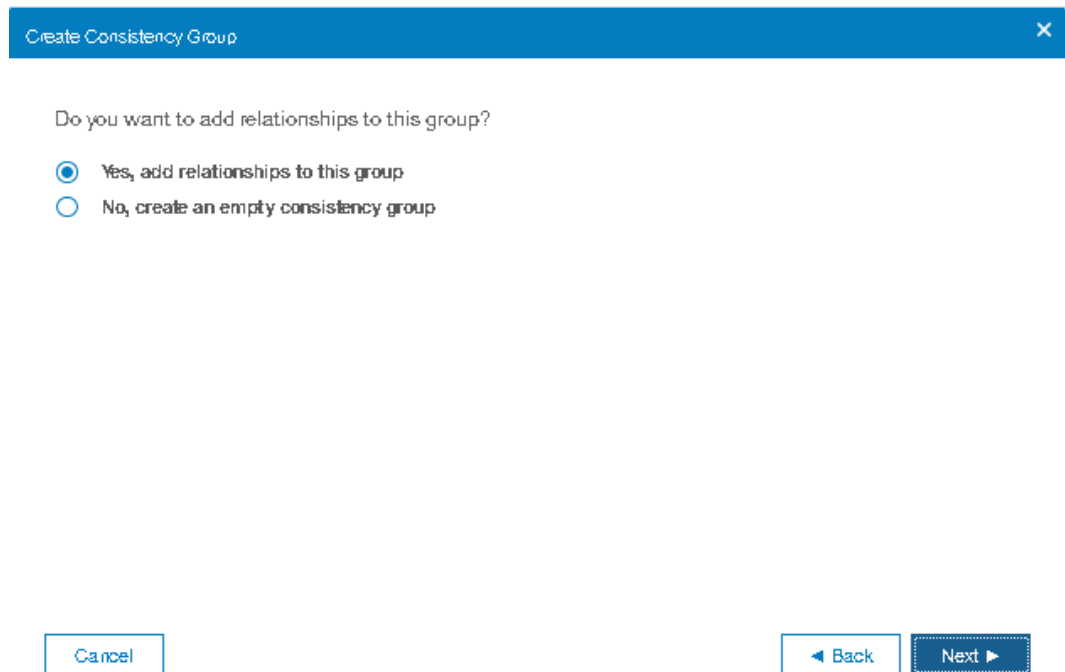
☒ On another system

mcr-tb5-cluster-123 ▼

Cancel < Back Next >

Figure 10-114 Location of auxiliary volumes

5. Select whether you want to add relationships to this group, as shown in Figure 10-115 on page 572. The following options are available:
 - If you select Yes, click **Next** to continue the wizard and go to step 6.
 - If you select No, click **Finish** to create an empty Consistency Group that can be used later.



Create Consistency Group

Do you want to add relationships to this group?

☒ Yes, add relationships to this group

☐ No, create an empty consistency group

Cancel

Back

Next

Figure 10-115 Add relationships to this group

6. Select one of the following types of relationships to create, as shown in Figure 10-116 on page 573:
 - Metro Mirror
 - Global Mirror
 - Global Mirror with Change Volumes

Note: For Metro Mirror or regular Global Mirror, also indicate whether to add consistency protection.

Click **Next**.

Create Consistency Group

Select the type of copy that you want to create:

☒ Metro Mirror

☐ Global Mirror

☒ Add Consistency Protection

☐ Global Mirror with Change Volumes

Cancel Back Next

Figure 10-116 Select the type of relationship that you want to create

7. As shown in Figure 10-117 on page 574, you can optionally select existing relationships to add to the group. Click **Next**.

Note: To select multiple relationships, hold down Ctrl and click the entries that you want to include.

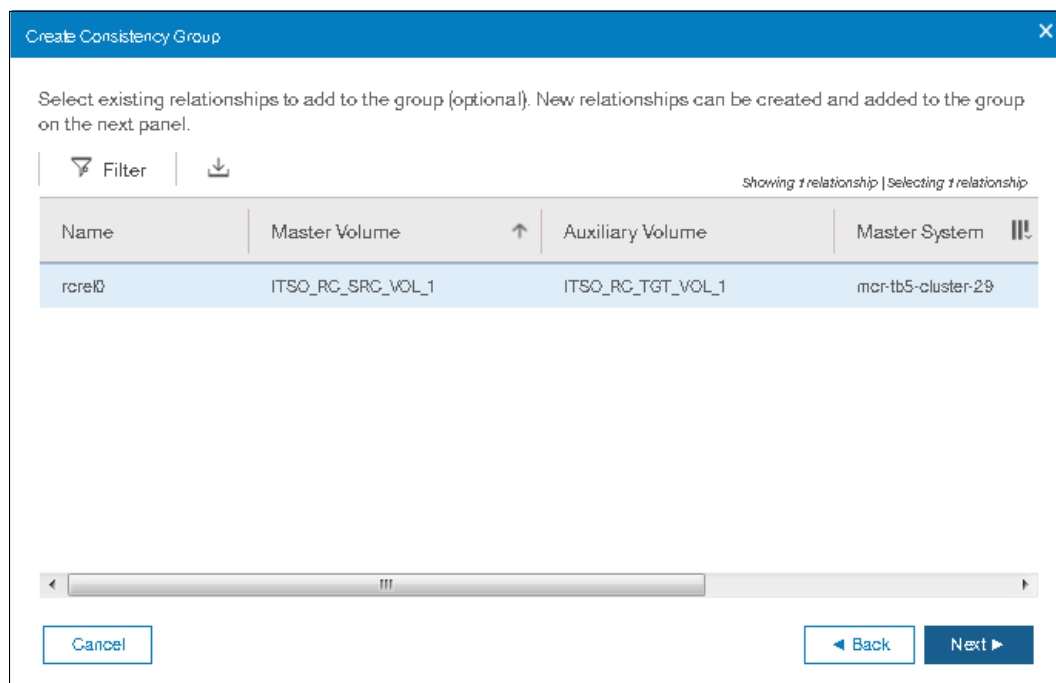


Figure 10-117 Select existing relationships to add to the group

8. In the window that is shown in Figure 10-120 on page 576, you can create relationships.
 - a. Select a volume in the Master drop-down list.
 - b. Then, select a volume in the Auxiliary drop-down list for this master.
 - c. Click **Add**.

Note: When the first relationship is added to an empty group, the group takes on the same state, primary (copy direction), type (Metro Mirror or Global Mirror), and cycling mode as the relationship. Subsequent relationships must have the same state, copy direction, and type as the group in order to be added to it. A relationship can belong to only one consistency group.

- d. As shown in Figure 10-118 on page 575, select whether you would like to add a change volume or not.




Figure 10-118 Add a change volume

- e. Click **Next**.

Note: Selecting to add a change volume, will not add a change volume to the remote system. Creation of change volume on the remote system for the auxiliary has to be done manually.

- f. Select whether you want to create a new master change volume or use an existing one as shown in Figure 10-119.




Figure 10-119 Master change volume selection

- g. Click **Finish**.
 - h. Repeat this step to create other relationships, if needed.
9. To remove a relationship that was created, click **X** (Figure 10-120). After all of the relationships that you want to create are displayed, click **Next**.

Create Consistency Group

Select the master and auxiliary volumes for new remote copy relationships to add to the remote-copy consistency group. (optional)

Master

Auxiliary

ITSO_RC_SRC_VOL_2 → ITSO_RC_TGT_VOL_2

Cancel Back Next

Figure 10-120 Create relationships between the master and auxiliary volumes

10. Specify whether the volumes are already synchronized. Then, click **Next** (Figure 10-121).

Create Consistency Group

Are the volumes already synchronized?

☐ Yes, the volumes are already synchronized.

☒ No, the volumes are not synchronized.

Cancel Back Next

Figure 10-121 Volumes are already synchronized

11. In the last window, select whether you want to start to copy the data. Then, click **Finish**, as shown in Figure 10-122.

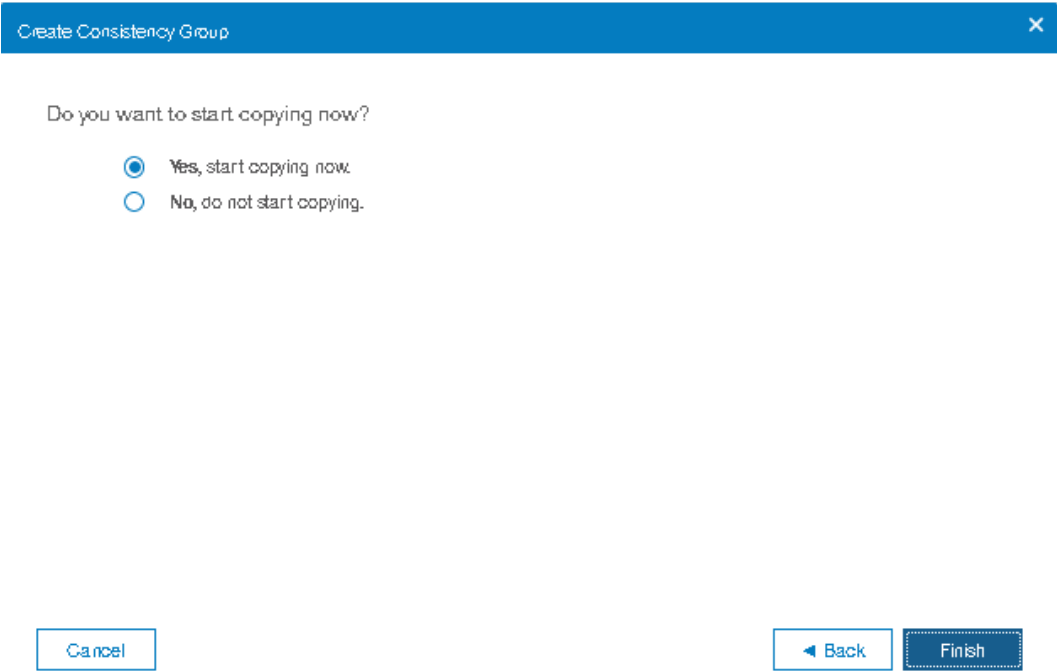


Figure 10-122 Synchronize now

12. The relationships are visible in the Remote Copy pane. If you selected to copy the data, you can see that the status of the relationships is Inconsistent copying. You can check the copying progress in the Running Tasks status area, as shown in Figure 10-123.

Name	State	Master Volume	Auxiliary Volume
Not in a Group			
rcp0	Inconsistent Copying	mcn-b5-cluster-29	mcn-b5-cluster-123
rcp0	Inconsistent Copying	ITSO_RC_SRC_VOL_1	ITSO_RC_TGT_VOL_1
rcp1	Inconsistent Copying	ITSO_RC_SRC_VOL_2	ITSO_RC_TGT_VOL_2

Figure 10-123 Consistency Group created with relationship in copying and synchronized status

After the copies are completed, the relationships and the Consistency Group change to the Consistent Synchronized status.

10.10.4 Renaming Consistency Group

To rename a Consistency Group, complete the following steps:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. In the pane, select the Consistency Group that you want to rename. Then, select **Actions** → **Rename**, as shown in Figure 10-124 on page 578.




Figure 10-124 Renaming a Consistency Group

Note: You can also right-click a remote copy consistency group and select **Rename**.

3. Enter the new name that you want to assign to the Consistency Group and click **Rename**, as shown in Figure 10-125.

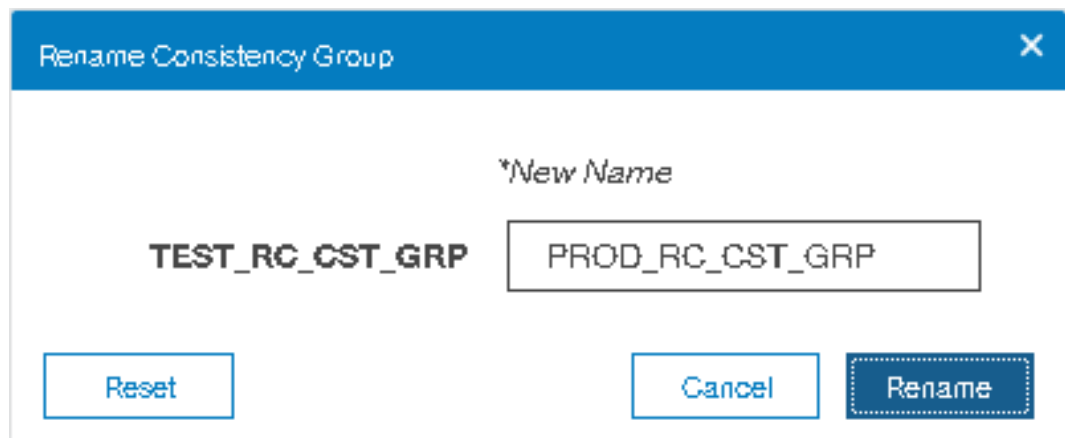


Figure 10-125 Changing the name for a Consistency Group

The new Consistency Group name is displayed on the Remote Copy pane.

10.10.5 Renaming remote copy relationship

Complete the following steps to rename a remote copy relationship:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. In the table, select the remote copy relationship mapping that you want to rename. Click **Actions** → **Rename**, as shown in Figure 10-126.

Tip: You can also right-click a remote copy relationship and select **Rename**.



Figure 10-126 Rename remote copy relationship action

3. In the Rename Relationship window, enter the new name that you want to assign to the FlashCopy mapping and click **Rename**, as shown in Figure 10-127.

Figure 10-127 Renaming a remote copy relationship

Remote copy relationship name: You can use the letters A - Z and a - z, the numbers 0 - 9, and the underscore (_) character. The remote copy name can be 1 - 15 characters. No blanks are allowed.

10.10.6 Moving stand-alone remote copy relationship to Consistency Group

Complete the following steps to move a remote copy relationship to a Consistency Group:

1. From the main navigation pane, click **Copy Services** → **Remote Copy**.
2. Expand the **Not in a Group** column.
3. Select the relationship that you want to move to the Consistency Group.
4. Click **Actions** → **Add to Consistency Group**, as shown in Figure 10-128.

Tip: You can also right-click a remote copy relationship and select **Add to Consistency Group**.

Name	State	Master Volume	Auxiliary Volume
rcrel2	Consistent Synchronized	SRC_VOL_1	TGT_VOL_1
PRD_RC_CST_0	Consistent Synchronized	mcv-ib5-cluster-29 → mcr-ib5-cluster-123	
TST_RC_CST_0	Consistent Synchronized	mcv-ib5-cluster-29 → mcr-ib5-cluster-123	
rcrel1	Consistent Synchronized	ITSO_RC_SRC_VOL_2	ITSO_RC_TGT_VOL_2
rcrel0	Consistent Synchronized	ITSO_RC_SRC_VOL_1	ITSO_RC_TGT_VOL_1

Figure 10-128 Add to Consistency Group action

5. In the Add Relationship to Consistency Group window, select the Consistency Group for this remote copy relationship by using the drop-down list, as shown in Figure 10-129 on page 580. Click **Add to Consistency Group** to confirm your changes.

Add Relationship to Consistency Group
✕

Select the consistency group to move the relationship rcrel2

Consistency Group

ITSO_RC_CST_GRP

▼

Cancel

Add to Consistency Group

Figure 10-129 Adding a relationship to a Consistency Group

Note: The state of the remote copy consistency group and the recopy copy relationship that is being added, should match, otherwise you can not add that remote copy relationship into the existing remote copy consistency group.

10.10.7 Removing remote copy relationship from Consistency Group

Complete the following steps to remove a remote copy relationship from a Consistency Group:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. Select a Consistency Group.
3. Select the remote copy relationship that you want to remove from the Consistency Group.
4. Click **Actions** → **Remove from Consistency Group**, as shown in Figure 10-130.

Tip: You can also right-click a remote copy relationship and select **Remove from Consistency Group**.




Figure 10-130 Remove from Consistency Group action

5. In the Remove Relationship From Consistency Group window, click **Remove**, as shown in Figure 10-131 on page 581.

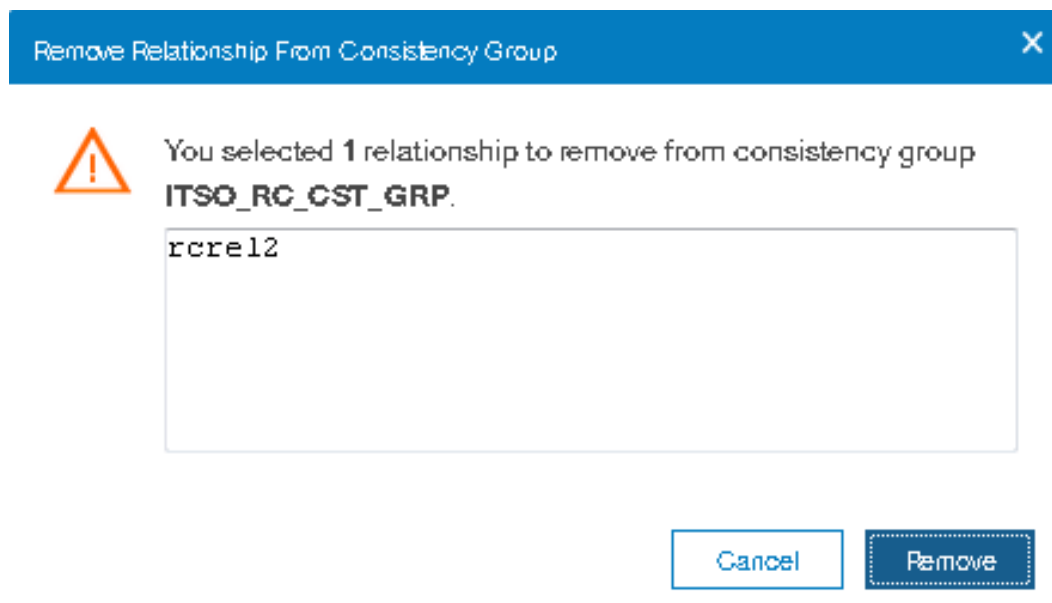


Figure 10-131 Remove a relationship from a Consistency Group

10.10.8 Starting remote copy relationship

When a remote copy relationship is created, the remote copy process can be started. Only relationships that are not members of a Consistency Group, or the only relationship in a Consistency Group, can be started individually.

Complete the following steps to start a remote copy relationship:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. Expand the **Not in a Group** column.
3. In the table, select the remote copy relationship that you want to start.
4. Click **Actions** → **Start** to start the remote copy process, as shown in Figure 10-132.

Tip: You can also right-click a relationship and select **Start** from the list.

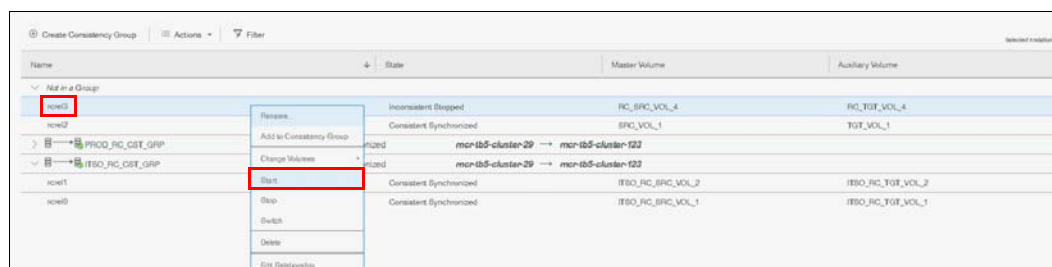


Figure 10-132 Starting the remote copy process

5. After the task is complete, the remote copy relationship status has a Consistent Synchronized state, as shown in Figure 10-133 on page 582.

Create Consistency Group Actions Filter					Selected relationship
Name	+	State	Master Volume	Auxiliary Volume	
Not in a Group					
rcw3		Consistent Synchronized	RC_SRC_VOL_4	RC_TGT_VOL_4	
rcw2		Consistent Synchronized	SRC_VOL_1	TGT_VOL_1	
PROD_RC_CST_GRP		Consistent Synchronized	mcv-b5-cluster-29 → mcr-b5-cluster-123		
ITSO_RC_CST_GRP		Consistent Synchronized	mcv-b5-cluster-29 → mcr-b5-cluster-123		
rcw1		Consistent Synchronized	ITSO_RC_SRC_VOL_2	ITSO_RC_TGT_VOL_2	
rcw0		Consistent Synchronized	ITSO_RC_SRC_VOL_1	ITSO_RC_TGT_VOL_1	

Figure 10-133 Consistent Synchronized remote copy relationship

10.10.9 Starting remote copy Consistency Group

All of the mappings in a Consistency Group are brought to the same state. To start the remote copy Consistency Group, complete the following steps:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. Select the Consistency Group that you want to start, as shown in Figure 10-134.

Create Consistency Group Actions Filter					Selected remote copy consistency group
Name	+	State	Master Volume	Auxiliary Volume	
Not in a Group					
rcw3		Consistent Synchronized	RC_SRC_VOL_4	RC_TGT_VOL_4	
rcw2		Consistent Synchronized	SRC_VOL_1	TGT_VOL_1	
PROD_RC_CST_GRP		Consistent Stopped	mcv-b5-cluster-29 → mcr-b5-cluster-123		
ITSO_RC_CST_GRP		Consistent Stopped	mcv-b5-cluster-29 → mcr-b5-cluster-123		
rcw1		Consistent Stopped	ITSO_RC_SRC_VOL_2	ITSO_RC_TGT_VOL_2	
rcw0		Consistent Stopped	ITSO_RC_SRC_VOL_1	ITSO_RC_TGT_VOL_1	

Figure 10-134 Remote Copy Consistency Groups view

3. Click **Actions** → **Start** (Figure 10-135) to start the remote copy Consistency Group.

Create Consistency Group Actions Filter					Selected remote copy consistency group
Name	+	State	Master Volume	Auxiliary Volume	
Not in a Group					
rcw3		Consistent Synchronized	RC_SRC_VOL_4	RC_TGT_VOL_4	
rcw2		Consistent Synchronized	SRC_VOL_1	TGT_VOL_1	
PROD_RC_CST_GRP		Consistent Stopped	mcv-b5-cluster-29 → mcr-b5-cluster-123		
ITSO_RC_CST_GRP		Consistent Stopped	mcv-b5-cluster-29 → mcr-b5-cluster-123		
rcw1		Consistent Stopped	ITSO_RC_SRC_VOL_2	ITSO_RC_TGT_VOL_2	
rcw0		Consistent Stopped	ITSO_RC_SRC_VOL_1	ITSO_RC_TGT_VOL_1	

Figure 10-135 Start action Figure 4

4. You can check the remote copy Consistency Group progress, as shown in Figure 10-136.

Create Consistency Group Actions Filter					Selected remote copy consistency group
Name	+	State	Master Volume	Auxiliary Volume	
Not in a Group					
rcw3		Consistent Synchronized	RC_SRC_VOL_4	RC_TGT_VOL_4	
rcw2		Consistent Synchronized	SRC_VOL_1	TGT_VOL_1	
DB2_TEST_RC		Consistent Synchronized	mcv-b5-cluster-29 → mcr-b5-cluster-123		
ITSO_RC_CST_GRP		Consistent Synchronized	mcv-b5-cluster-29 → mcr-b5-cluster-123		
rcw1		Consistent Synchronized	ITSO_RC_SRC_VOL_2	ITSO_RC_TGT_VOL_2	
rcw0		Consistent Synchronized	ITSO_RC_SRC_VOL_1	ITSO_RC_TGT_VOL_1	

Figure 10-136 Checking the remote copy Consistency Group progress

5. After the task completes, the Consistency Group and all of its relationships becomes in a Consistent Synchronized state.

10.10.10 Switching copy direction

When a remote copy relationship is in the Consistent synchronized state, the copy direction for the relationship can be changed. Only relationships that are not a member of a Consistency Group (or the only relationship in a Consistency Group) can be switched individually. These relationships can be switched from master to auxiliary or from auxiliary to master, depending on the case.

Complete the following steps to switch a remote copy relationship:

1. From the System pane, select **Copy Services** → **Remote Copy**.
2. Expand the **Not in a Group** column.
3. In the table, select the remote copy relationship that you want to switch.
4. Click **Actions** → **Switch** (Figure 10-137) to start the remote copy process.

Tip: You can also right-click a relationship and select **Switch**.




Figure 10-137 Switch copy direction action

5. The Warning window that is shown in Figure 10-138 opens. A confirmation is needed to switch the remote copy relationship direction. The remote copy is switched from the master volume to the auxiliary volume. Click **Yes**.

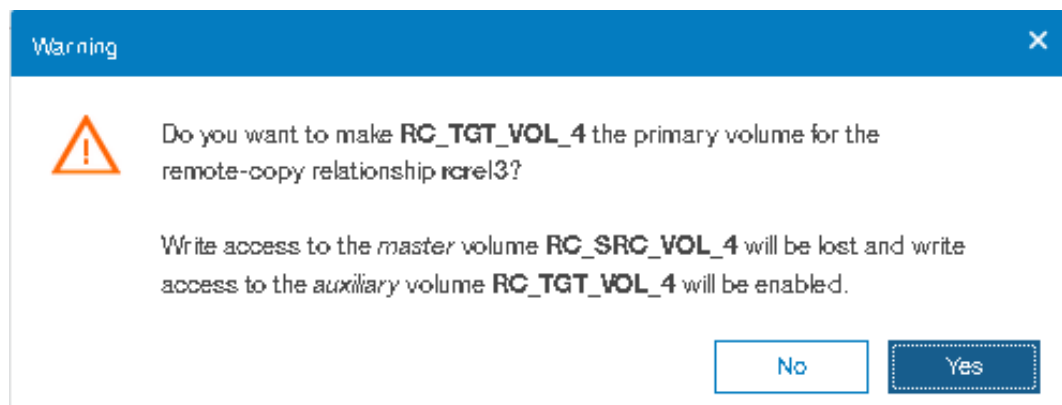


Figure 10-138 Warning window

Figure 10-139 on page 584 shows the command-line output about this task.

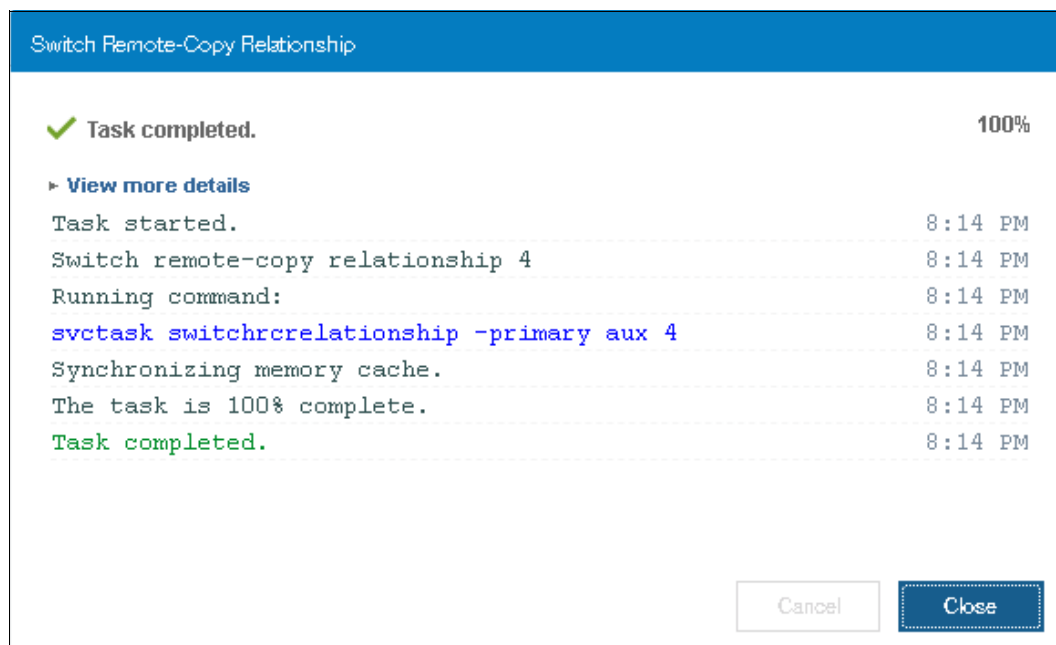


Figure 10-139 Command-line output for switch relationship action

The copy direction is now switched, as shown in Figure 10-140 with a red circle. The auxiliary volume is now accessible and shown as the primary volume. Also, the auxiliary volume is now synchronized to the master volume.

Create Consistency Group				Actions	Filter	Selected Relationship	
Name	State	Master Volume		Auxiliary Volume			
Not in a Group							
rcrc3	Consistent Synchronized	RC_SRC_VOL_4		RC_TGT_VOL_4			
rcrc2	Consistent Synchronized	SRC_VOL_1		TGT_VOL_1			
<div>ITSO_RC_CST_GRP</div> Consistent Synchronizedmcr-ib5-cluster-29 → mcr-ib5-cluster-123							
<div>PROD_RC_CST_GRP</div> Consistent Synchronizedmcr-ib5-cluster-29 → mcr-ib5-cluster-123							

Figure 10-140 Checking remote copy synchronization direction

10.10.11 Switching the copy direction for a Consistency Group

When a Consistency Group is in the Consistent Synchronized state, the copy direction for this Consistency Group can be changed.

Important: When the copy direction is switched, it is crucial that no outstanding I/O exists to the volume that changes from primary to secondary because all of the I/O is inhibited to that volume when it becomes the secondary. Therefore, careful planning is required before you switch the copy direction for a Consistency Group.

Complete the following steps to switch a Consistency Group:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. Select the Consistency Group that you want to switch.
3. Click **Actions** → **Switch** (as shown in Figure 10-141 on page 585) to start the remote copy process.

Tip: You can also right-click a relationship and select **Switch**.

Name	State	Master Volume	Auxiliary Volume
Not in a Group			
rcw03	Consistent Synchronized	RC_SRC_VOL_4	RC_TGT_VOL_4
rcw02	Consistent Synchronized	SPC_VOL_1	TGT_VOL_1
PROD_RC_CST_GRP	Consistent Synchronized	mcr-tb5-cluster-29 → mcr-tb5-cluster-123	
DB2_TEST_RC	Consistent Synchronized	ITSO_RC_SRC_VOL_3	ITSO_RC_TGT_VOL_3
ITSO_RC_CST_GRP	Consistent Synchronized	mcr-tb5-cluster-29 → mcr-tb5-cluster-123	

Figure 10-141 Switch action

- The warning window that is shown in Figure 10-142 opens. A confirmation is needed to switch the Consistency Group direction. In the example that is shown in here, the Consistency Group is switched from the master group to the auxiliary group. Click **Yes**.

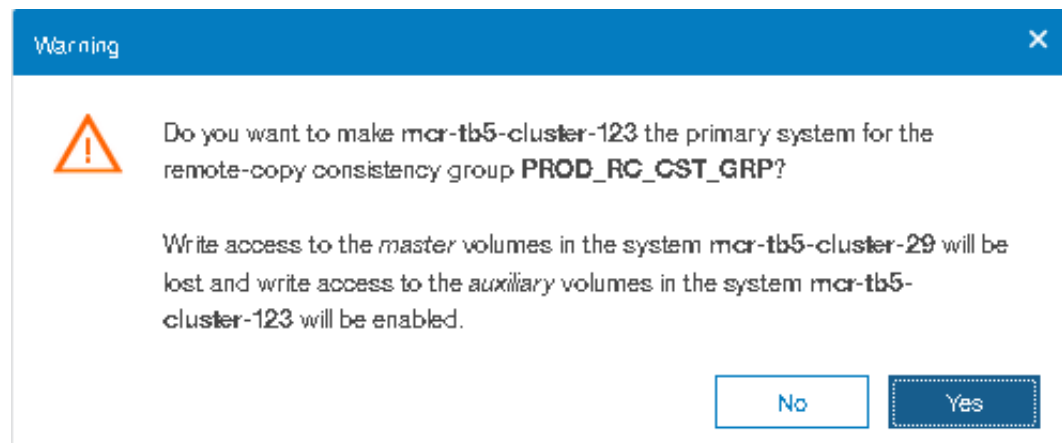


Figure 10-142 Warning window before switching the relationship

- The remote copy direction is now switched as shown in Figure 10-143. The auxiliary volume is now accessible and shown as a primary volume.

Name	State	Master Volume	Auxiliary Volume
Not in a Group			
rcw03	Consistent Synchronized	RC_SRC_VOL_4	RC_TGT_VOL_4
rcw02	Consistent Synchronized	SPC_VOL_1	TGT_VOL_1
PROD_RC_CST_GRP	Consistent Synchronized	mcr-tb5-cluster-123 → mcr-tb5-cluster-29	
DB2_TEST_RC	Consistent Synchronized	ITSO_RC_SRC_VOL_3	ITSO_RC_TGT_VOL_3
ITSO_RC_CST_GRP	Consistent Synchronized	mcr-tb5-cluster-29 → mcr-tb5-cluster-123	

Figure 10-143 Relationship after switch

10.10.12 Stopping a remote copy relationship

After it is started, the remote copy process can be stopped, if needed. Only relationships that are not a member of a Consistency Group (or the only relationship in a Consistency Group) can be stopped individually. You can also use this procedure to enable write access to a consistent secondary volume.

Complete the following steps to stop a remote copy relationship:

- From the main navigation pane, select **Copy Services** → **Remote Copy**.
- Expand the **Not in a Group** column.
- In the table, select the remote copy relationship that you want to stop.

4. Click **Actions** → **Stop** (as shown in Figure 10-144) to stop the remote copy process.

Tip: You can also right-click a relationship and select **Stop** from the list.




Figure 10-144 Stop action

5. The Stop Remote Copy Relationship window opens, as shown in Figure 10-145. To allow secondary read/write access, select **Allow secondary read/write access**. Then, click **Stop Relationship**.

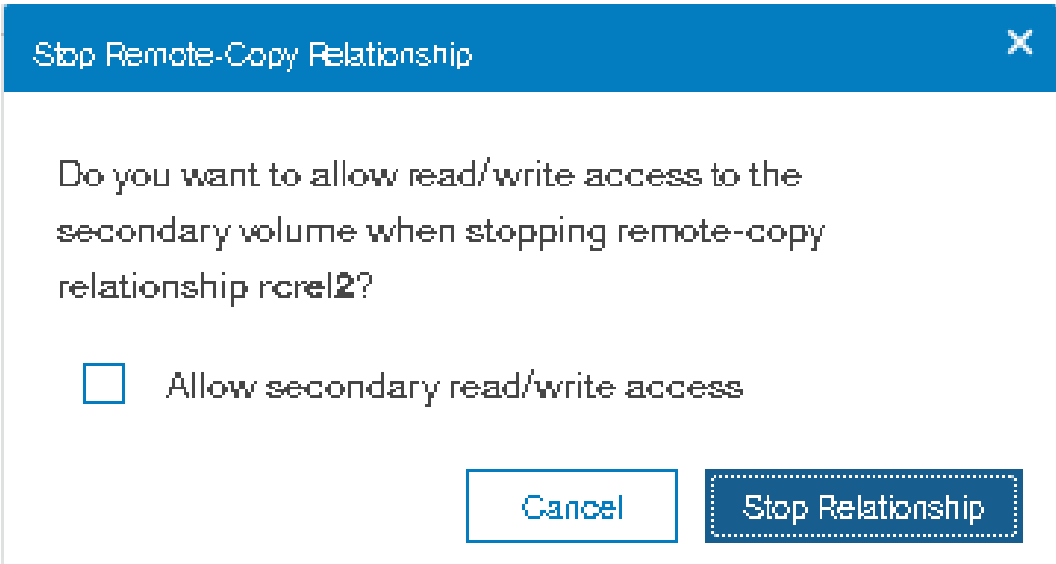


Figure 10-145 Stop Remote Copy Relationship window

6. Figure 10-146 on page 587 shows the command-line output for the stop remote copy relationship.



Figure 10-146 Stop remote copy relationship command-line output

The new relationship status can be checked, as shown in Figure 10-147. The relationship is now Consistent Stopped.

The screenshot shows a table titled "Consistency Groups" with columns: Name, State, Master Volume, and Auxiliary Volume. The "State" column for the "rcs2" relationship is circled in red and labeled "Consistent Stopped".

Name	State	Master Volume	Auxiliary Volume
rcs2	Consistent Stopped	RC_SRC_VOL_4	RC_TGT_VOL_4
DB2_TEST_RC	Consistent Synchronized	DB2_SRC_VOL_1	DB2_TGT_VOL_1
DB2_TEST_RC_GRP	Consistent Synchronized	DB2_SRC_VOL_1	DB2_TGT_VOL_1
DB2_TEST_RC_GRP	Consistent Synchronized	DB2_SRC_VOL_1	DB2_TGT_VOL_1
rcs1	Consistent Synchronized	RC_SRC_VOL_2	RC_TGT_VOL_2
rcs0	Consistent Synchronized	RC_SRC_VOL_1	RC_TGT_VOL_1

Figure 10-147 Checking remote copy synchronization status

10.10.13 Stopping Consistency Group

After it is started, the Consistency Group can be stopped, if necessary. You can also use this task to temporarily write access to consistent secondary volumes.

Perform the following steps to stop a Consistency Group:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. In the table, select the Consistency Group that you want to stop.
3. Click **Actions** → **Stop** (as shown in Figure 10-148 on page 588) to stop the remote copy Consistency Group.

Tip: You can also right-click a relationship and select **Stop** from the list.

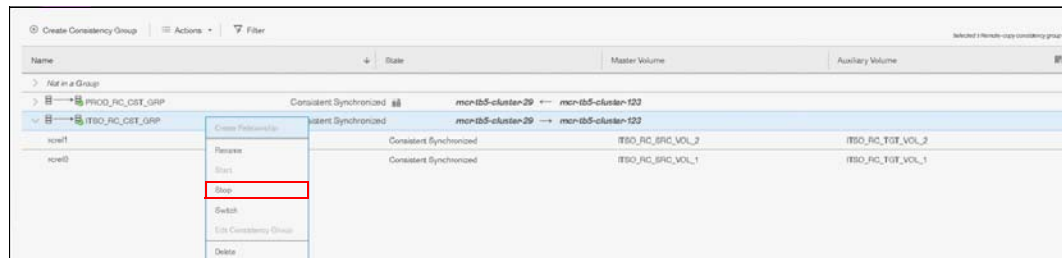


Figure 10-148 Selecting the Stop option

4. The Stop Remote Copy Consistency Group window opens, as shown in Figure 10-149. To allow secondary read/write access, select **Allow secondary read/write access**. Then, click **Stop Consistency Group**.

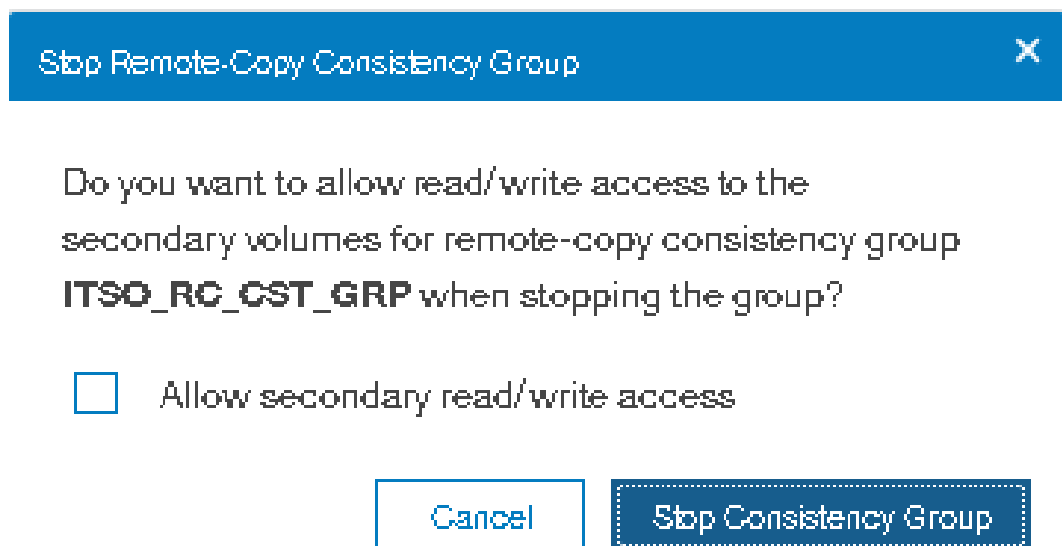


Figure 10-149 Stop Remote Copy Consistency Group window

The new relationship status can be checked, as shown in Figure 10-150. The relationship is now Consistent Stopped.




Figure 10-150 Checking remote copy synchronization status

10.10.14 Deleting stand-alone remote copy relationships

Complete the following steps to delete a stand-alone remote copy mapping:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.
2. In the table, select the remote copy relationship that you want to delete.

Multiple remote copy mappings: To select multiple remote copy mappings, hold down Ctrl and click the entries that you want.

3. Click **Actions** → **Delete**, as shown in Figure 10-151.

Tip: You can also right-click a remote copy mapping and select **Delete**.

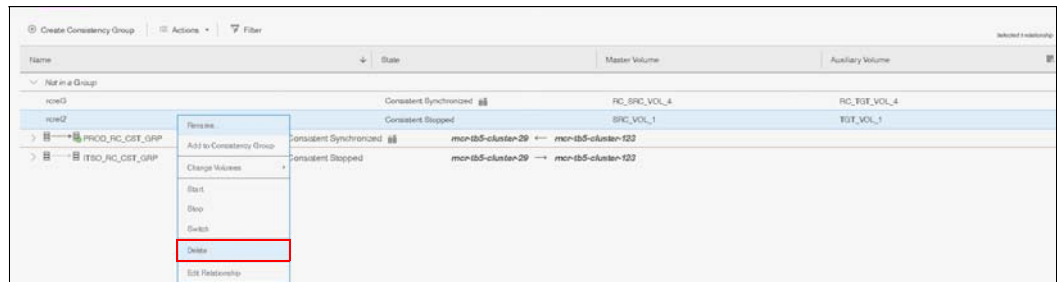


Figure 10-151 Selecting the Delete Relationship option

4. The Delete Relationship window opens (Figure 10-152). In the “Verify the number of relationships that you are deleting” field, enter the number of volumes that you want to remove. This verification was added to help to avoid deleting the wrong relationships. Click **Delete**.

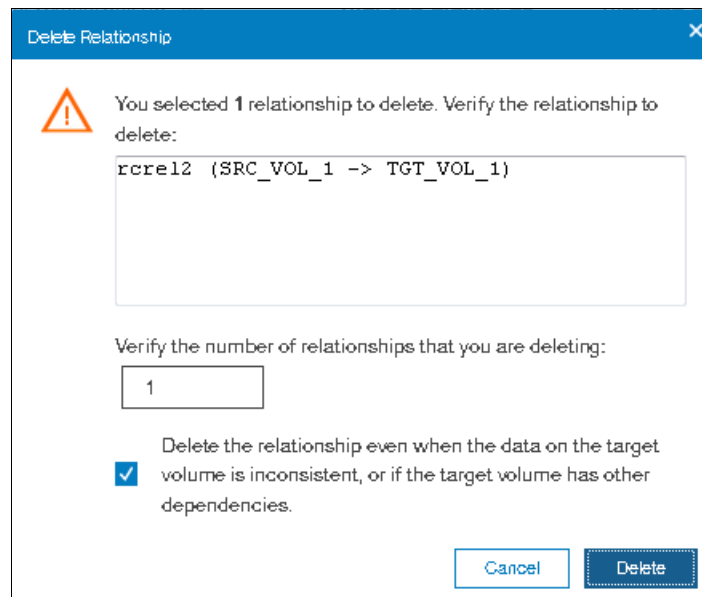


Figure 10-152 Delete remote copy relationship

10.10.15 Deleting Consistency Group

Important: Deleting a Consistency Group does not delete its remote copy mappings.

Complete the following steps to delete a Consistency Group:

1. From the main navigation pane, select **Copy Services** → **Remote Copy**.

2. In the left column, select the Consistency Group that you want to delete.
3. Click **Actions** → **Delete**, as shown in Figure 10-153.




Figure 10-153 Selecting the Delete Consistency Group option

4. The warning window that is shown in Figure 10-154 opens. Click **Yes**.

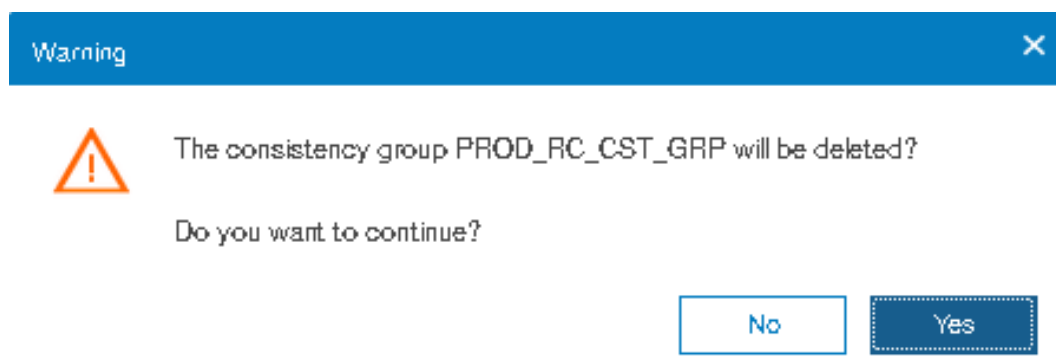


Figure 10-154 Confirmation message

10.11 Troubleshooting remote copy

Remote copy (Metro Mirror and Global Mirror) has two primary error codes that are displayed: 1920 or 1720. A 1920 is a congestion error. This error means that the source, the link between the source and target, or the target cannot keep up with the requested copy rate. A 1720 error is a heartbeat or system partnership communication error. This error often is more serious because failing communication between your system partners involves extended diagnostic time.

10.11.1 1920 error

A 1920 error (event ID 050010) can have several triggers, including the following probable causes:

- ▶ Primary Spectrum Virtualize system or SAN fabric problem (10%)
- ▶ Primary Spectrum Virtualize system or SAN fabric configuration (10%)
- ▶ Secondary Spectrum Virtualize system or SAN fabric problem (15%)
- ▶ Secondary Spectrum Virtualize system or SAN fabric configuration (25%)
- ▶ Intercluster link problem (15%)
- ▶ Intercluster link configuration (25%)

In practice, the most often overlooked cause is latency. Global Mirror has a round-trip-time tolerance limit of 80 or 250 milliseconds, depending on the firmware version and the hardware model. See Figure 10-82 on page 535. A message that is sent from your source Lenovo storage V series system to your target system and the accompanying acknowledgment must have a total time of 80 or 250 milliseconds round trip. In other words, it must have up to 40 or 125 milliseconds latency each way.

The primary component of your round-trip time is the physical distance between sites. For every 1000 kilometers (621.4 miles), you observe a 5-millisecond delay each way. This delay does not include the time that is added by equipment in the path. Every device adds a varying amount of time depending on the device, but a good rule is 25 microseconds for pure hardware devices.

For software-based functions (such as compression that is implemented in applications), the added delay tends to be much higher (usually in the millisecond plus range.) Next, we describe an example of a physical delay.

Company A has a production site that is 1900 kilometers (1180.6 miles) away from its recovery site. The network service provider uses a total of five devices to connect the two sites. In addition to those devices, Company A employs a SAN FC router at each site to provide Fibre Channel over IP (FCIP) to encapsulate the FC traffic between sites.

Now, there are seven devices, and 1900 kilometers (1180.6 miles) of distance delay. All the devices are adding 200 microseconds of delay each way. The distance adds 9.5 milliseconds each way, for a total of 19 milliseconds. Combined with the device latency, the delay is 19.4 milliseconds of physical latency minimum, which is under the 80-millisecond limit of Global Mirror until you realize that this number is the best case number.

The link quality and bandwidth play a large role. Your network provider likely ensures a latency maximum on your network link. Therefore, be sure to stay as far beneath the Global Mirror round-trip-time (RTT) limit as possible. You can easily double or triple the expected physical latency with a lower quality or lower bandwidth network link. Then, you are within the range of exceeding the limit if high I/O occurs that exceeds the existing bandwidth capacity.

When you get a 1920 event, always check the latency first. If the FCIP routing layer is not properly configured, it can introduce latency. If your network provider reports a much lower latency, you might have a problem at your FCIP routing layer. Most FCIP routing devices have built-in tools to enable you to check the RTT. When you are checking latency, remember that TCP/IP routing devices (including FCIP routers) report RTT using standard 64-byte ping packets.

In Figure 10-155 on page 592, you can see why the effective transit time must be measured only by using packets that are large enough to hold an FC frame, or 2148 bytes (2112 bytes of payload and 36 bytes of header). Allow estimated resource requirements to be a safe amount, because various switch vendors have optional features that might increase this size. After you verify your latency by using the proper packet size, proceed with normal hardware troubleshooting.

Before we proceed, we look at the second largest component of your RTT, which is *serialization delay*. Serialization delay is the amount of time that is required to move a packet of data of a specific size across a network link of a certain bandwidth. The required time to move a specific amount of data decreases as the data transmission rate increases.

Figure 10-155 on page 592 shows the orders of magnitude of difference between the link bandwidths. It is easy to see how 1920 errors can arise when your bandwidth is insufficient. Never use a TCP/IP ping to measure RTT for FCIP traffic.

Packet Size	Link Size	Serialization Delay (Time Required to Send Data)	Unit
64	256 Kbps	2.0E+03	microseconds
64	1.5 Mbps	3.4E+02	microseconds
64	100 Mbps	5.1E+00	microseconds
64	155 Mbps	3.3E+00	microseconds
64	622 Mbps	8.2E-01	microseconds
64	1 Gbps	5.1E-04	microseconds
64	10 Gbps	5.1E-05	microseconds
1500	256 Kbps	4.7E+04	microseconds
1500	1.5 Mbps	8.0E+03	microseconds
1500	100 Mbps	1.2E+02	microseconds
1500	155 Mbps	7.7E+01	microseconds
1500	622 Mbps	1.9E+01	microseconds
1500	1 Gbps	1.2E+01	microseconds
1500	10 Gbps	1.2E+00	microseconds
2148	256 Kbps	6.7E+04	microseconds
2148	1.5 Mbps	1.1E+04	microseconds
2148	100 Mbps	1.7E+02	microseconds
2148	155 Mbps	1.1E+02	microseconds
2148	622 Mbps	2.8E+01	microseconds
2148	1 Gbps	1.7E+01	microseconds
2148	10 Gbps	1.7E-03	microseconds

Figure 10-155 Effect of packet size (in bytes) versus the link size

In Figure 10-155, the amount of time in microseconds that is required to transmit a packet across network links of varying bandwidth capacity is compared. The following packet sizes are used:

- ▶ 64 bytes: The size of the common ping packet
- ▶ 1500 bytes: The size of the standard TCP/IP packet
- ▶ 2148 bytes: The size of an FC frame

Finally, your path maximum transmission unit (MTU) affects the delay that is incurred to get a packet from one location to another location. An MTU might cause fragmentation or be too large and cause too many retransmits when a packet is lost.

10.11.2 1720 error

The 1720 error (event ID 050020) is the other problem remote copy might encounter. The amount of bandwidth that is needed for system-to-system communications varies based on the number of nodes. It is important that it is not zero. When a partner on either side stops communication, you see a 1720 appear in your error log. According to the product documentation, there are no likely field-replaceable unit breakages or other causes.

The source of this error is most often a fabric problem or a problem in the network path between your partners. When you receive this error, check your fabric configuration for zoning of more than one host bus adapter (HBA) port for each node per I/O Group if your fabric has more than 64 HBA ports zoned. One port for each node per I/O Group per fabric that is associated with the host is the suggested zoning configuration for fabrics.

For those fabrics with 64 or more host ports, this recommendation becomes a rule. Therefore, you see four paths to each volume discovered on the host because each host needs to have at least two FC ports from separate HBA cards, each in a separate fabric. On each fabric, each host FC port is zoned to two of node ports where each port comes from one node canister. This gives four paths per host volume. More than four paths per volume are supported but not recommended.

Improper zoning can lead to SAN congestion, which can inhibit remote link communication intermittently. Checking the zero buffer credit timer from IBM Virtual Storage Center and comparing against your sample interval reveals potential SAN congestion. If a zero buffer credit timer is above 2% of the total time of the sample interval, it might cause problems.

Next, always ask your network provider to check the status of the link. If the link is acceptable, watch for repeats of this error. It is possible in a normal and functional network setup to have occasional 1720 errors, but multiple occurrences could indicate a larger problem.

If you receive multiple 1720 errors, recheck your network connection and then check the system partnership information to verify its status and settings. Then, proceed to perform diagnostics for every piece of equipment in the path between the two Lenovo storage V series systems. It often helps to have a diagram that shows the path of your replication from both logical and physical configuration viewpoints.

If your investigations fail to resolve your remote copy problems, contact your Lenovo Support representative for a more complete analysis.

10.12 HyperSwap

The *HyperSwap* high availability function allows business continuity in a hardware failure, power failure, connectivity failure, or disasters, such as fire or flooding. It is available on the IBM Storwize V7000 for Lenovo and Lenovo Storage V5030.

The HyperSwap function provides highly available volumes that are accessible through two sites at up to 300 km (186.4 miles) apart. A fully independent copy of the data is maintained at each site. When data is written by hosts at either site, both copies are synchronously updated before the write operation is completed. The HyperSwap function automatically optimizes itself to minimize data that is transmitted between sites and to minimize host read and write latency.

If the nodes go offline or the storage at either site goes offline, leaving an online and accessible up-to-date copy, the HyperSwap function can automatically fail over access to the online copy. The HyperSwap function also automatically resynchronizes the two copies when possible.

HyperSwap capability enables each volume to be presented by two I/O groups. The configuration tolerates combinations of node and site failures, by using the same flexible choices of host multipathing driver interoperability that are available for the Lenovo storage V-series systems. The use of FlashCopy helps maintain a golden image during automatic resynchronization.

Important: Because Remote Mirroring is used to support the HyperSwap capability, Remote Mirroring licensing is a requirement for using HyperSwap.

The HyperSwap function uses a hyperswap topology by spreading the control enclosure of the system across two sites, with storage at a third site that acts as a tie-breaking quorum device:

- ▶ The HyperSwap topology requires at least one control enclosure in each main data site. Therefore, to get a volume that is resiliently stored on both sites, one control enclosure with the necessary storage capacity is required.
- ▶ The HyperSwap topology uses additional system resources to support a fully independent cache on each site, providing full performance even if one site is lost.
- ▶ The HyperSwap function can be configured by using the GUI or command-line interface (CLI).
- ▶ The hosts, Lenovo Storage V5030 control enclosures and storage enclosures are in one of two failure domains or sites. External virtualized storage capacity can also be used.
- ▶ Volumes are visible as a single object across both sites (the Lenovo Storage V5030 control enclosure).

At least two control enclosures are required for HyperSwap. System scalability depends on the hardware details, as shown in Table 10-12.

Table 10-12 HyperSwap support

	Lenovo V3700 V2	Lenovo V3700 V2 XP	Lenovo Storage V5030	IBM V7000 for Lenovo
Maximum number of I/O groups	1	1	2	4
Support for HyperSwap	No	No	Yes	Yes
Support for Stretched Cluster	No	No	No	No

A V5030 HyperSwap cluster is always restricted to a single control enclosure per site. The IBM Storwize V7000 for Lenovo and SVC can provide more scalability and offer more flexibility.

10.12.1 Introduction to HyperSwap volumes

The HyperSwap function is built on the Remote Copy features that include Metro Mirror, Global Mirror, and Global Mirror with Change Volumes.

The HyperSwap function works with the standard multipathing drivers that are available on various host types. No additional host support is required to access the highly available volume. Where multipathing drivers support Asymmetric Logical Unit Access (ALUA), the storage system informs the multipathing driver about the nodes that are in the same site and the nodes that need to be used to minimize I/O latency.

The host and Lenovo Storage V5030 site attributes must be configured to enable this optimization and to enable HyperSwap functionality. A three-site setup is required. Two sites are used as the main data center to provide two independent data copies. A quorum disk or an IP-based quorum can be used as a quorum device. However, the quorum device must be placed in a third, independent site.

The quorum disk must be supported as an “extended quorum device”. The connection can be implemented by using Fibre Channel, Fibre Channel through wavelength-division multiplexing (WDM), synchronous digital hierarchy (SDH) and synchronous optical network (SONET), or FCIP. The minimum bandwidth is 2 MBps.

The IP quorum substitutes the active quorum disk’s tiebreaker role. Redundancy can be implemented by using multiple quorum apps, similar to multiple quorum disks. However, only one app is active at a time. The other apps are available if the active quorum device app fails. For more information about quorum devices, see the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 Lenovo Information Center:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/tbrd_clstrcli_4892pz.html

Because HyperSwap is running as a single cluster that is distributed across two main data centers, one Lenovo Storage V5030 control enclosure is required in each site. Both control enclosures must be added to the same cluster. Only the Lenovo Storage V5030 supports the clustering of two control enclosures, so two Lenovo Storage V5030 control enclosures are required for HyperSwap. Metro Mirror is used to keep both data copies in sync.

The host accesses both I/O groups, as shown in Figure 10-156 on page 596. The original Metro Mirror target ID is not used for host access. Instead, HyperSwap presents the Metro Mirror source ID for the target volume to the host. From the host perspective, the same volume is available on both I/O groups, although the Lenovo Storage V5030 volumes are connected through Metro Mirror.




Figure 10-156 HyperSwap

A site attribute must be set for any host, Lenovo Storage V5030 storage systems, and external virtualized storage system. The host uses the local I/O group (same site attribute) for data access, as shown in Figure 10-157.

The continuous blue line shows the host default access path to the volume at the same site. The dotted blue line shows the non-preferred access path that is used if the preferred access path is not available. Accessing both I/O groups doubles the number of paths from host to volume. Take note of the limited number of supported paths for your multipath device driver and limit the number of paths to an acceptable level.




Figure 10-157 Data access

Data flow

The host reads and writes data to the local I/O group within the same site. The HyperSwap system sends the data to the remote site by using internal Metro Mirror, as shown in Figure 10-158 on page 598.




Figure 10-158 Data flow for a single volume

If a host accesses the volume on the Metro Mirror target site, all read and write requests can be forwarded to the I/O group that acts as the Metro Mirror source volume, as shown in Figure 10-159 on page 599. All host-related traffic must be handled from the remote I/O group, which increases the long-distance data traffic.




Figure 10-159 Data flow from the “wrong” site

Access to the Metro Mirror target volume is measured and HyperSwap triggers a switch of the Metro Mirror copy direction, if the workload on the “wrong” site is significantly higher than the workload on the Metro Mirror source site over any length of time, as shown in Figure 10-160 on page 600. This copy direction switch reduces the additional host-related long-distance traffic and provides better response time.




Figure 10-160 Data flow: Switch copy direction

The duration until the copy direction switches depends on the load distribution on the volume in both sites. Although a HyperSwap volume can be accessed concurrently for read and write I/O from any host in any site, all I/O is forwarded to one I/O group in one site. Usage of the wrong site increases the long-distance traffic. The HyperSwap cluster monitors the workload and it can switch the copy direction if the highest workload is arriving on the wrong site.

Applications with an equal workload pattern to the same volume by using both I/O groups (Oracle Real Application Clusters (RAC) and VMware vMotion) are not optimal for local HyperSwap.

10.12.2 Failure scenarios

If one node fails, the other node in the same I/O group takes over the responsibility for all volumes that are owned by the affected I/O group, as shown in Figure 10-161 on page 601. The system can deactivate the write cache, which might influence the overall performance. The multipath driver can switch the active paths to the named node. The Metro Mirror relationship continues to operate and provides a synchronous, consistent copy at the remote I/O group.

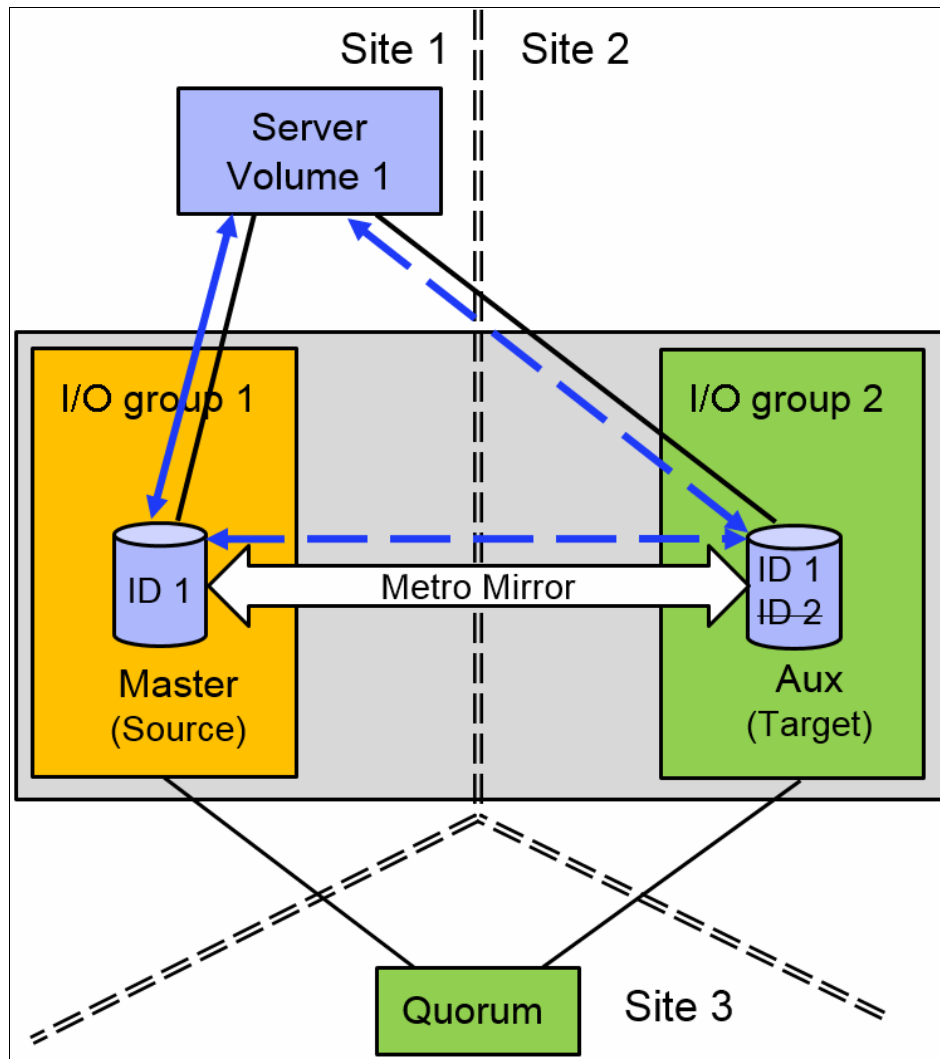


Figure 10-161 Single node failure in an I/O group

If an I/O group fails, the host can use the second I/O group at the remote site, as shown in Figure 10-162 on page 602. The remote I/O group handles all volume-related traffic, but HyperSwap cannot keep both copies in sync anymore because of an inactive I/O group.




Figure 10-162 I/O group failure

As soon as the failed I/O group is back, the system can automatically resynchronize both copies in the background. Before the resynchronization, the Lenovo Storage V5030 can perform a FlashCopy on HyperSwap source and target volumes, as shown in Figure 10-163 on page 603. Each change volume requires two FlashCopy relationships, one relationship in each direction. So, four FlashCopy relationships are required for each HyperSwap volume.

To provide an easy to use GUI, those relationships and FlashCopy volumes are not shown in the GUI. They are only visible and manageable by using the CLI.




Figure 10-163 Resynchronization

After successful resynchronization, the host switches automatically to the I/O group at the same site for the best performance and limited inter-switch link (ISL) usage, as shown in Figure 10-164 on page 604.




Figure 10-164 Data flow after resynchronization

HyperSwap uses Metro Mirror technology, which enables the usage of Metro Mirror consistency groups that are described in “Remote Copy Consistency Groups” on page 526.

10.12.3 Current HyperSwap limitations

HyperSwap has the following limitations:

- ▶ Cluster internal Metro Mirror is used for replication, so the size of a HyperSwap volume cannot be changed by using `expandvdisksize` and `shrinkvdisksize`.
- ▶ A cascaded Remote Copy is not available. HyperSwap volumes cannot be replicated to a second, independent storage system by using Remote Copy functionality.

Number of HyperSwap volumes supported per Lenovo Storage V5030 system is 1024. Additional FlashCopy requirements can reduce the number of possible HyperSwap volumes.

- ▶ Hosts that access a HyperSwap volume through Internet Small Computer System Interface (iSCSI) or serial-attached SCSI (SAS) cannot take advantage of the high availability function.
- ▶ FlashCopy usage can be complicated because the Metro Mirror source volume and target volume can switch during daily operation. Because of this possibility, the identification of the copy direction is required for a successful FlashCopy.
- ▶ The Remote Copy relationship must be removed first for a reverse FlashCopy operation. After a reverse FlashCopy, all HyperSwap functions must be implemented manually again (Remote Mirror + FlashCopy relationships).
- ▶ IBM FlashCopy Manager is not supported by HyperSwap volumes.

External storage virtualization

This chapter describes how to incorporate external storage systems into the virtualized world of the Lenovo Storage V5030. A key feature of the Lenovo Storage V5030 is its ability to consolidate disk controllers from various vendors into storage pools. By virtualizing vendors disk controllers, the storage administrator can manage and provision storage to applications from a single user interface and use a common set of advanced functions across all of the storage systems under the control of the Lenovo Storage V5030.

A distinction must be made between virtualizing external storage and importing existing data into the Lenovo Storage V5030. Virtualizing external storage means the creation of logical units with no data on them and the addition of these logical units to storage pools under the Lenovo Storage V5030 control. In this way, the external storage can benefit from the Lenovo Storage V5030 features, such as Easy Tier and Copy Services.

When existing data needs to be put under the control of the Lenovo Storage V5030, it must first be imported as an *image mode volume*. It is strongly recommended to copy the existing data onto internal or external storage that is under the control of the Lenovo Storage V5030 instead of letting the data within an image mode volume, so the data can benefit from the Lenovo Storage V5030 features.

Note: External storage virtualization is available on the Lenovo Storage V5030 model only. It is not available on the Lenovo V3700 V2 or Lenovo Storage V3700 V2 XP. However, these models can still import data from external storage systems. For more information about storage migration, see Chapter 7, “Storage migration” on page 323.

Specifically, this chapter provides information about the following topics:

- ▶ 11.1, “Planning for external storage virtualization” on page 608
- ▶ 11.2, “Working with external storage” on page 611

11.1 Planning for external storage virtualization

This section describes how to plan for virtualizing external storage with the Lenovo Storage V5030. Virtualizing the storage infrastructure with the Lenovo Storage V5030 makes your storage environment more flexible, cost-effective and easy to manage. The combination of the Lenovo Storage V5030 and an external storage system allows more storage capacity benefits from the powerful software functions within the Lenovo Storage V5030.

The external storage systems that are incorporated into the Lenovo Storage V5030 environment can be new systems or existing systems. Any data on the existing storage systems can be easily migrated to an environment that is managed by the Lenovo Storage V5030, as described in Chapter 7, “Storage migration” on page 323.

11.1.1 License for external storage virtualization

From a licensing standpoint, when external storage systems are virtualized by the Lenovo Storage V5030, a per-enclosure External Virtualization license is required.

Migration: If the Lenovo Storage V5030 is used as a general management tool, you must order the correct External Virtualization licenses. The only exception is if you want to migrate existing data from external storage systems to Lenovo Storage V5030 internal storage and then remove the external storage. You can temporarily configure your External Storage license for a 45-day period. For more than a 45-day migration requirement, the correct External Virtualization license must be ordered.

You can configure the Lenovo Storage V5030 licenses by clicking the **Settings** icon and then **System** → **Licensed Functions**. For more information about setting licenses on the IBM Storwize V5030, see Chapter 2, “Initial configuration” on page 35.

For assistance with licensing questions or to purchase any of these licenses, contact your IBM account team or IBM Business Partner.

11.1.2 SAN configuration planning

External virtualization is only supported by using Fibre Channel or Fibre Channel over Ethernet (FCoE). Therefore, it is a prerequisite to install a pair of the optional 16 Gb Fibre Channel adapter cards (or a pair of the optional 10 GbE adapter cards if you use FCoE).

External storage controllers that are virtualized by the Lenovo Storage V5030 must be connected through storage area network (SAN) switches. A direct connection between the Lenovo Storage V5030 and external storage controllers is not supported.

Ensure that the switches or directors are at the firmware levels that are supported by the Lenovo Storage V5030 and that the port login maximums that are listed in the restriction document are not exceeded. The configuration restrictions are listed on the IBM Support home page, which is available at this web page:

<https://ibm.biz/BdjGMJ>

The suggested SAN configuration is based in a dual fabric solution. The ports on external storage systems and the Lenovo Storage V5030 ports must be evenly split between the two fabrics to provide redundancy if one of the fabrics goes offline.

After the Lenovo Storage V5030 and external storage systems are connected to the SAN fabrics, zoning on the switches must be configured. In each fabric, create a zone with the four Lenovo Storage V5030 worldwide port names (WWPNs), two from each node canister with up to a maximum of eight WWPNs from each external storage system.

Ports: The Lenovo Storage V5030 supports a maximum of 16 ports or WWPNs from an externally virtualized storage system.

Figure 11-1 shows an example of how to cable devices to the SAN. Refer to this example as we describe the zoning. For this example, we used an IBM Storwize V3700 for Lenovo as our external storage.

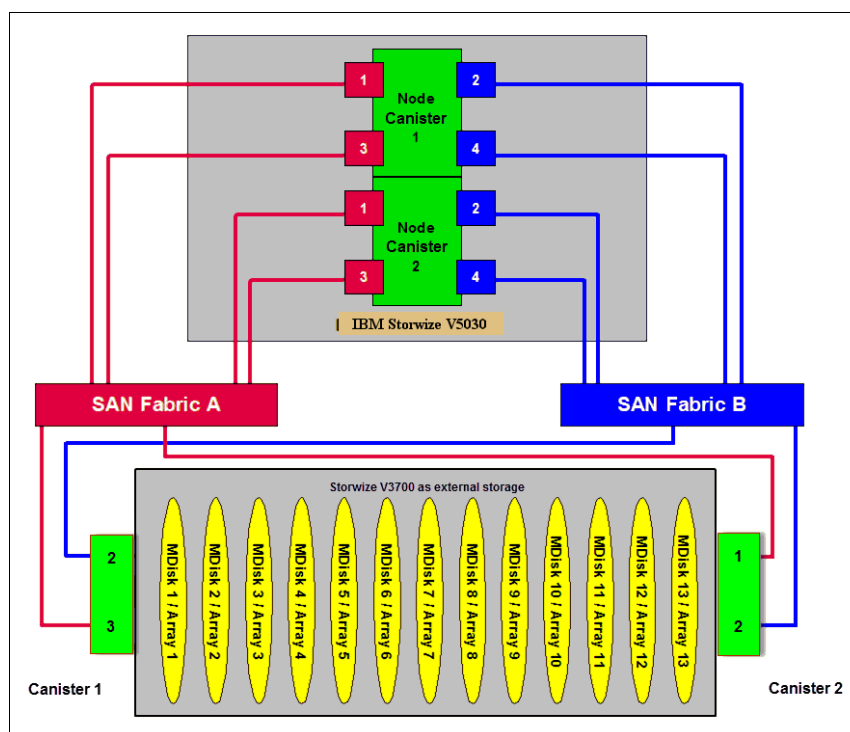


Figure 11-1 SAN cabling and zoning example

Create an Lenovo Storage V5030 external storage zone for each storage system to be virtualized, as shown in the following examples:

- ▶ Zone the external IBM Storwize V3700 for Lenovo canister 1 port 2 with the Lenovo Storage V5030 canister 1 port 2 and canister 2 port 2 in the blue fabric.
- ▶ Zone the external IBM Storwize V3700 for Lenovo canister 2 port 2 with the Lenovo Storage V5030 canister 1 port 4 and canister 2 port 4 in the blue fabric.
- ▶ Zone the external IBM Storwize V3700 for Lenovo canister 1 port 3 with the Lenovo Storage V5030 canister 1 port 1 and canister 2 port 1 in the red fabric.
- ▶ Zone the external IBM Storwize V3700 for Lenovo canister 2 port 1 with the Lenovo Storage V5030 canister 1 port 3 and canister 2 port 3 in the red fabric.

11.1.3 External storage configuration planning

Logical units that are created on the external storage system must provide redundancy through various RAID levels, preventing a single physical disk failure from causing a managed disk (MDisk), storage pool, or associated host volume from getting offline. To minimize the risk of data loss, virtualize storage systems only where logical unit numbers (LUNs) are configured by using a RAID level other than RAID 0 (RAID 1, RAID 10, RAID 0+1, RAID 5, RAID 6, Distributed RAID 5, or Distributed RAID 6).

Verify that the storage controllers to be virtualized by the Lenovo Storage V5030 meet the configuration restrictions, which are listed on the IBM Support home page, at this web page:

<https://ibm.biz/BdjGMJ>

Ensure that the firmware or microcode levels of the storage controllers to be virtualized are supported by the Lenovo Storage V5030. See the Interoperability matrix web page for more details:

<https://datacentersupport.lenovo.com/tw/en/products/storage/lenovo-storage/v5030/6536/documentation>

The Lenovo Storage V5030 must have exclusive access to the LUNs from the external storage system that are presented to it. LUNs cannot be shared between the Lenovo Storage V5030 and other storage virtualization platforms or between an Lenovo Storage V5030 and hosts. However, different LUNs can be mapped from the same external storage system to an Lenovo Storage V5030 and other hosts in the SAN through different storage ports.

Ensure that the external storage subsystem LUN masking is configured to map all LUNs to all of the WWPNs in the Lenovo Storage V5030 storage system.

Ensure that you check the Lenovo Information Center and review the “Configuring and servicing storage system” topic before you prepare the external storage systems for discovery from the Lenovo Storage V5030 system. This Knowledge Center topic is at this web page:

<https://ibm.biz/BdjGMJ>

11.1.4 Guidelines for virtualizing external storage

When external storage is virtualized by using the Lenovo Storage V5030, the following guidelines must be followed:

- ▶ Avoid splitting arrays into multiple LUNs at the external storage system level. When possible, create a single LUN per array for mapping to the Lenovo Storage V5030.
- ▶ Use 6 - 8 disks per RAID group when you create the external LUNs, more than that and it may result in a longer rebuild time in case of a single disk failure, affecting the performance of the LUN and exposing it to complete failure if a second disk fails during the rebuild. Additionally, the smaller the number of disks, the more likely it is that write operations span an entire stripe (stripe size, which is multiplied by the number of members, minus one). In this case, write performance is improved.
- ▶ Except for Easy Tier, do not mix MDisk that vary in performance or reliability in the same storage pool. Put only MDisk of the same size and performance into the same storage pool. Likewise, group MDisk from different arrays into different pools. For more information about Easy Tier, see Chapter 9, “Advanced features for storage efficiency” on page 403.

- Do not leave volumes in image mode. Use image mode only to import or export existing data into or out of the Lenovo Storage V5030. Migrate data from image mode volumes and associated MDisks to other storage pools to benefit from storage virtualization and the enhanced benefits of the Lenovo Storage V5030, such as Easy Tier.

11.2 Working with external storage

This section describes how to manage external storage by using a Lenovo Storage V5030.

The basic concepts of managing an external storage system are the same as the concepts for managing internal storage. The Lenovo Storage V5030 discovers LUNs from the external storage system as one or more MDisks. These MDisks are added to a storage pool in which volumes are created and mapped to hosts, as needed.

11.2.1 Adding external storage

To add external storage systems to the Lenovo Storage V5030 virtualized environment, complete the following steps:

1. Zone a minimum of two and a maximum of 16 Fibre Channel ports from the external storage system with all eight Fibre Channel ports on the Lenovo Storage V5030 system. For more information about zoning, see 11.1.2, “SAN configuration planning” on page 608. Because the Lenovo Storage V5030 is virtualizing your storage, hosts need to be zoned with the Lenovo Storage V5030 controller’s WWPNs.
2. By using the storage partitioning or LUN masking feature of the external storage system, create a group that includes all eight Lenovo Storage V5030 WWPNs.
3. Create equal size arrays on the external system by using any RAID level except zero.
4. Create a single LUN per RAID array.

- Map the LUNs to all eight Fibre Channel ports on the Lenovo Storage V5030 system by assigning them to the group that was created in step 2 on page 611.
- Verify that the Lenovo Storage V5030 discovered the LUNs as unmanaged MDisks. To get to the external storage pool panel, select **External Storage** from the Pools menu as shown in Figure 11-2.

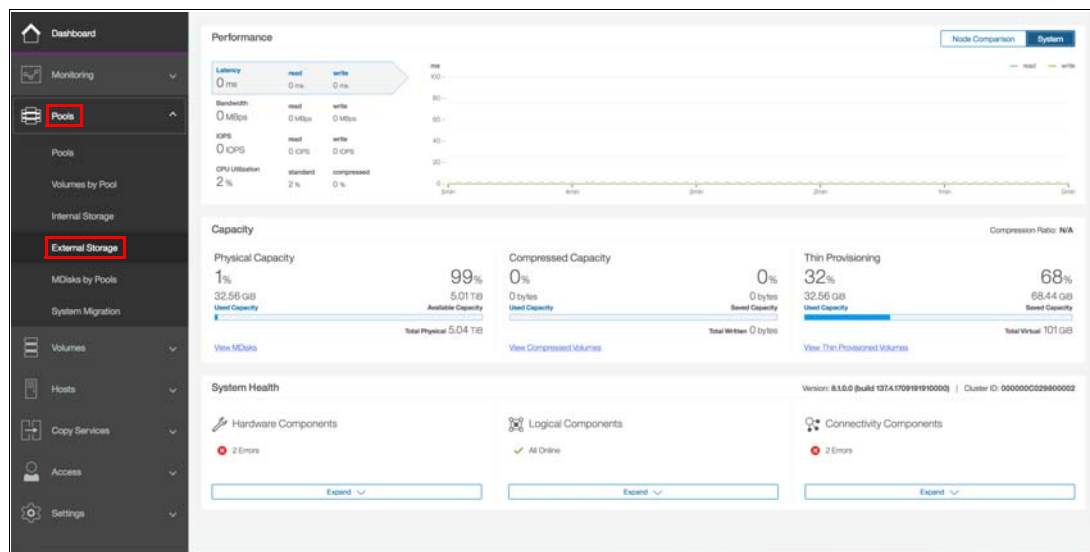


Figure 11-2 Selecting the external storage

If the external storage does not show up automatically, click **Discover storage** from the Actions menu on the External Storage panel, as shown in Figure 11-3.

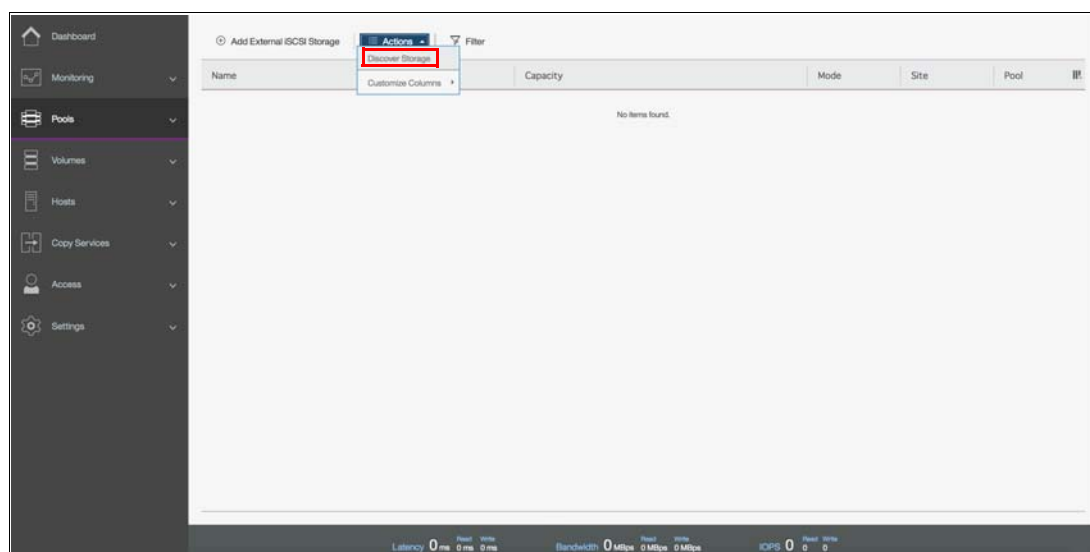


Figure 11-3 Discover storage

- The MDisks are unassigned and need to be assigned to the correct storage tiers. It is important to set the tiers correctly if you plan to use the Easy Tier feature. For more information about storage tiers, see Chapter 9, “Advanced features for storage efficiency” on page 403.
- Select the MDisks to assign and either use the Actions drop-down menu or right-click and select **Modify Tier**, as shown in Figure 11-4 on page 613.

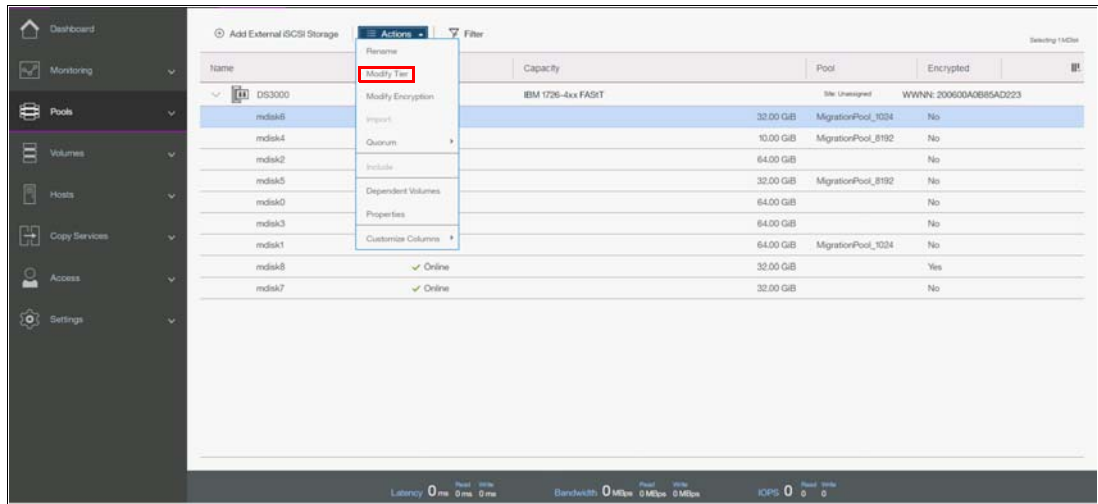


Figure 11-4 Modify Tier option

9. Ensure that the correct MDisk tier is chosen, as shown in Figure 11-5. Click **Modify** to change the tier setting.

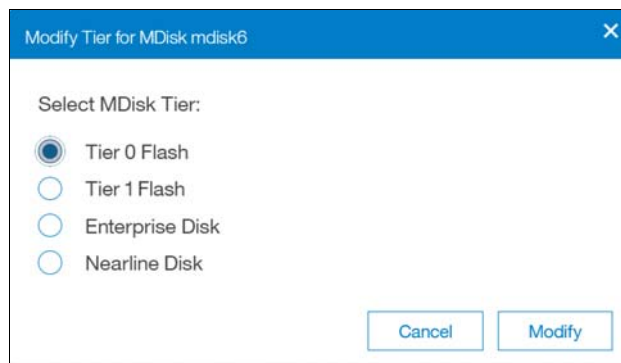


Figure 11-5 Choosing a storage tier

10. After the tier is assigned, add the MDisks to an existing pool or create a new pool to include them. Figure 11-6 shows how to add selected MDisks to an existing storage pool. Click **Assign** under the Actions menu.

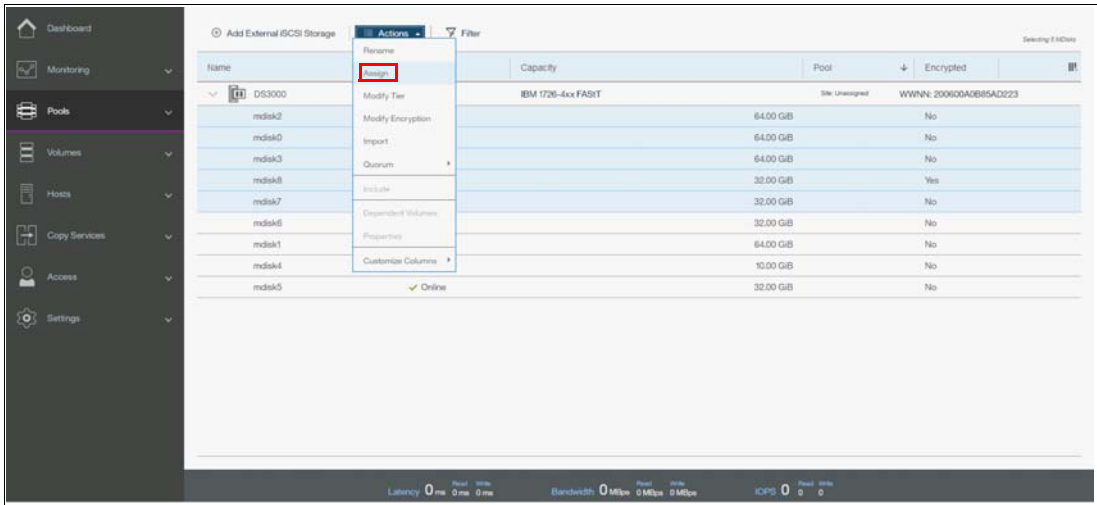


Figure 11-6 Adding MDisks to a pool

- If the storage pool does not exist, follow the procedure that is outlined in Chapter 4, “Storage pools” on page 139.
11. Add the MDisks to the pool. Select the pool to which the MDisks are going to be assigned and click **Assign**, as shown in Figure 11-7. After the task completes, click **Close**.

Important: If the external storage volumes to virtualize behind the Lenovo Storage V5030 contain data and this data needs to be retained, *do not* use the “Assign to pool” option to manage the MDisks. *This option can destroy the data on the disks.* Instead, use the Import option. For more information, see 11.2.2, “Importing image mode volumes” on page 615.

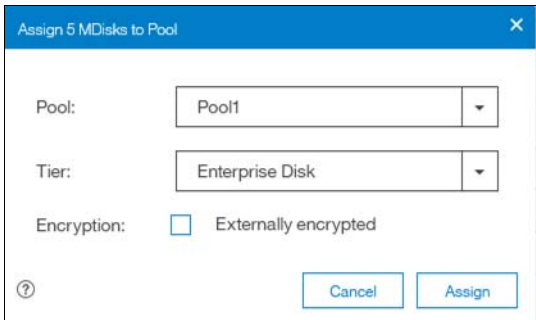


Figure 11-7 Selecting the storage pool to assign the MDisks to the pool

12. The external MDisks that are assigned to a pool within Lenovo Storage V5030 are displayed under the MDisks by Pools panel as shown in Figure 11-8 on page 615. Create volumes from the storage pool and map them to hosts, as needed. See Chapter 6, “Volume configuration” on page 269 to learn how to create and map volumes to hosts.




Figure 11-8 External MDisk displayed on MDisk by Pools panel

11.2.2 Importing image mode volumes

If the external storage systems are not new systems and data exists on the LUNs that must be kept after virtualization, you must import the existing LUNs. The process of importing existing data on external volumes is simplified by using the storage migration wizard, which is described in Chapter 7, “Storage migration” on page 323.

To manually import volumes, they must not be assigned to a storage pool and they must be unmanaged managed disks (MDisks). Hosts that access data from these external storage system LUNs can continue to access data, but the hosts must be rezoned and mapped to the Lenovo Storage V5030 to use these external storage system LUNs after they are presented through the Lenovo Storage V5030.

Figure 11-9 shows how to import an unmanaged MDisk. Select the unmanaged MDisk and click **Import** from the Actions drop-down menu. Multiple MDisks can be selected by using the Ctrl key.

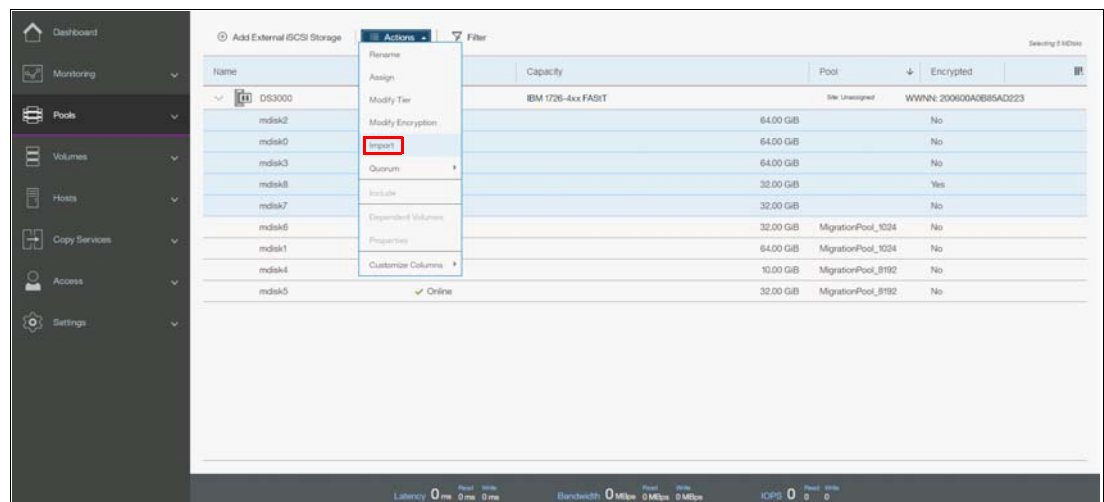


Figure 11-9 Import MDisk option

Selecting the Import option opens a new panel that requires additional volume information, as shown in Figure 11-10.

Import 5 MDisks As Volumes

Volume names:

A_DS3000_000000000000

B_DS3000_000000000000

C_DS3000_000000000000

D_DS3000_000000000000

E_DS3000_000000000000

Import method:

☒ Import to temporary pool as image-mode volume

☐ Migrate to an existing pool

Extent Size:

1.00 GiB

☐ Copy Services on the external storage system are used with this volume

?

Cancel Import

Figure 11-10 Import MDisks as Volumes panel

You can change the default volume names to more meaningful names by editing the Volume names text boxes.

You can choose between importing the volume to a temporary pool as an *image mode volume*, which the Lenovo Storage V5030 can create and name for you, or migrating the volume to an existing pool.

An image mode volume has a direct block-for-block translation from the imported MDisk and the external LUN. Therefore, the existing data is preserved. In this state, the Lenovo Storage V5030 is acting as a proxy and the image mode volume is simply a “pointer” to the existing external LUN. Because of the way that virtualization works on the Lenovo Storage V5030, the external LUN is presented as an MDisk, but we cannot map an MDisk directly to a host. Therefore, the Spectrum Virtualize software must create the image mode volume, to allow hosts to perform the mapping through the Lenovo Storage V5030.

If you choose a temporary pool, you must first select the extent size for the pool. The default value for extents is 1 GB. If you plan to migrate this volume to another pool later, ensure that the extent size matches the extent size of the prospective target pool. For more information about extent sizes, see Chapter 4, “Storage pools” on page 139.

If an existing storage pool is chosen, the Lenovo Storage V5030 can perform a migration task. The external LUN can be imported into a temporary migration pool and a migration task can run in the background to copy data to MDisks that are in the target storage pool. At the end of the migration, the external LUN and its associated MDisk can be in the temporary pool and show as managed, but they can be removed from the Lenovo Storage V5030.

Figure 11-11 on page 617 shows how to select an existing pool for volume migration. The pool must have enough available capacity to store the volumes that are being imported.

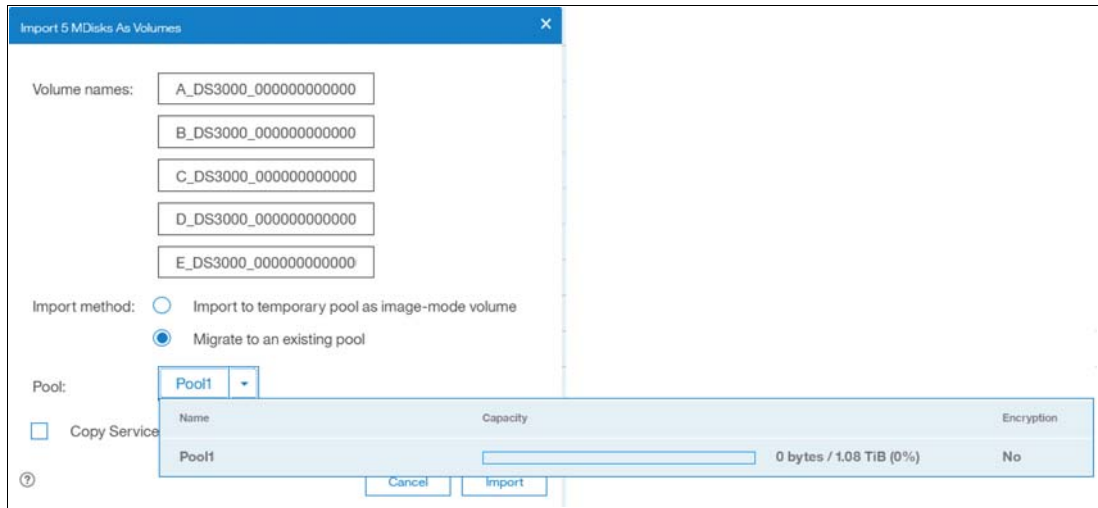


Figure 11-11 Import volumes into an existing pool

Select **Copy Services** if copy services (replication functionality) are used on the external storage system that hosts the LUN. Click **Import** to confirm your selections and to start the import process.

Note: Only pools with sufficient capacity are shown because the import of an MDisk to an existing storage pool can migrate storage. This storage can be migrated only if sufficient capacity exists in the target pool to create a copy of the data on its own MDisk. The external MDisk can be imported as an image mode volume into a temporary migration pool and a volume migration can take place in the background to create the volume in the target pool.

A migration task starts and can be tracked through the **System Migration** panel within the **Pools** menu, as shown in Figure 11-12. The actual data migration begins after the MDisk is imported successfully.

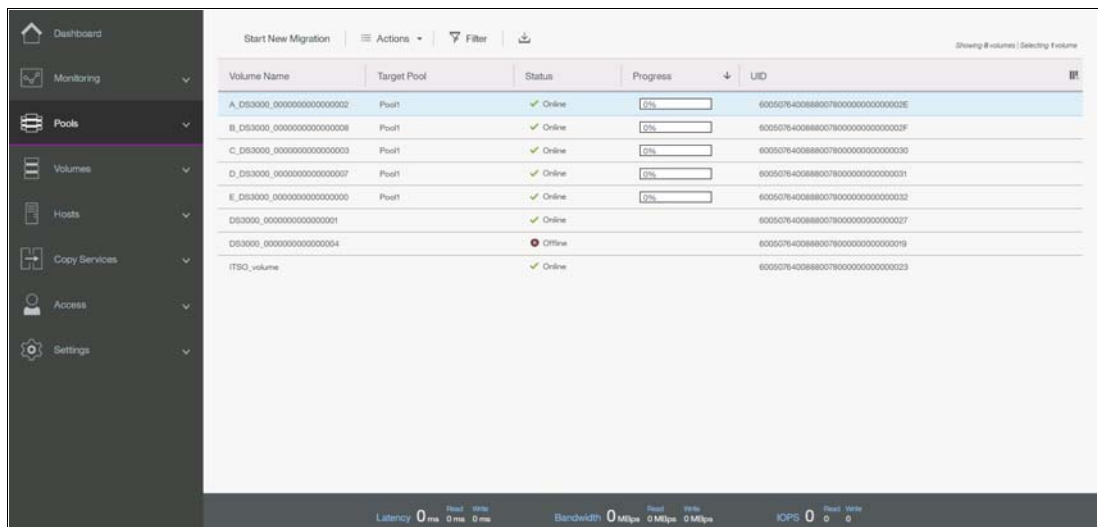


Figure 11-12 Checking the migration status

When the migration completes, the migration status disappears and the volume is displayed in the target pool, as shown in Figure 11-13 on page 618.

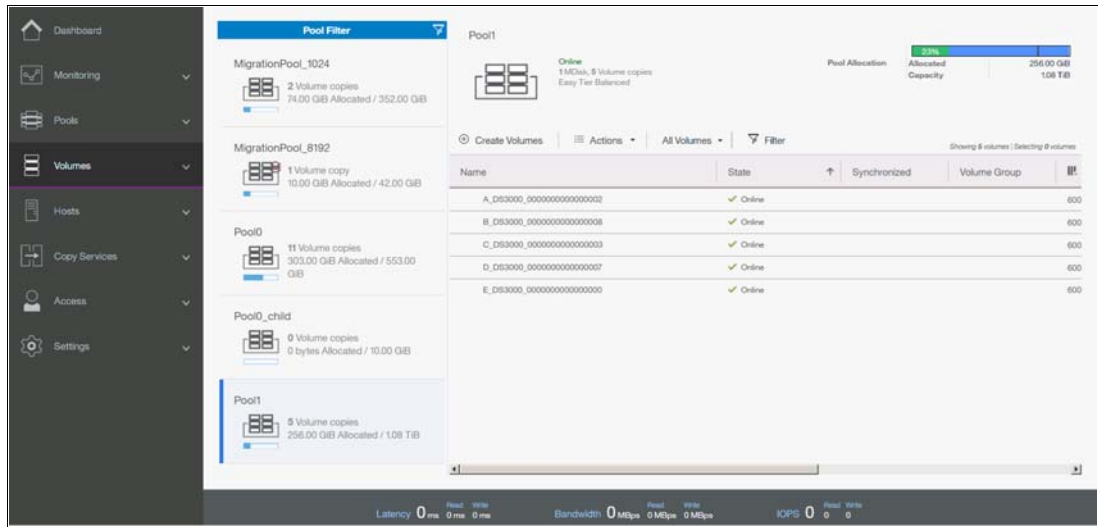


Figure 11-13 Volume is assigned to the target pool

After the migration completes, the image mode volume is automatically deleted, but the external LUN exists as a managed MDisk in the temporary storage pool. It is unassigned from the pool and listed as an unassigned MDisk. Later, you can retire the external LUN and remove it completely from the Lenovo Storage V5030 by unmapping the volume at the external storage and by clicking **Detect MDisks** on the Lenovo Storage V5030. For more information about removing external storage, see 11.2.4, “Removing external storage” on page 623.

If you choose to import a volume as an image mode volume, the external LUN appears as an MDisk with an associated image mode volume name and can be listed as shown in Figure 11-14.

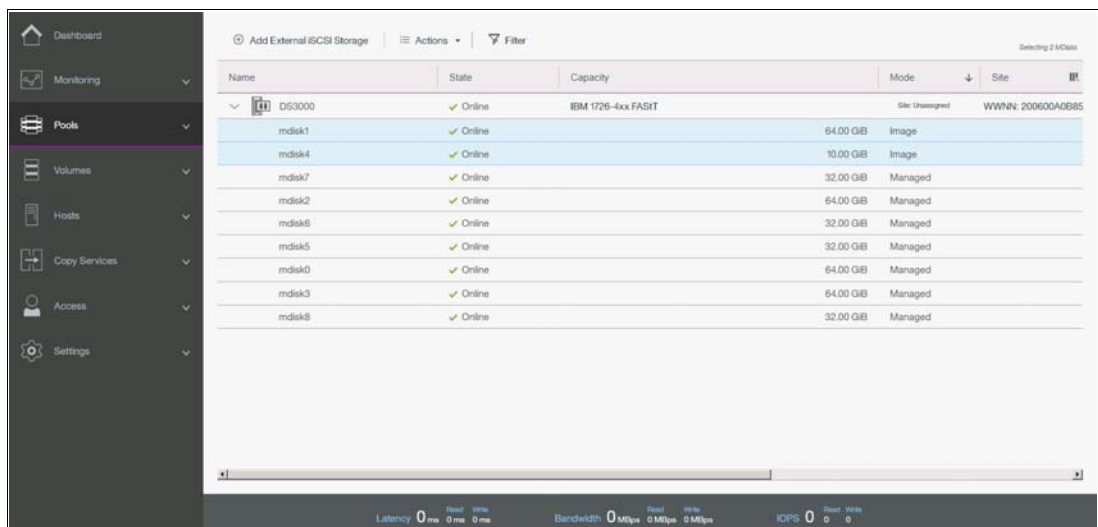


Figure 11-14 Image mode volumes

The volume is also listed in the **System Migration** panel because the Lenovo Storage V5030 expects you to migrate these volumes later, as shown in Figure 11-15.

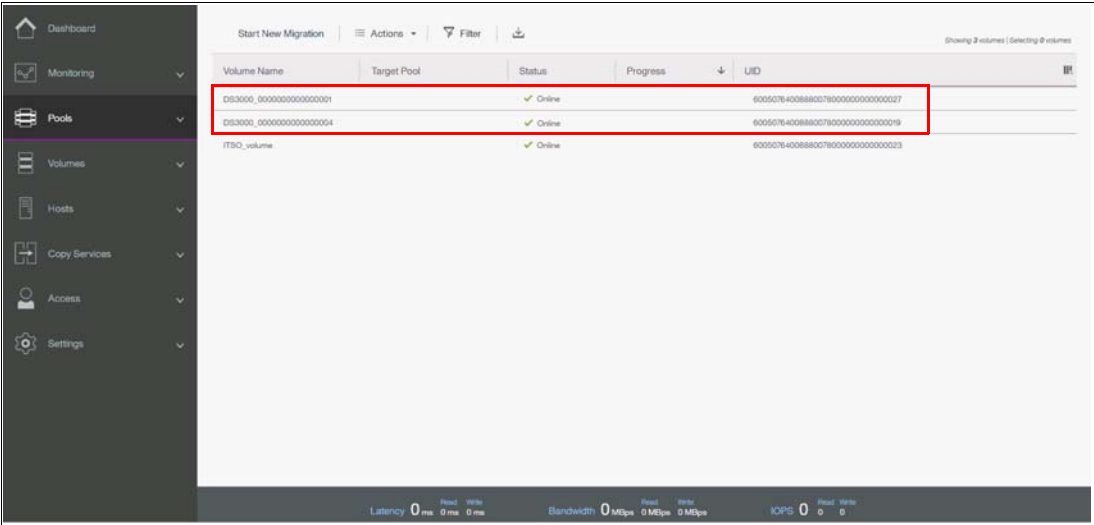


Figure 11-15 Migrations are available for image mode volumes

At the end of this process, the volume can be mapped to a host.

11.2.3 Managing external storage

The Lenovo Storage V5030 provides an individual external panel for managing external storage systems.

You can access the external panel by clicking **Pools** → **External Storage**, as shown in Figure 11-2 on page 612. Extended help information for external storage is available by clicking the help (?) icon and selecting **External Storage**, as shown in Figure 11-16.

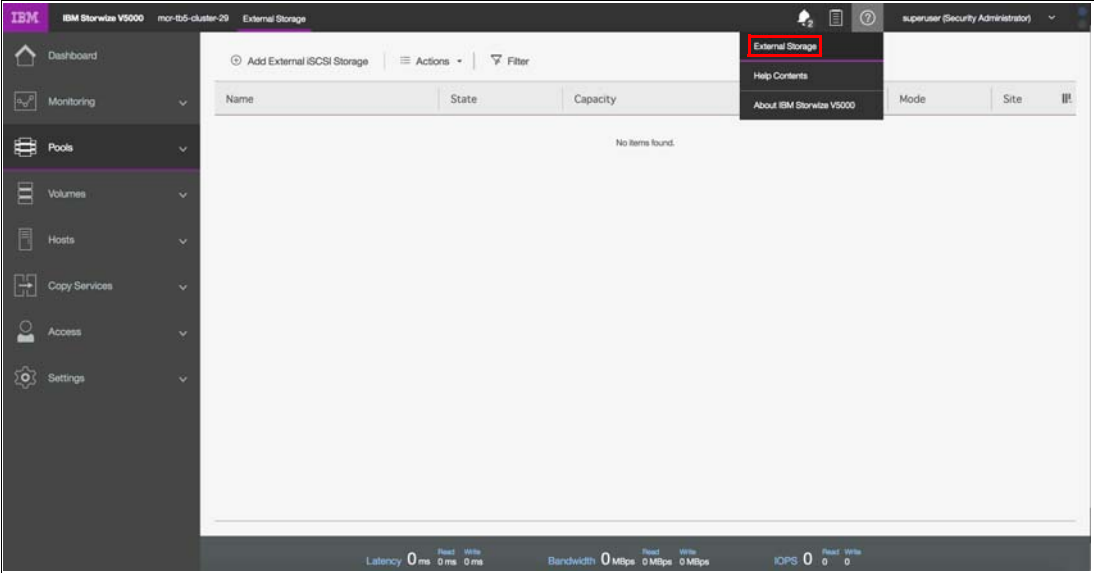


Figure 11-16 Extended help for external storage

The External Storage window that is shown in Figure 11-17 provides an overview of all of your external storage systems. The list shows the external storage systems that are managed by the Lenovo Storage V5030. With the help of the filter, you can show only the external storage systems that you want to work with. Clicking the arrow sign that precedes each of the external storage controllers provides more detailed information, including all of the MDisks that are mapped from it.

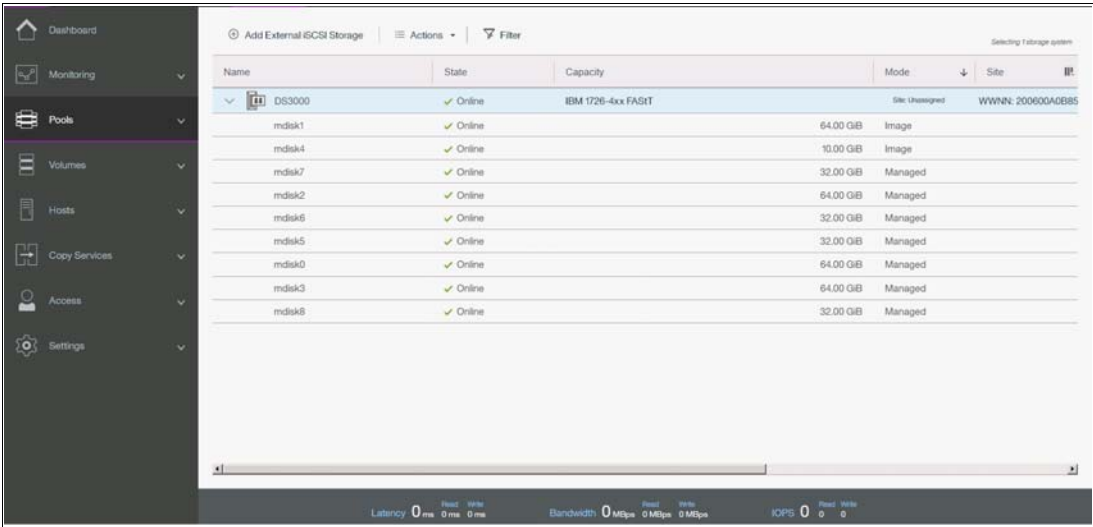


Figure 11-17 External Storage window

In the External Storage panel, there are options in the Actions menu that can be applied to external storage controllers, as shown in Figure 11-18. Select the external controller and click **Actions** to display the available options. Alternatively, right-click the external controller.

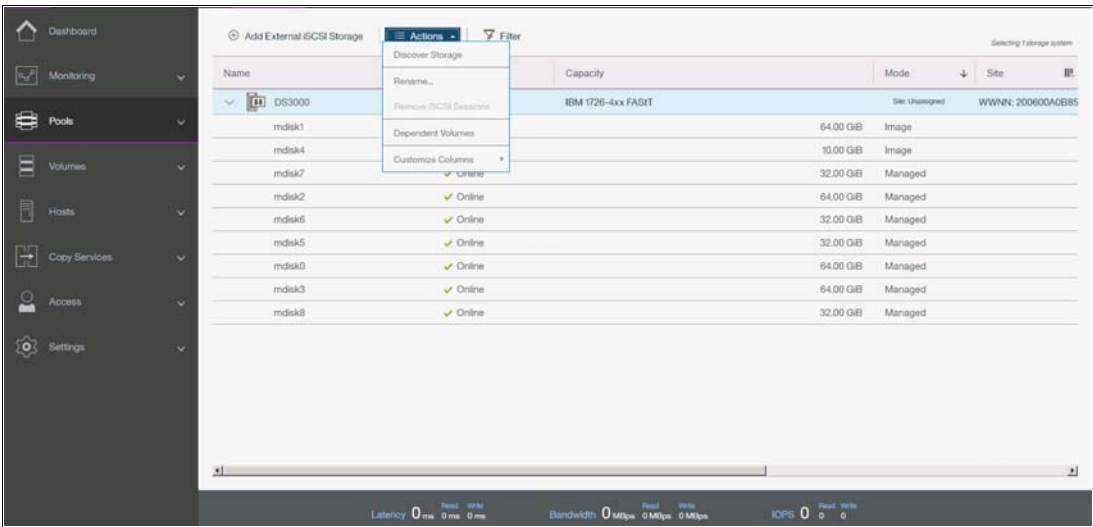
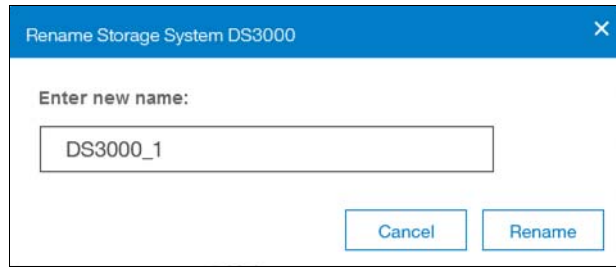


Figure 11-18 External controllers options under Actions menu

You can change the name of any external storage system by right-clicking the controller and selecting **Rename**. Alternatively, use the Actions drop-down menu and select **Rename**. In the Rename Storage System panel, define the storage controller name and click **Rename** as shown in Figure 11-19 on page 621.



Rename Storage System DS3000

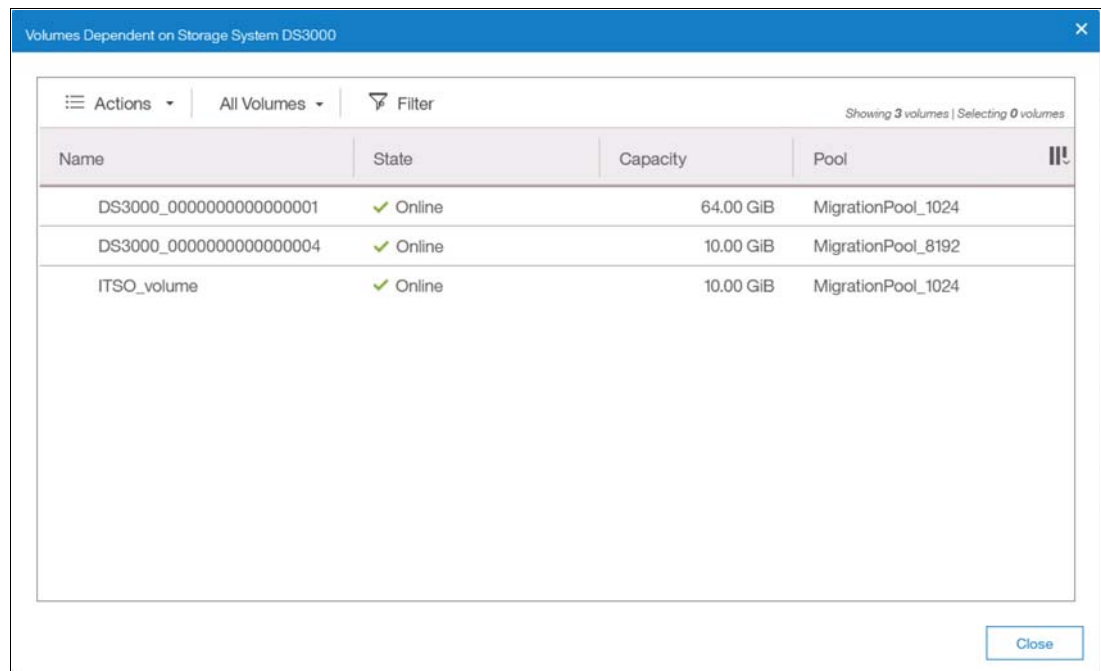
Enter new name:

DS3000_1

Cancel Rename

Figure 11-19 Rename Storage System panel

Click **Show Dependent Volumes** to display the logical volumes that depend on the selected external storage system, as shown in Figure 11-20.



Volumes Dependent on Storage System DS3000

Actions All Volumes Filter Showing 3 volumes | Selecting 0 volumes

Name	State	Capacity	Pool
DS3000_0000000000000001	✓ Online	64.00 GiB	MigrationPool_1024
DS3000_0000000000000004	✓ Online	10.00 GiB	MigrationPool_8192
ITSO_volume	✓ Online	10.00 GiB	MigrationPool_1024

Close

Figure 11-20 Volumes that depend on the external storage

From the **Volumes Dependent on Storage System** panel, multiple volume actions are available, as shown in Figure 11-21 on page 622.

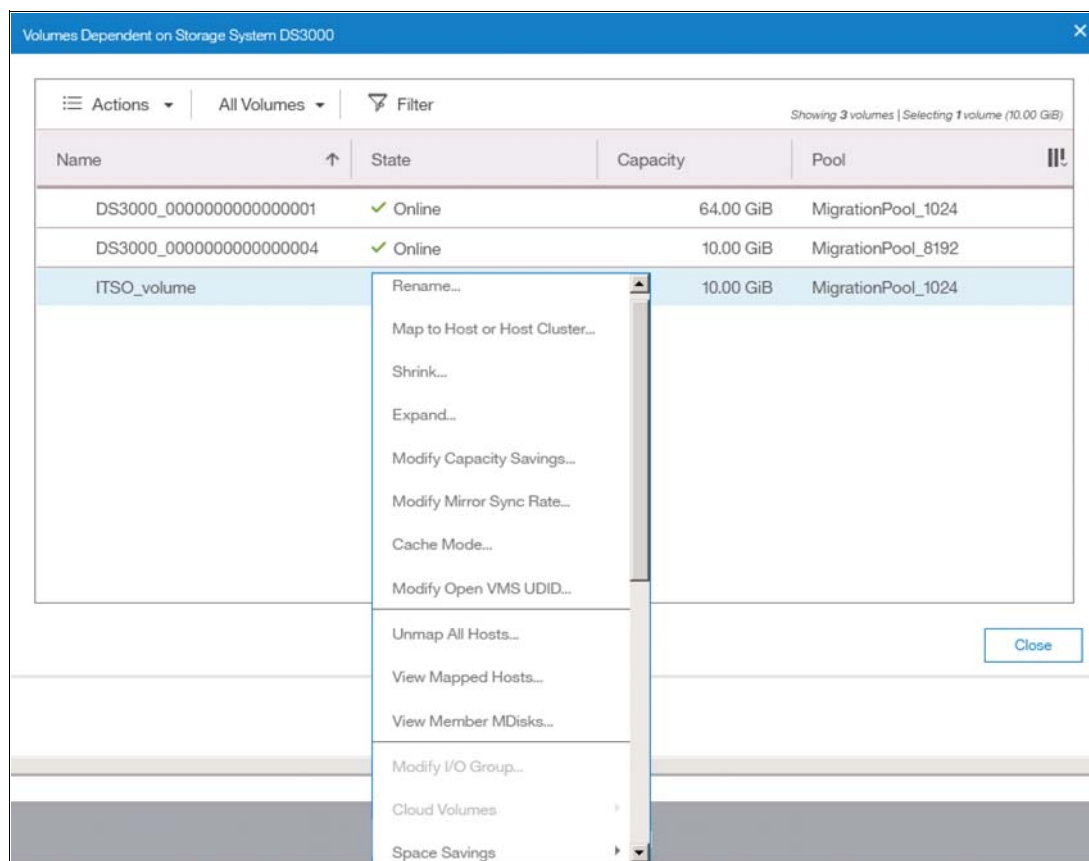


Figure 11-21 Volume actions

In the Lenovo Storage V5030 virtualization environment, you can migrate your application data nondisruptively from one internal or external storage pool to another, simplifying storage management with reduced risk.

Volume copy is another key feature that you can benefit from by using Lenovo Storage V5030 virtualization. Two copies can be created to enhance availability for a critical application. A volume copy can be also used to generate test data or for data migration.

For more information about the volume actions of the Lenovo Storage V5030 storage system, see Chapter 8, “Advanced host and volume administration” on page 349.

In the **External Storage** panel you can also right-click an **MDisk** (or use the Actions drop-down menu) to display the available options for a selected MDisk, as shown in Figure 11-22 on page 623.

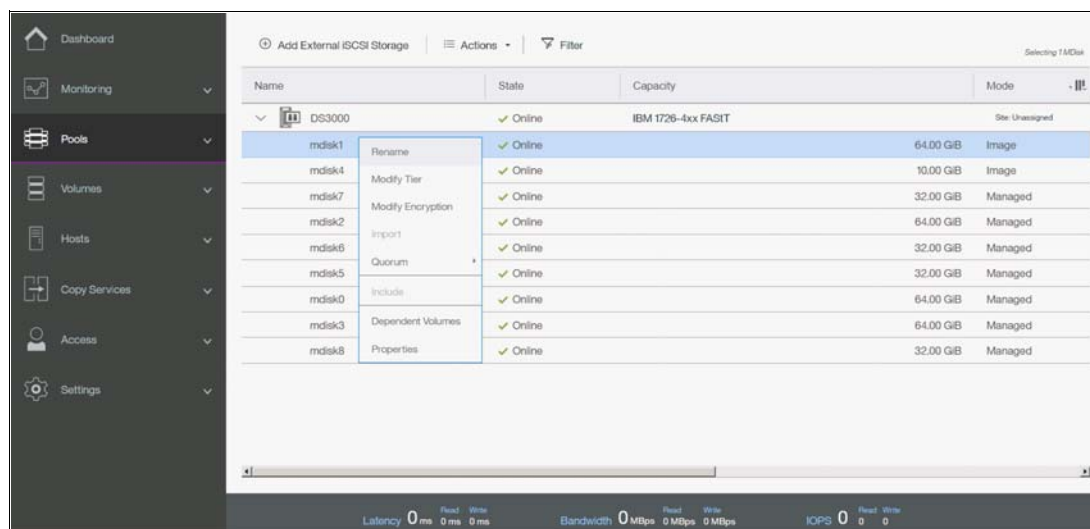


Figure 11-22 MDisk Actions menu in the External Storage window

11.2.4 Removing external storage

If you want to remove the external storage systems from the Lenovo Storage V5030 virtualized environment, the following options are available:

- To remove the external storage systems and discard the data on them, complete the following steps:
 - a. Stop any host I/O on the volumes.
 - b. Remove the volumes from the host file systems, logical volume, or volume group and remove the volumes from the host device inventory.
 - c. Remove the host mapping of volumes and the volumes themselves on the Lenovo Storage V5030.
 - d. Remove the storage pools to which the external storage systems belong, or you can keep the storage pool and remove the MDisks of the external storage from the storage pools.
 - e. Unzone and disconnect the external storage systems from the Lenovo Storage V5030.
 - f. Click **Detect MDisks** for the Lenovo Storage V5030 to discover the removal of the external storage systems.
- To remove the external storage systems and keep the volumes and their data on the Lenovo Storage V5030, complete the following steps:
 - a. Migrate volumes and their data to other internal or external storage pools that are on the Lenovo Storage V5030.
 - b. Remove the storage pools to which the external storage systems belong, or you can keep the storage pools and remove the MDisks of the external storage from the storage pools.
 - c. Unzone and disconnect the external storage systems from the Lenovo Storage V5030.
 - d. Click **Detect MDisks** for the Lenovo Storage V5030 to discover the removal of the external storage systems.

- To remove the external storage systems from the Lenovo Storage V5030 control and keep the volumes and their data on other external storage systems, complete the following steps:
 - a. Migrate volumes and their data to other internal or external storage pools on the Lenovo Storage V5030, as described in Chapter 7, “Storage migration” on page 323.
 - b. Remove the storage pools to which the original external storage systems belong, or you can keep the storage pools and remove the MDisks of that external storage from the storage pools.
 - c. Export the volumes that were migrated in step a to image mode with the new MDisks on the target external storage systems. For more information about the restrictions and prerequisites for migration, see Chapter 7, “Storage migration” on page 323.

You must record pre-migration information, for example, the original Small Computer System Interface (SCSI) identifiers (IDs) that the volumes used when they were mapped to hosts. Certain operating systems do not support a change of the SCSI ID during migration. Unzone and disconnect the external storage systems from the Lenovo Storage V5030.
 - d. Click **Detect MDisks** for the Lenovo Storage V5030 to discover the removal of the external storage systems.

RAS, monitoring, and troubleshooting

This chapter describes the reliability, availability, and serviceability (RAS) features and ways to monitor and troubleshoot the Lenovo Storage V3700 V2, V3700 V2 XP and V5030.

Specifically, this chapter provides information about the following topics:

- ▶ 12.1, “Reliability, availability, and serviceability features” on page 626
- ▶ 12.2, “System components” on page 627
- ▶ 12.3, “Configuration backup” on page 645
- ▶ 12.4, “System update” on page 650
- ▶ 12.5, “Monitoring” on page 666
- ▶ 12.6, “Audit log” on page 670
- ▶ 12.7, “Event log” on page 671
- ▶ 12.8, “Support Assistance” on page 679
- ▶ 12.9, “Collecting support information” on page 689
- ▶ 12.10, “Powering off the system and shutting down the infrastructure” on page 699

12.1 Reliability, availability, and serviceability features

This section describes the reliability, availability, and serviceability (RAS) features of the Lenovo Storage V3700 V2, V3700 V2 XP and V5030, as well as monitoring and troubleshooting. RAS features are important concepts in the design of the Lenovo Storage V3700 V2, V3700 V2 XP and V5030. Hardware and software features, design considerations, and operational guidelines all contribute to make the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 reliable.

Fault tolerance and a high level of availability are achieved with the following features:

- ▶ The RAID capabilities of the underlying disk subsystems
- ▶ The software architecture that is used by The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 nodes
- ▶ Auto-restart of nodes that are stopped
- ▶ Battery units to provide cache memory protection in a site power failure
- ▶ Host system multipathing and failover support

High levels of serviceability are achieved with the following features:

- ▶ Cluster error logging
- ▶ Asynchronous error notification
- ▶ Dump capabilities to capture software-detected failures
- ▶ Concurrent diagnostic procedures
- ▶ Directed maintenance procedures
- ▶ Concurrent log analysis and memory dump data recovery tools
- ▶ Concurrent maintenance of all of the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 components
- ▶ Concurrent upgrade of Lenovo Storage V3700 V2, V3700 V2 XP and V5030 software and microcode of drives
- ▶ Concurrent addition or deletion of a node canister in a cluster
- ▶ Software recovery through the Service Assistant Tool
- ▶ Automatic software version correction when a node is replaced
- ▶ Detailed status and error conditions that are displayed through the Service Assistant Tool
- ▶ Error and event notification through Simple Network Management Protocol (SNMP), syslog, and email
- ▶ Access to the Service Assistant Tool through the tech port for network connection problems
- ▶ Remote support personnel is able to access the system to complete troubleshooting and maintenance tasks

At the core of the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 is a redundant pair of *node canisters*. The two canisters share the load of transmitting and receiving data between the attached hosts and the disk arrays.

12.2 System components

This section describes each of the components that make up the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems. The components are described in terms of location, function, and serviceability.

12.2.1 Enclosure midplane

The *enclosure midplane* connects the node or expansion canisters to the power supply units and to the drives. The midplane is part of the enclosure midplane assembly, which consists of the midplane and the front section of the enclosure.

During the basic system configuration, vital product data (VPD) is written to the enclosure midplane. On a control enclosure midplane, the VPD contains information, such as worldwide node name (WWNN) 1, WWNN 2, machine type and model, machine part number, and serial number. On an expansion enclosure midplane, the VPD contains information, such as machine type and model, machine part number, and serial number.

The enclosure midplane is initially generic and it is configured as a control enclosure midplane or expansion enclosure midplane only when the VPD is written. After the VPD is written, a control enclosure midplane is no longer interchangeable with an expansion enclosure midplane and vice versa.

Important: The enclosure midplane must be replaced only by a trained representative.

For information about the midplane replacement process, see the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 Information Center at:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/tbrd_rmvrplparts_1955wm.html

For a complete overview of maintenance tasks see:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.common.nav.doc/overview_storage_vseries.html

12.2.2 Node canisters

Two node canister slots are on the top of the unit. The left slot is canister 1, and the right slot is canister 2.

Figure 12-1 shows the rear view of a fully equipped control enclosure.




Figure 12-1 Rear view of a control enclosure with two node canisters (the Storwize V5020)

USB ports

Each node canister has one USB port. The location of the port is the same on every model, and no indicators are associated with it.

Figure 12-2 shows the location of the USB port.




Figure 12-2 Node canister USB port (the Lenovo Storage V3700 V2)

The USB flash drive is not required to initialize the system configuration. However, it can be used for other functions. Using the USB flash drive is required in the following situations:

- ▶ When you cannot connect to a node canister in a control enclosure by using the service assistant or the technician port, and you want to see the status of the node or re-enable the technician port.
- ▶ When you do not know, or cannot use, the service IP address for the node canister in the control enclosure and must set the address.
- ▶ When you have forgotten the superuser password and must reset the password.

Ethernet ports

The Lenovo Storage V3700 V2 and Lenovo Storage V3700 V2 XP node canisters have two 100/1000 Mbps Ethernet ports. Both ports can be used for management, Internet Small Computer System Interface (iSCSI) traffic, and Internet Protocol (IP) replication. Additionally, port 2 can be used as a technician port (the white box with “T” in the center of the box) for system initialization and servicing. After initialization the technician port will be disabled. It is possible to reactivate the technician port later again via CLI commands.

Figure 12-3 shows the Ethernet ports on the Lenovo Storage V3700 V2.

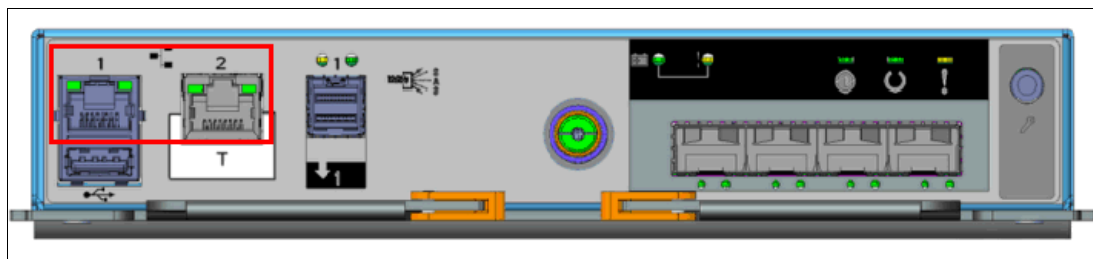


Figure 12-3 Lenovo Storage V3700 V2 Ethernet ports

Figure 12-4 on page 629 shows the Ethernet ports on the Lenovo Storage V3700 V2 XP.




Figure 12-4 Lenovo Storage V3700 V2 XP Ethernet ports

Each Lenovo Storage V5030 node canister has two 1/10 Gbps Ethernet ports and one Ethernet technician port. Port 1 and 2 can be used for management, iSCSI traffic, and IP replication. Port T can be used as a technician port for system initialization and service only.

Figure 12-5 shows the Ethernet ports on the Lenovo Storage V5030.

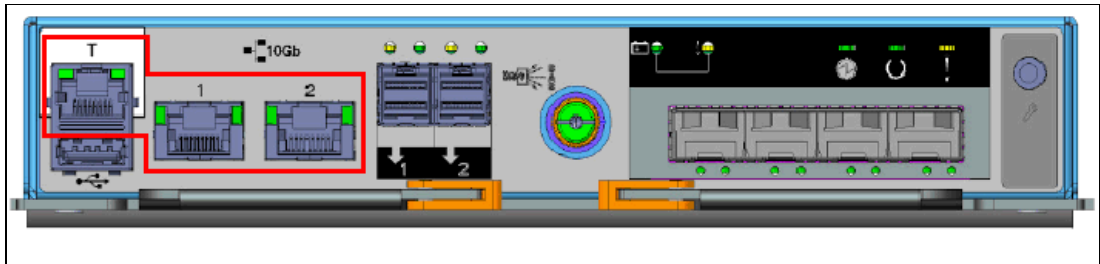


Figure 12-5 Lenovo Storage V5030 Ethernet ports

Each port has two LEDs that display the status of its activity. Their meanings are shown in Table 12-1.

Table 12-1 Ethernet port status LEDs

Name and position	Color	State	Meaning
Activity (left)	Green	Flashing	The link is active.
		Off	The link is inactive.
Link speed (right)	Green	Solid	A connection exists to a remote device at 1 Gbps or more.
		Off	No connection exists to a remote device, or the link is connected at less than 1 Gbps.

Serial-attached SCSI ports

Each Lenovo V3700 V2 node canister uses one 12 Gbps serial-attached SCSI (SAS) port to connect optional expansion enclosures. This port does not support host attachment.

Figure 12-6 on page 630 shows the SAS ports on the Lenovo Storage V3700 V2.

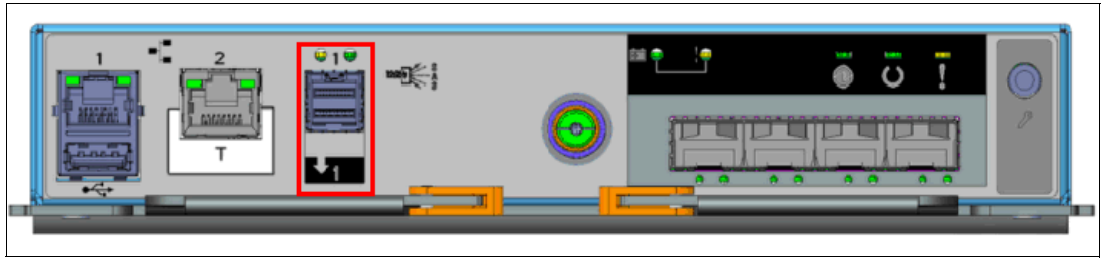


Figure 12-6 Storwize V5010 SAS ports

Each Lenovo Storage V3700 V2 XP node canister has three 12 Gbps SAS ports. Port 1 can be used to connect optional expansion enclosures, and ports 2 and 3 can be used for host attachment.

Figure 12-7 shows the SAS ports on the Lenovo Storage V3700 V2 XP.




Figure 12-7 Lenovo Storage V3700 V2 SAS ports

Each Lenovo Storage V5030 node canister has two 12 Gbps SAS ports to connect optional expansion enclosures. This port does not support host attachment.

Figure 12-8 shows the SAS ports on the Lenovo Storage V5030.

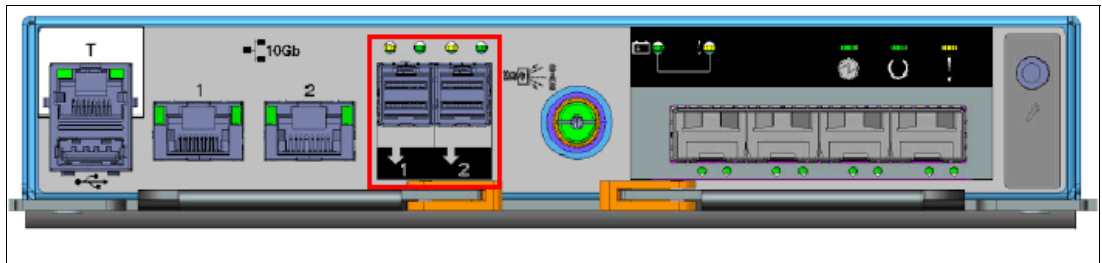


Figure 12-8 Lenovo Storage V5030 SAS ports

Each port has two LEDs that display the status of its activity. Their meanings are shown in Table 12-2 on page 631.

Table 12-2 SAS port status LEDs

Name and position	Color	State	Meaning
Fault (left)	Amber	Solid	<ul style="list-style-type: none"> ▶ One of the following conditions has occurred: ▶ One or more, but not all, of the 4 lanes are up. (If no lanes are up, the activity light will be off.) ▶ One or more of the lanes is running at a different speed to the others. ▶ One or more of the up lanes are attached to a different address to the others. ▶ An unsupported device is plugged into this SAS port.
		Off	No fault exists. All four lanes (phys) have a connection.
Link (right)	Green	Solid	A connection exists on at least one lane (phy).
		Off	None of the SAS connections are working.

Battery status

Each node canister houses a battery, the status of which is displayed by two LEDs on the back of the unit, as shown in Figure 12-9.




Figure 12-9 Battery status LEDs (the Storwize V5020)

The meaning of each LED is described in Table 12-3.

Table 12-3 Battery status LEDs

Name and position	Color	State and Meaning
Battery status (left)	Green	<ul style="list-style-type: none"> ▶ FAST BLINK The battery is charging. It does not have a sufficient charge to perform a “fire hose” dump. ▶ BLINK The battery has sufficient charge to perform one “fire hose” dump. ▶ ON The battery is fully charged and has sufficient charge to perform two “fire hose” dumps. ▶ OFF The battery is not available for use.
Fault (right)	Amber	<ul style="list-style-type: none"> ▶ OFF If the LED is off, no known conditions are preventing normal operation, unless the battery status LED is also on. ▶ ON An active condition or fault could compromise normal operation. ▶ SLOW BLINK There is a non-critical fault with the battery.

Name and position	Color	State and Meaning
Battery in use	Green	<ul style="list-style-type: none"> ▶ OFF The battery is not being used to power the canister. ▶ FAST BLINK The battery is currently providing power for a “fire hose” dump.

Canister status

The status of each canister is displayed by three LEDs on the back of the unit, as shown in Figure 12-10.




Figure 12-10 Node canister status LEDs (the Lenovo Storage V3700 V2 XP)

The meaning of each LED is described in Table 12-4.

Table 12-4 Canister status LEDs

Name and position	Color	State and Meaning
Power (left)	Green	<ul style="list-style-type: none"> ▶ OFF No power is available or power is coming from the battery. ▶ SLOW BLINK Power is available but the main CPU is not running; the system is in standby mode. ▶ FAST BLINK System is in self test. ▶ ON Power is available and the system code is running.
Status (middle)	Green	<ul style="list-style-type: none"> ▶ OFF Indicates one of the following conditions: <ul style="list-style-type: none"> — No power to the canister — Canister is in standby mode or self test — Operating system is loading ▶ BLINK The canister is in candidate or service state. It is not performing I/O. It is safe to remove the node. ▶ BLINK FAST The canister is carrying out a fire hose dump. ▶ ON The canister is active, able to perform I/O, or starting. The system is part of a cluster.

Name and position	Color	State and Meaning
Canister Fault (right)	Amber	<ul style="list-style-type: none"> ► OFF The node is in candidate or active state. Any error that has been detected is not severe enough to stop the node participating in a cluster or performing I/O. ► BLINK The canister is being identified. There might or might not be a fault condition. ► ON The node is in service state or an error exists that might be stopping the system code from starting (node error 550). The node canister cannot become active in the system until the problem is resolved. The problem is not necessarily related to a hardware component.

Replaceable components

The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 node canister contains the following field-replaceable (client-replaceable) components:

- Host Interface Card
- Memory
- Battery

Figure 12-11 shows the location of these parts within the node canister.




Figure 12-11 Node canister client-replaceable components

Note: Because these components are inside the node canister, their replacement leads to a redundancy loss until the replacement is complete.

Host Interface Card replacement procedure

For information about the Host Interface Card (HIC) replacement process, see the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 Information Center at this web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/v3700_rplc_hic.html

Figure 12-12 shows a HIC replacement.




Figure 12-12 HIC replacement

Memory replacement procedure

For information about the memory replacement process, see the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 Information Center at this web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/v3700_rplc_nodecan_dimm.html

Figure 12-13 on page 635 shows the location of the memory modules.




Figure 12-13 Location of memory modules

Figure 12-14 on page 636 shows a memory replacement.

Note: The memory modules do not stand up. They lie in a cascading fashion.



Figure 12-14 Memory replacement

Battery Backup Unit replacement procedure

Attention: The battery is a lithium ion battery. To avoid a possible explosion, do not incinerate the battery. Exchange the battery only with the part that is approved by Lenovo.

Because the Battery Backup Unit (BBU) replacement leads to a redundancy loss until the replacement is complete, we advise that you replace the BBU only when you are instructed to replace it. We advise you to follow the Directed Maintenance Procedure (DMP).

During the procedure, while you lift and lower the battery, grasp the blue handle on each end of the battery and keep the battery parallel to the canister system board, as shown in Figure 12-15 on page 637.



Figure 12-15 BBU replacement

Important: During the replacement, the battery must be kept parallel to the canister system board while the battery is removed or replaced. Keep equal force, or pressure, on each end.

For more information about the BBU replacement process, see the Lenovo Information Center at this web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/v3700_rplc_batt_nodecan.html

More replacement procedures can be found on Information Center web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/tbrd_rmvrplparts_1955wm.html

12.2.3 Expansion canisters

Two expansion canister slots are on the top of the unit. As with the control enclosure, the left slot is canister 1 and the right slot is canister 2.

Figure 12-16 shows the rear view of a fully equipped expansion enclosure.




Figure 12-16 Rear view of an expansion enclosure with two expansion canisters

SAS ports

SAS ports are used to connect the expansion canister to the node canister or to an extra expansion canister in the chain. Figure 12-17 shows the SAS ports that are on the expansion canister.




Figure 12-17 Expansion canister SAS ports

Each port has two LEDs that display the status of its activity. Their meanings are shown in Table 12-5.

Table 12-5 SAS port status LEDs

Name and position	Color	State	Meaning
Fault (left)	Amber	Solid	One of the following errors exists: <ul style="list-style-type: none"> ► Only 1, 2, or 3 lanes (phys) have a connection. ► Not all of the lanes (phys) that have a connection are running at the same speed. ► Not all of the lanes (phys) that have a connection are attached to the same address. ► An unsupported device is connected to the port.
		Off	No fault exists. All four lanes (phys) have a connection.
Link (right)	Green	Solid	A connection exists on at least one lane (phy).
		Off	No connection exists on any lane (phy).

Canister status

The status of each expansion canister is displayed by three LEDs on the back of the unit, as shown in Figure 12-18.




Figure 12-18 Enclosure canister status LEDs

The meaning of each LED is described in Table 12-6 on page 639.

Table 12-6 Expansion canister status LEDs

Name and position	Color	State	Meaning
Power (left)	Green	Solid	The canister is receiving power.
		Off	No power is available, or the power is coming from the battery.
Status (middle)	Green	Solid	The canister is running normally.
		Blinking	The canister is unable to read data from the midplane.
		Off	The system is off, in standby, or running a self-test, or the operating system is loading.
Fault (right)	Amber	Solid	A fault requires part replacement, or the canister is still starting.
		Blinking	The canister is being identified. A fault might or might not exist.
		Off	The canister has no faults that require part replacement.

12.2.4 Disk subsystem

This section describes the parts of the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 disk subsystem, which is made up of control and expansion enclosures.

The Lenovo V3700 V2 and Lenovo V3700 V2 XP can have one control enclosure. The Lenovo Storage V5030 can consist of 1 or 2 control enclosures.

Each Lenovo Storage V3700 V2 and Lenovo Storage V3700 V2 XP control enclosure can attach up to 10 expansion enclosures. Each Lenovo Storage V5030 control enclosure can attach up to 20 expansion enclosures.

SAS cabling

Expansion enclosures are attached to control enclosures and between each other by using SAS cables.

A set of correctly interconnected enclosures is called a *chain*. Each chain is made up of two *strands*. A strand runs through the canisters that are in the same position in each enclosure in the chain. Canister 1 of an enclosure is cabled to canister 1 of the downstream enclosure. Canister 2 of an enclosure is cabled to canister 2 of the downstream enclosure.

Each strand consists of 4 phys, and each phy operates at 12 Gbps, therefore a strand has a usable speed of 48 Gbps.

A strand starts with a SAS initiator chip inside an Lenovo Storage V3700 V2, V3700 V2 XP and V5030 node canister and progresses through SAS expanders, which connect to the disk drives. Each canister contains an *expander*. Each drive has two ports, each of which is connected to a different expander and strand. This configuration means that both nodes directly access each drive, and no single point of failure exists.

At system initialization, when devices are added to or removed from strands (and at other times), the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 software performs a discovery process to update the state of the drive and enclosure objects.

The Lenovo Storage V3700 V2 supports one SAS chain for each control enclosure, and up to 10 expansion enclosures can be attached to this chain. The node canister uses SAS port 1 for expansion enclosures.

Figure 12-19 shows the SAS cabling on a Lenovo Storage V3700 V2 with three attached expansion enclosures.




Figure 12-19 SAS expansion cabling on the Lenovo Storage V3700 V2

The Lenovo Storage V3700 V2 XP supports one SAS chain for each control enclosure, and up to 10 expansion enclosures can be attached to this chain. The node canister uses SAS port 1 for expansion enclosures.

Figure 12-20 on page 641 shows the SAS cabling on a Lenovo Storage V3700 V2 XP with three attached expansion enclosures.




Figure 12-20 SAS expansion cabling on the Lenovo Storage V3700 V2 XP

The Lenovo Storage V5030 supports two SAS chains for each control enclosure, and up to 10 expansion enclosures can be attached to each chain. The node canister uses SAS port 1 for expansion enclosures.

Figure 12-21 on page 642 shows the SAS cabling on a Lenovo Storage V5030 with six attached expansion enclosures (three enclosures in each chain).




Figure 12-21 SAS expansion cabling on the Lenovo Storage V5030

Important: When a SAS cable is inserted, ensure that the connector is oriented correctly by confirming that the following conditions are met:



- ▶ The pull tab must be below the connector.
- ▶ Insert the connector gently until it clicks into place. If you feel resistance, the connector is probably oriented the wrong way. Do not force it.
- ▶ When the connector is inserted correctly, the connector can be removed only by pulling the tab.
- ▶ Cabling is done from the controller view top → down. Top/down button up is **not** supported.

Drive slots

The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 have different types of enclosures, depending on the model, warranty, and number of drive slots.

Table 12-7 shows the drive slots on each enclosure type.

Table 12-7 Drive slots for each enclosure type

Enclosure type	Drive slots
<ul style="list-style-type: none"> ▶ Control enclosure 2077/2078-112 ▶ Control enclosure 2077/2078-212 ▶ Control enclosure 2077/2078-312 ▶ Expansion enclosure 2077/2078-12F 	12 x 3.5-inch slots 
<ul style="list-style-type: none"> ▶ Expansion enclosure 2077/2078-92F 	92 x 3.5-inch slots (usage of 2.5-inch drives possible with carriers)
<ul style="list-style-type: none"> ▶ Control enclosure 2077/2078-124 ▶ Control enclosure 2077/2078-224 ▶ Control enclosure 2077/2078-324 ▶ Expansion enclosure 2077/2078-24F 	24 x 2.5-inch slots 

Drive replacement procedure

You can reseal or replace a failed drive in a Lenovo Storage V3700 V2, V3700 V2 XP or V5030 by removing it from its enclosure and replacing it with the correct new drive without requiring the Directed Maintenance Procedure to supervise the service action.

The system can automatically perform the drive hardware validation tests and can promote the drive into the configuration if these tests pass, automatically configuring the inserted drive as a spare. The status of the drive after the promotion can be recorded in the event log either as an informational message or an error if a hardware failure occurs during the system action.

For more information about the drive replacement process, see the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 Information Center at this web pages:

- ▶ Replacing a 3.5-inch drive assembly
http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/v3700_rplc_35_drv_assembly.html
- ▶ Replacing a 2.5-inch drive assembly:
http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/v3700_rplc_25_drv_assembly.html

12.2.5 Power supply units

All enclosures require two power supply units (PSUs) for normal operation. A single PSU can power the entire enclosure for redundancy. We advise that you supply AC power to each PSU from different power distribution units (PDUs).

Figure 12-22 on page 644 shows a fully equipped control enclosure with two supply units. The PSUs are identical between the control and expansion enclosures.




Figure 12-22 Power supply units

The left PSU is numbered 1, and the right PSU is numbered 2.

Power supplies in both control and expansion enclosures are hot-swappable and replaceable without a need to shut down a node or cluster. If the power is interrupted in one node canister for less than 2.5 seconds, the canister cannot perform a fire hose dump and continues operation from battery.

PSU status

Each PSU has three LEDs that display the status of its activity. The LEDs are the same for the control and expansion units.

Figure 12-23 shows the PSU status LEDs.



Figure 12-23 PSU status LEDs

The meaning of each LED is shown in Table 12-8 on page 645.

Table 12-8 PSU status LEDs

Name and position	Color	State	Meaning
Input status (top)	Green	Solid	Input power is available.
		Off	No input power is available.
Output status (middle)	Green	Solid	PSU is providing DC output power.
		Off	PSU is not providing DC output power.
Fault (bottom)	Amber	Solid	A fault exists with the PSU.
		Blinking	The PSU is being identified. A fault might exist.
		Off	No fault is detected.

PSU replacement procedure

For information about the PSU replacement process, see the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 Information Center at this web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/v3700_rplc_pwrsupply.html

12.3 Configuration backup

The configuration backup file must be used if a serious failure occurs that requires the system configuration to be restored. The file contains configuration data of arrays, pools, volumes, and so on (but no client data).

The configuration backup file can be downloaded and saved by using the graphical user interface (GUI) or the command-line interface (CLI). The CLI option requires you to log in to the system and download the file by using Secure Copy Protocol (SCP). It is a preferred practice for an automated backup of the configuration.

Important: Save the configuration files of the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 regularly. The best approach is to save daily and automate this task. Always perform the additional manual backup before you perform any critical maintenance task, such as an update of the microcode or software version.

The backup file is updated by the cluster every day and stored in the /dumps directory. Even so, it is important to start a manual backup after you change your system configuration.

To successfully perform the configuration backup, follow the prerequisites and requirements:

- ▶ All nodes must be online.
- ▶ No independent operations that change the configuration can be running in parallel.
- ▶ No object name can begin with an underscore.

Important: You can perform an ad hoc backup of the configuration only from the CLI. However, the output of the command can be downloaded from both the CLI and the GUI.

12.3.1 Generating a manual configuration backup by using the CLI

You can use the CLI to trigger a configuration backup either manually on an ad hoc basis or by an automatic process regularly. The **svcconfig backup** command generates a new backup file. Triggering a backup by using the GUI is not possible, but you can save the output from the GUI.

Example 12-1 shows the output of the **svcconfig backup** command.

Example 12-1 Triggering a backup by using the CLI

```
>svcconfig backup
.....
.....
.....
CMMVC6155I SVCCONFIG processing completed successfully
```

The **svcconfig backup** command creates three files that provide information about the backup process and cluster configuration. These files are created in the /dumps directory on the configuration node and can be retrieved by using SCP. Use the **lsdumps** command to list them, as shown in Example 12-2.

Example 12-2 Listing the backup files by using the CLI

```
>lsdumps
id filename
...
48 svc.config.backup.xml_781000E-1
49 svc.config.backup.sh_781000E-1
50 svc.config.backup.log_781000E-1
...
```

The three files that are created by the backup process are described in Table 12-9.

Table 12-9 Files that are created by the backup process

File name	Description
svc.config.backup.xml_<serial>	This file contains the cluster configuration data.
svc.config.backup.sh_<serial>	This file contains the names of the commands that were issued to create the backup of the cluster.
svc.config.backup.log_<serial>	This file contains details about the backup, including any error information that might be reported.

12.3.2 Downloading a configuration backup by using the GUI

The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 do not offer an option to initiate a backup from the GUI. However, you can download existing daily backups or manual backups that were triggered from the CLI.

To download a configuration backup file by using the GUI, complete the following steps:

1. Browse to **Settings** → **Support** → **Support Package** and select **Manual Upload Instructions**

See Figure 12-24 on page 647.

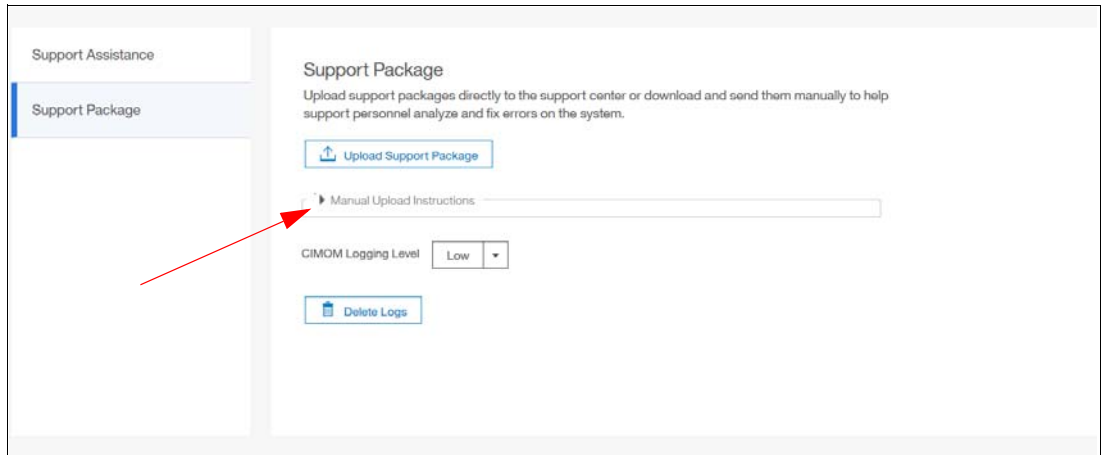


Figure 12-24 Manual Upload Instructions

When you select **Manual Upload Instructions**, a window opens (Figure 12-25).

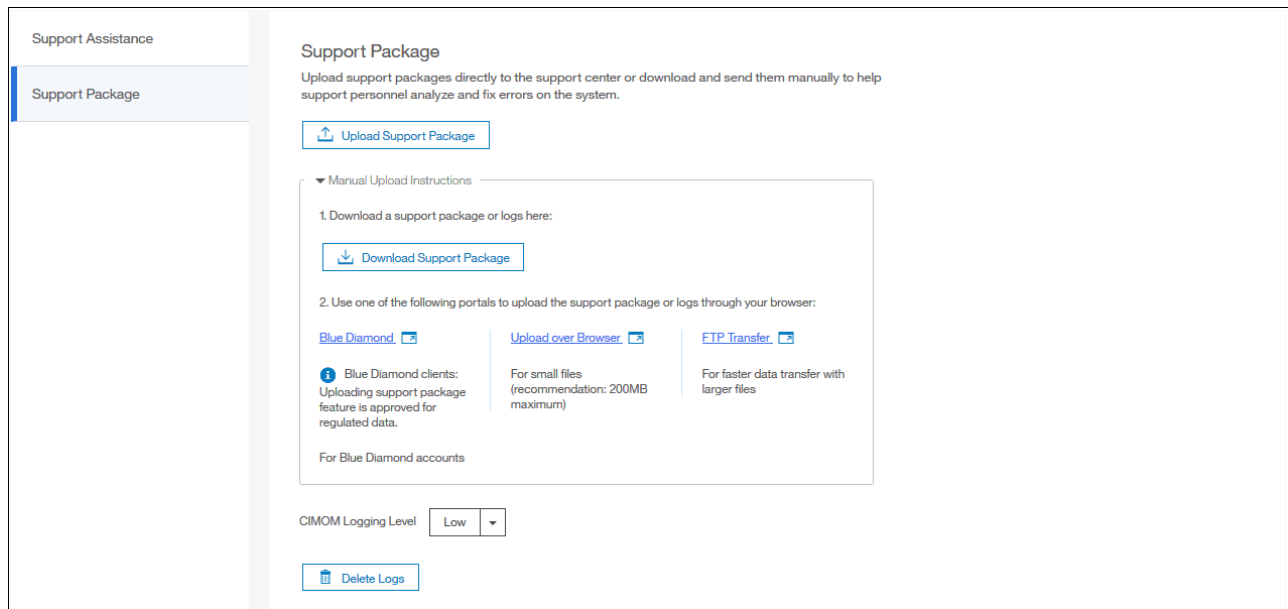


Figure 12-25 Download Support Package

Pressing the Button **Download Support Package** brings you to the next option, where you can select the different kinds of Support packages, see Figure 12-26 on page 648 for details.

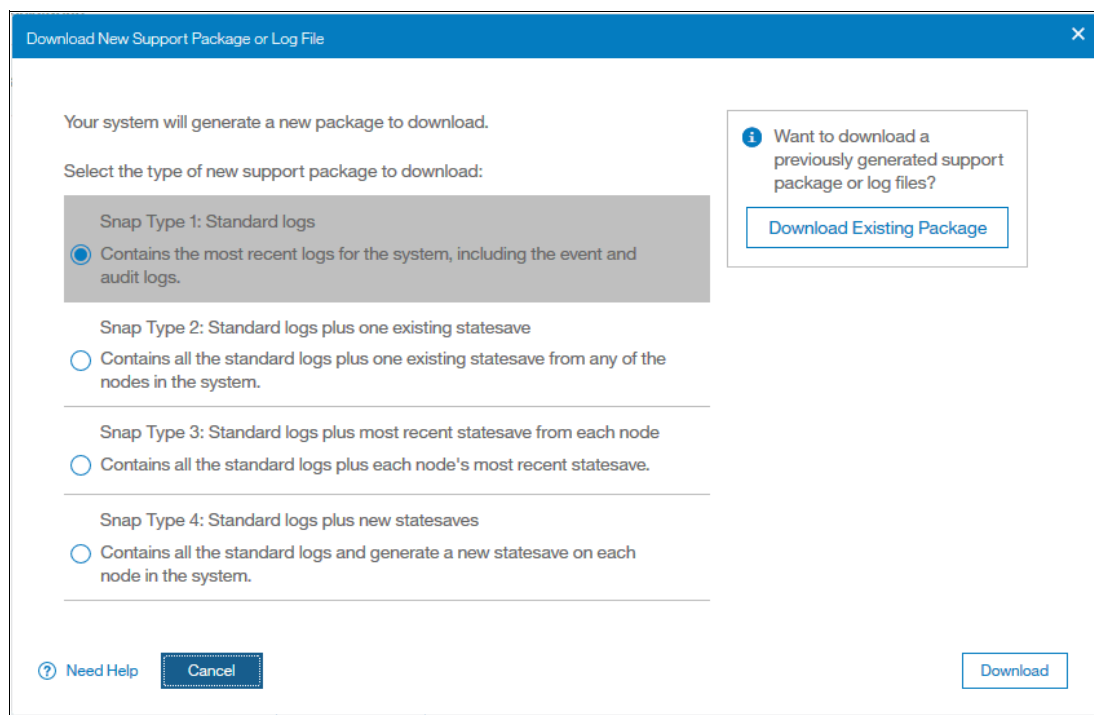


Figure 12-26 Download Support Package

Select **Download Existing Package** to get a list of all the available log files that are stored on the configuration node, as shown in Figure 12-27.

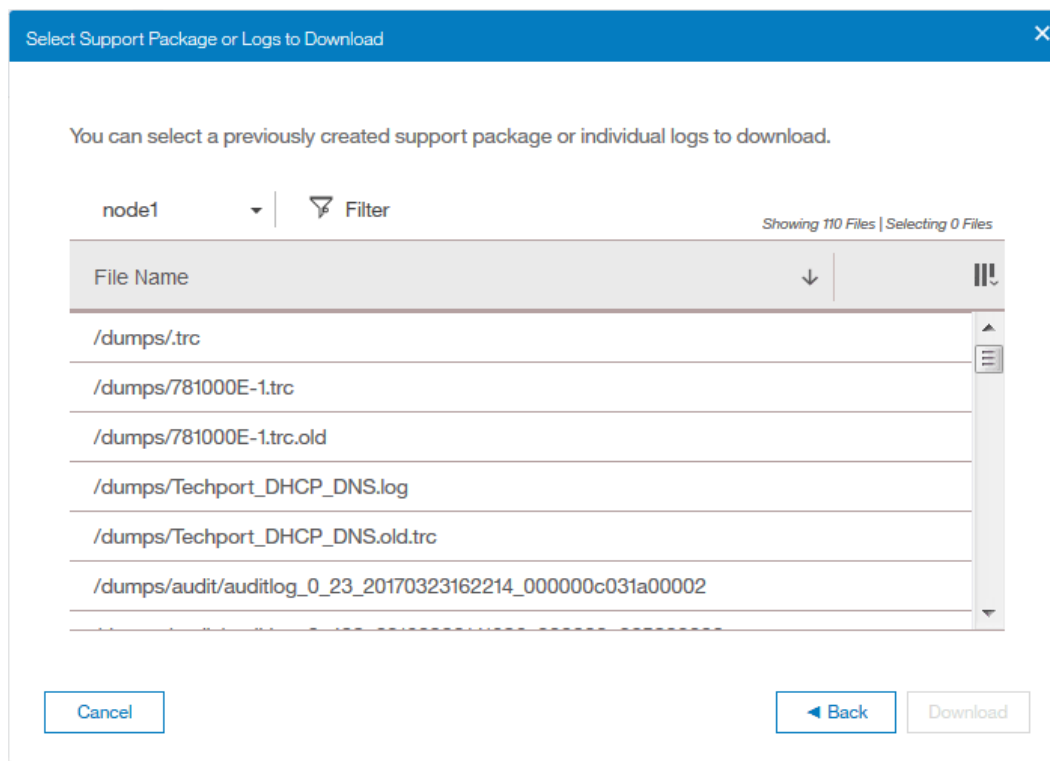


Figure 12-27 Full log listing option

2. Search for the files that are named `svc.config.backup.xml_*`, `svc.config.backup.sh_*`, and `svc.config.backup.log_*`. Select the files, right-click, and select **Download**, as shown in Figure 12-28.

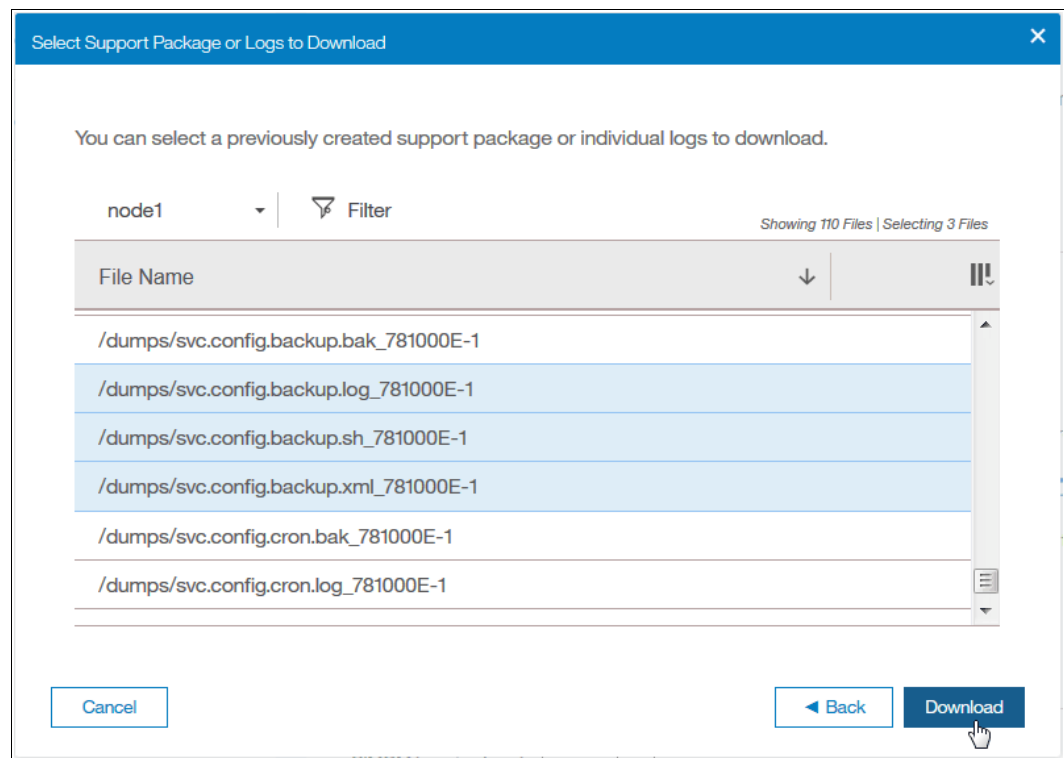


Figure 12-28 Backup files download

Even though the configuration backup files are updated automatically daily, it might be useful to verify the time stamp of the actual file. Open the `svc.config.backup.xml_xx` file with a text editor and search for the string `timestamp=`, which is near the top of the file. Figure 12-29 shows the file and the timestamp information.

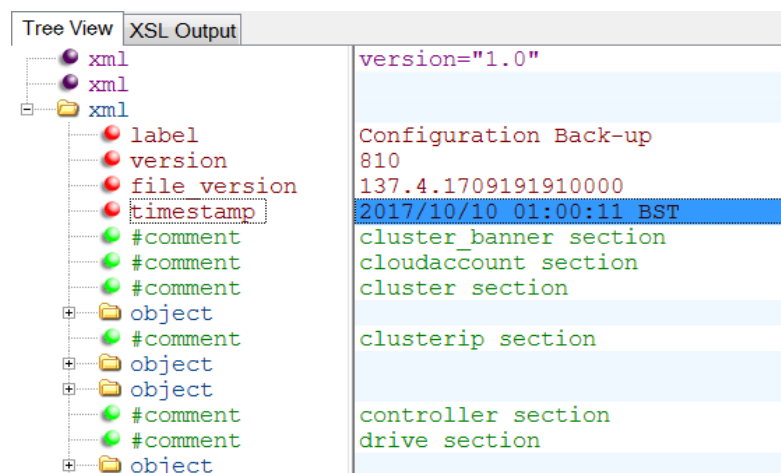


Figure 12-29 Timestamp in the backup xml file

12.4 System update

The system update process involves updating the entire Lenovo Storage V3700 V2, V3700 V2 XP and V5030 environment.

The node canister software and the drive firmware are updated separately so these tasks are described in different topics.

Note: IBM Storwize V5000 for Lenovo hardware is not supported by V8.1 or later. The V7.7.1 and V7.8.1 code streams will continue to be updated with critical fixes for this hardware.

12.4.1 Updating node canister software

For information about the latest software and to download the software package, go to the following web page:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1004336>

The GUI also shows whether a software update is available and the latest software level when you navigate to **Settings** → **System** → **Update System**, as shown in Figure 12-30.

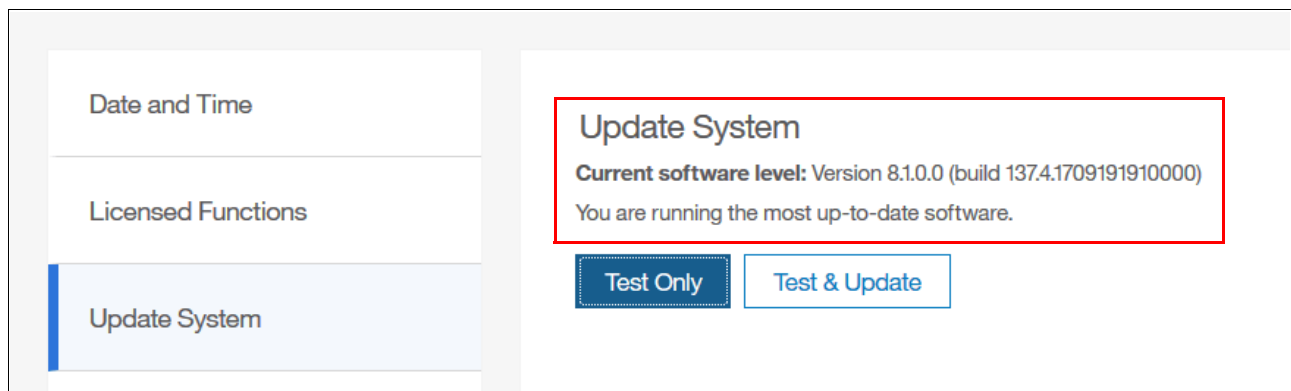


Figure 12-30 Latest software level available

Important: Certain levels of code support updates only from specific previous levels. If you update to more than one level above your current level, you might be required to install an intermediate level. For information about on update compatibility, see this web page:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1004336>

Preparing for the update

Allow sufficient time to plan your tasks, review your preparatory update tasks, and complete the update of the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 environment. The update procedures can be divided into the following general update tasks, as shown in Table 12-10 on page 651.

Table 12-10 Software update tasks

Sequence	Upgrade tasks
1	Decide whether you want to update automatically or manually. During an automatic update procedure, the clustered system updates each of the nodes systematically. The automatic method is the preferred procedure for updating software on nodes. However, you can update each node manually.
2	Ensure that Common Information Model (CIM) object manager (CIMOM) clients are working correctly. When necessary, update these clients so that they can support the new version of the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 code. Examples can be operating system (OS) versions and options, such as FlashCopy Manager or VMware plug-ins.
3	Ensure that multipathing drivers in the environment are fully redundant. If you experience failover issues with multipathing driver support, resolve these issues before you start normal operations.
4	Update other devices in the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 environment. Examples might include updating the hosts and switches to the correct levels.
5	Update your Lenovo Storage V3700 V2, V3700 V2 XP and V5030.

Important: Ensure that no unfixed errors are in the log and that the system date and time are correctly set before you start the update.

The amount of time that it takes to perform a node canister update can vary depending on the amount of preparation work that is required and the size of the environment. Generally, to update the node software, allow 20 - 40 minutes for each node canister and a single 30-minute wait when the update is halfway complete. One node in each I/O group can be upgraded to start, then the system can wait 30 minutes before it upgrades the second node in each I/O group. The 30-minute wait allows the recently updated node canister to come online and be confirmed as operational, and it allows time for the host multipath to recover.

The software update can be performed concurrently with normal user I/O operations. After the updating node is unavailable, all I/O operations fail to that node and the failed I/O operations are directed to the partner node of the working pair. Applications do not see any I/O failures.

The maximum I/O rate that can be sustained by the system might degrade while the code is uploaded to a node, the update is in progress, the node is rebooted, and the new code is committed because write caching is disabled during the node canister update process.

Important: Ensure that the multipathing drivers are fully redundant with every available path and online. You might see errors that are related to the paths, which can go away (failover) and the error count can increase during the update. When the paths to the nodes return, the nodes fall back to become a fully redundant system.

When new nodes are added to the system, the upgrade package is automatically downloaded to the new nodes from the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 systems.

Update test utility

The Lenovo Storage V3700 V2, V3700 V2 XP and V5030 update test utility checks for known issues that can cause problems during a software update. You can download the utility and read more about it at information center web page:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/svc_updatetestutility.html

The software update test utility can be downloaded in advance of the update process, or it can be downloaded and run directly during the software update, as guided by the update wizard. You can run the utility multiple times on the same system to perform a readiness check-in preparation for a software update.

The installation and use of this utility is non disruptive, and it does not require a restart of any node. Therefore, host I/O is not interrupted. The utility is only installed on the current configuration node.

System administrators must continue to check whether the version of code that they plan to install is the latest version.

Updating the software automatically by using the GUI

Complete the following steps to automatically update the node canister software by using the GUI:

1. Browse to **Settings** → **System** → **Update System** and select **Test and Update**, as shown in Figure 12-31.




Figure 12-31 Update system panel

Alternatively, you can run only the test utility by selecting **Test Only**.

2. Select the test utility and update package files by clicking the folder icons, as shown in Figure 12-32. The code levels are entered automatically.




Figure 12-32 File selection

Alternatively, for the **Test Only** option, upload only the test utility and enter the code level manually.

3. Select **Automatic update** and click **Next** to come to the next question regarding paused update, as shown in Figure 12-34 on page 654. The **Automatic update** option is the default and advised choice.

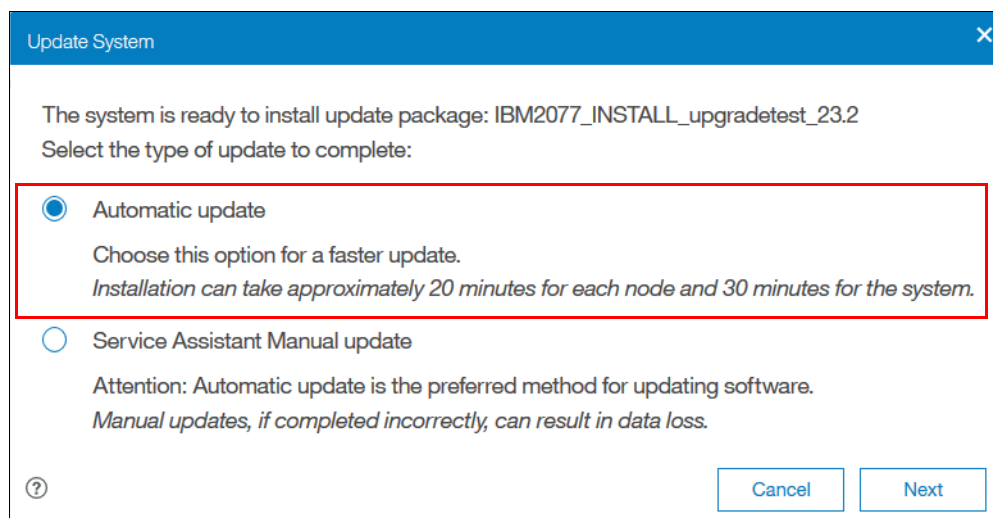


Figure 12-33 Automatic update selection

Shown in Figure 12-34 on page 654 you can choose if you want to pause the update or not. Default is **Fully automatic**. Click **Finish** to start the update

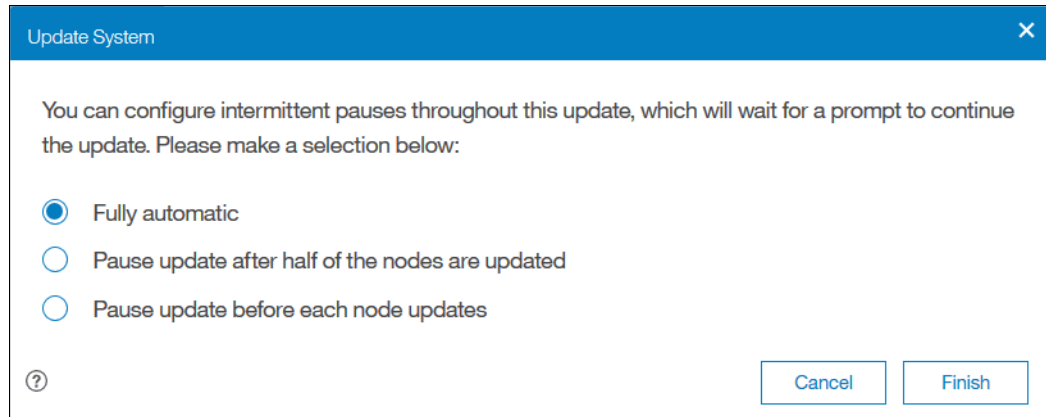


Figure 12-34 Fully automatic

4. Wait for the test utility and update package to upload to the system, as shown in Figure 12-35.




Figure 12-35 File upload

5. After the files upload, the test utility is automatically run, as shown in Figure 12-36. The test utility verifies that no issues exist with the current system environment, such as failed components and drive firmware that is not at the latest level.




Figure 12-36 State while the test utility runs

If the test utility discovers any warnings or errors, a window opens to inform the user, as shown in Figure 12-37 on page 655. Click **Read more** to get more information.

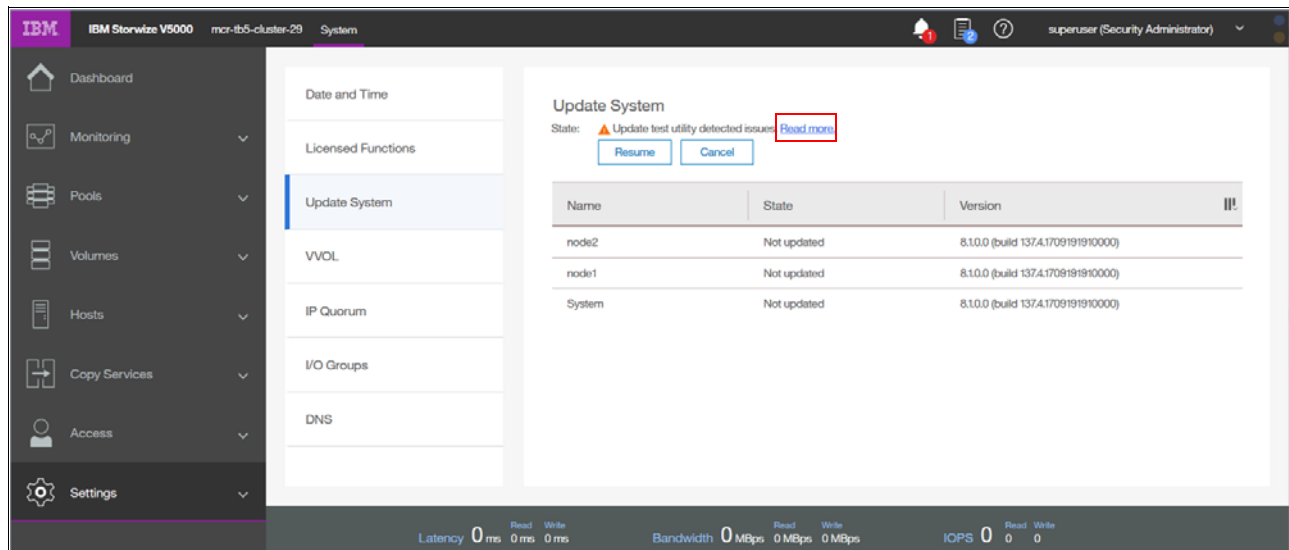


Figure 12-37 Warning about the issues that were detected

Figure 12-38 shows that in this example the test utility identified one warning.

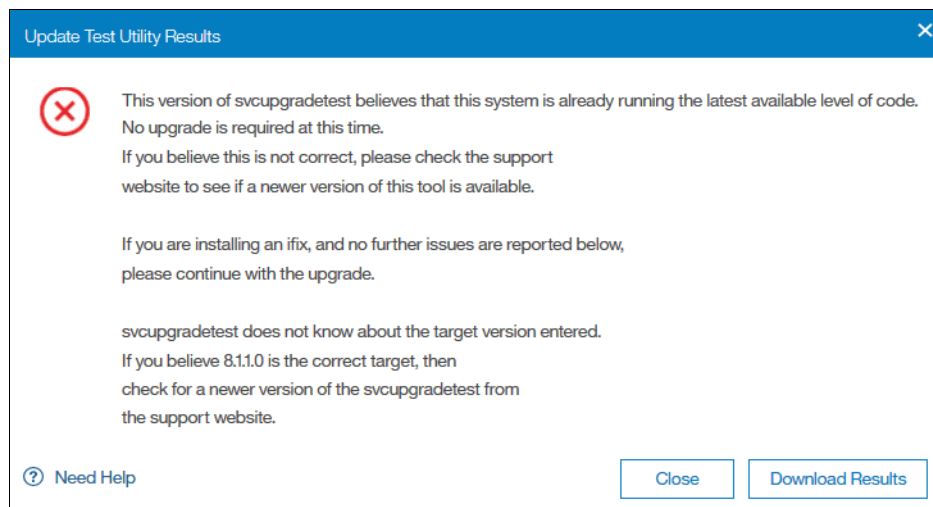


Figure 12-38 Test utility results

Warnings do not prevent the software update from continuing, even if the recommended procedure is to fix all issues before you proceed.

Close the window and select either **Resume** or **Cancel**, as shown in Figure 12-39 on page 656. Clicking **Resume** continues the software update. Clicking **Cancel** cancels the software update so that the user can correct any issues.

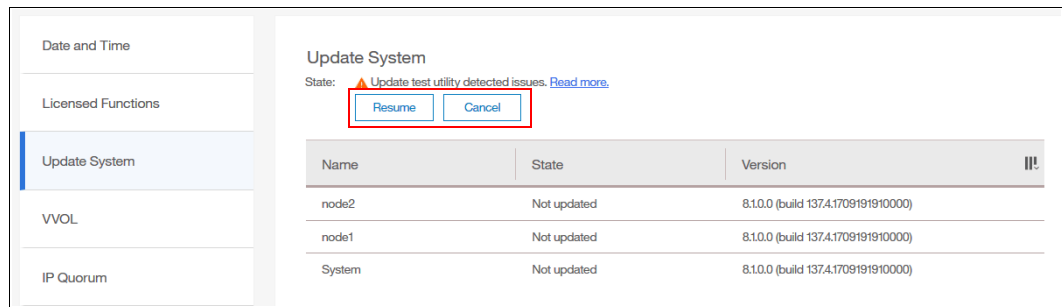


Figure 12-39 State after you run the test utility

Selecting **Resume** prompts the user to confirm the action, as shown in Figure 12-40.

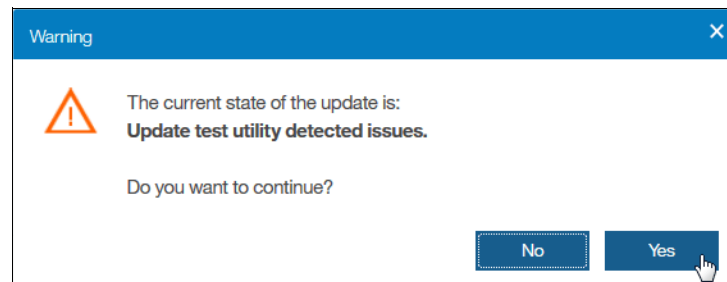


Figure 12-40 Resume confirmation window

- Wait for each node to be updated and rebooted, one at a time until the update process is complete. The GUI displays the overall progress of the update and the current state of each node, as shown in Figure 12-41.




Figure 12-41 Automatic update progress

During the update process, a node fails over and you can temporarily lose connection to the GUI. After this situation happens, a warning is displayed, as shown in Figure 12-42 on page 657. Select **Yes**.

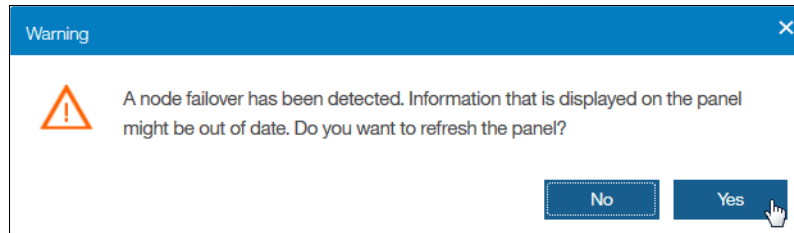


Figure 12-42 Configuration node failover warning

Updating the software manually by using the GUI and SAT

Important: We advise that you update the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 automatically by following the update wizard. If a manual update is used, ensure that you do not skip any steps.

Complete the following steps to manually update the software by using the GUI and Service Assistant Tool (SAT):

1. Browse to **Settings** → **System** → **Update System** and select **Update and Test**, as shown in Figure 12-43.




Figure 12-43 Update system panel

Alternatively, you can run the test utility by selecting **Test Only**.

2. Select the test utility and update package files by clicking the folder icons, as shown in Figure 12-44. The code levels are entered automatically.




Figure 12-44 File selection

Alternatively, for the **Test Only** option, upload only the test utility and enter the code level manually.

3. Select **Service Assistant Manual update** and click **Finish**, as shown in Figure 12-45.

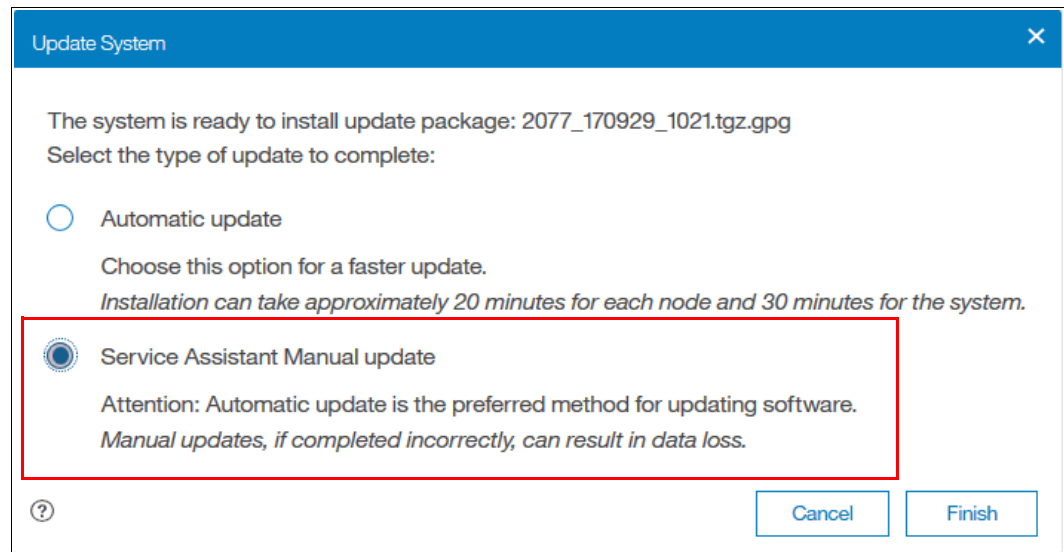


Figure 12-45 Manual update selection

4. Wait for the test utility and update package to upload to the system, as shown in Figure 12-46.

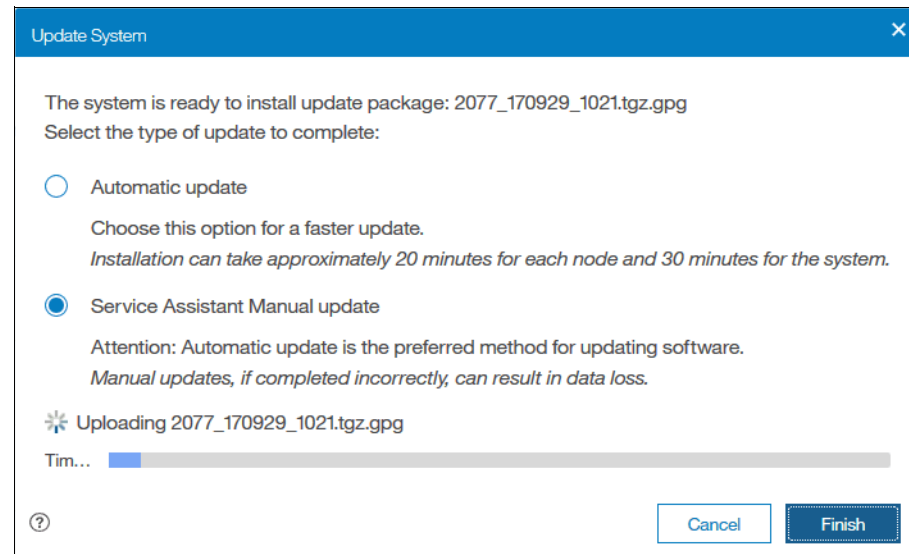


Figure 12-46 File upload

5. After the files upload, the test utility is automatically run, as shown in Figure 12-47 on page 659. The test utility verifies that no issues exist with the current system environment, such as failed components and drive firmware that is not at the latest level.

Update System

State: Running update test utility.

Resume

Cancel

Name	State	Version
node2	Not updated	7.6.1.0 (build 124.4.1602111021000)
node1	Not updated	7.6.1.0 (build 124.4.1602111021000)
System	Not updated	7.6.1.0 (build 124.4.1602111021000)

Figure 12-47 State while the test utility run

If the utility identifies no issues, the system is ready for the user to initiate the manual upgrade, as shown in Figure 12-48.

Date and Time

Licensed Functions

Update System

VVOL

IP Quorum

Update System

State: Ready to start manual upgrade. Open the Service Assistant and select Update Manually.

Resume

Cancel

Name	State	Version
node2	Not updated	8.1.0.1 (build 137.4.1710101636000)
node1	Not updated	8.1.0.1 (build 137.4.1710101636000)
System	Not updated	8.1.0.1 (build 137.4.1710101636000)

Figure 12-48 State while you wait for the manual upgrade to start

- Choose a node to update. Non-configuration nodes must be updated first. Update the configuration node last. Browse to **Monitoring** → **System** and hover over the canisters to confirm the nodes that are the non-configuration nodes, as shown in Figure 12-49 on page 660.

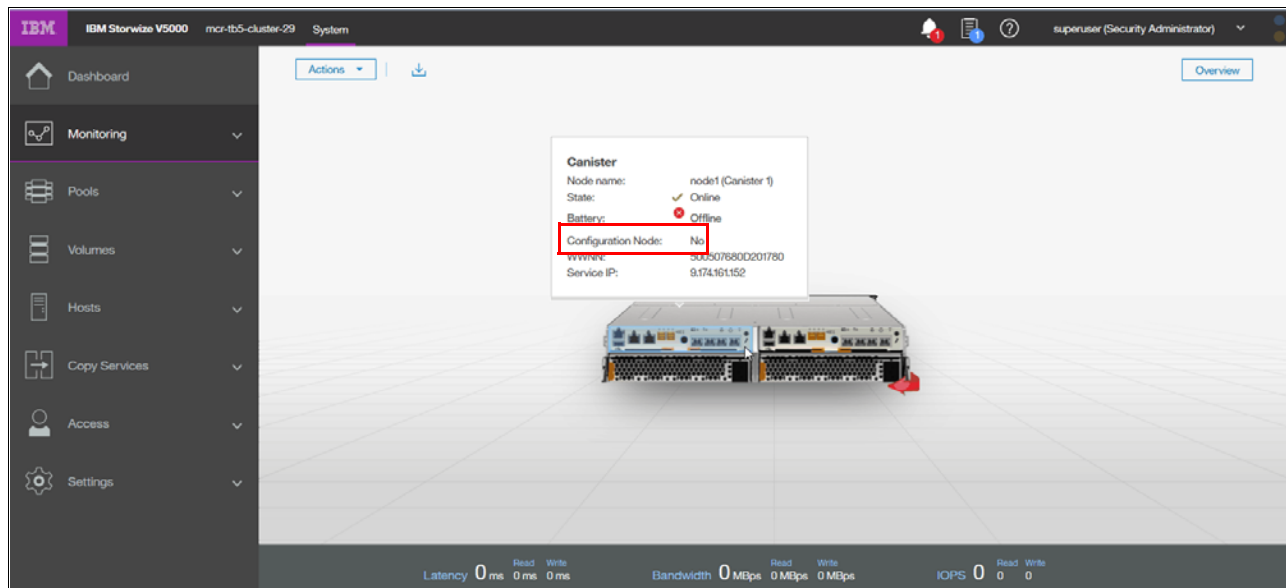


Figure 12-49 Checking the configuration node status

7. Right-click the canister that contains the node that you want to update and select **Remove**, as shown in Figure 12-50.

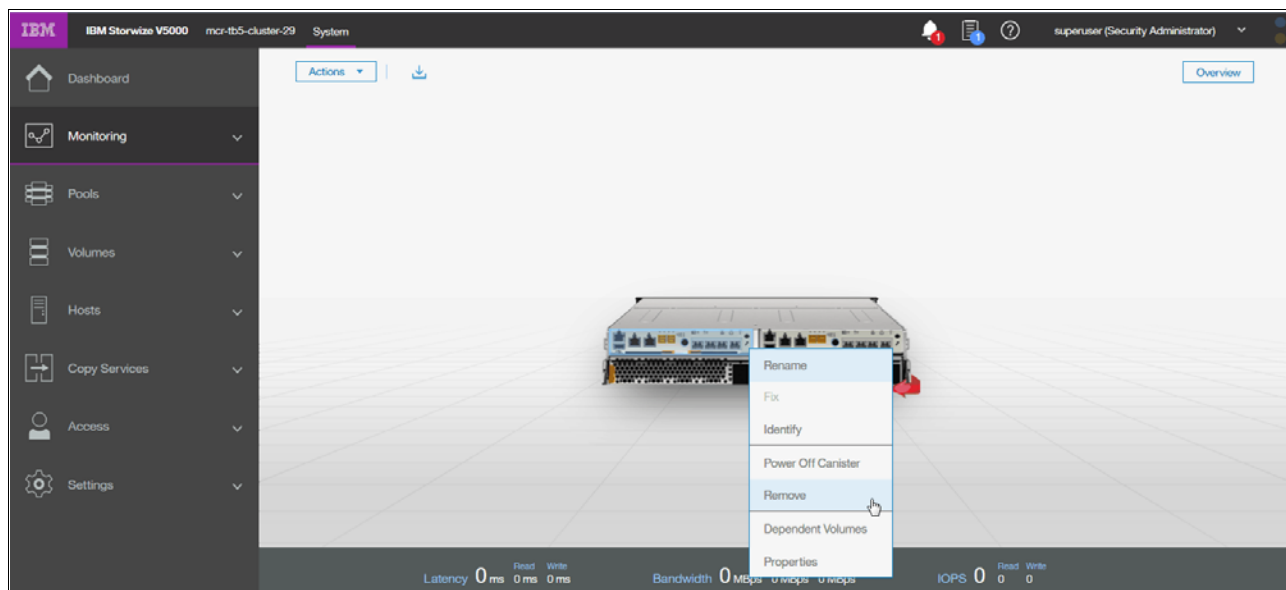


Figure 12-50 Removing a node canister

Important: Ensure that you select the non-configuration nodes first.

8. A warning message appears to ask whether you want to remove the node, as shown in Figure 12-51 on page 661. Click **Yes**.

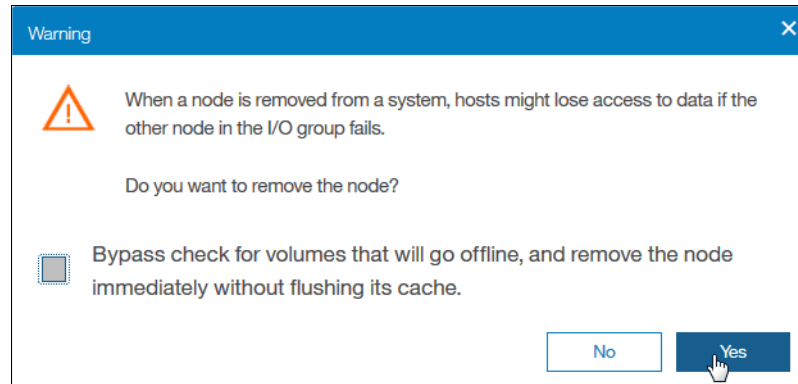


Figure 12-51 Node removal confirmation window

The non-configuration node is removed from the management GUI Update System panel and is shown as Unconfigured when you hover over the node after you select **Monitoring** → **System**.

- Open the Service Assistant Tool for the node that you removed. Enter the Service IP Address followed by /service into a browser window. Without /service the browser will open the associated GUI to this service IP. No HTTP:// or HTTPS:// is needed.

Example: 172.163.18.34/service

- In the Service Assistant Tool, ensure that the node that is ready for update is selected. The node can be in the Service status, display a 690 error, and show no available cluster information, as shown in Figure 12-52.




Figure 12-52 Node to update in the Service Assistant Tool

- In the Service Assistant Tool, select **Update Manually**, and choose the required node canister software upgrade file, as shown in Figure 12-53 on page 662.

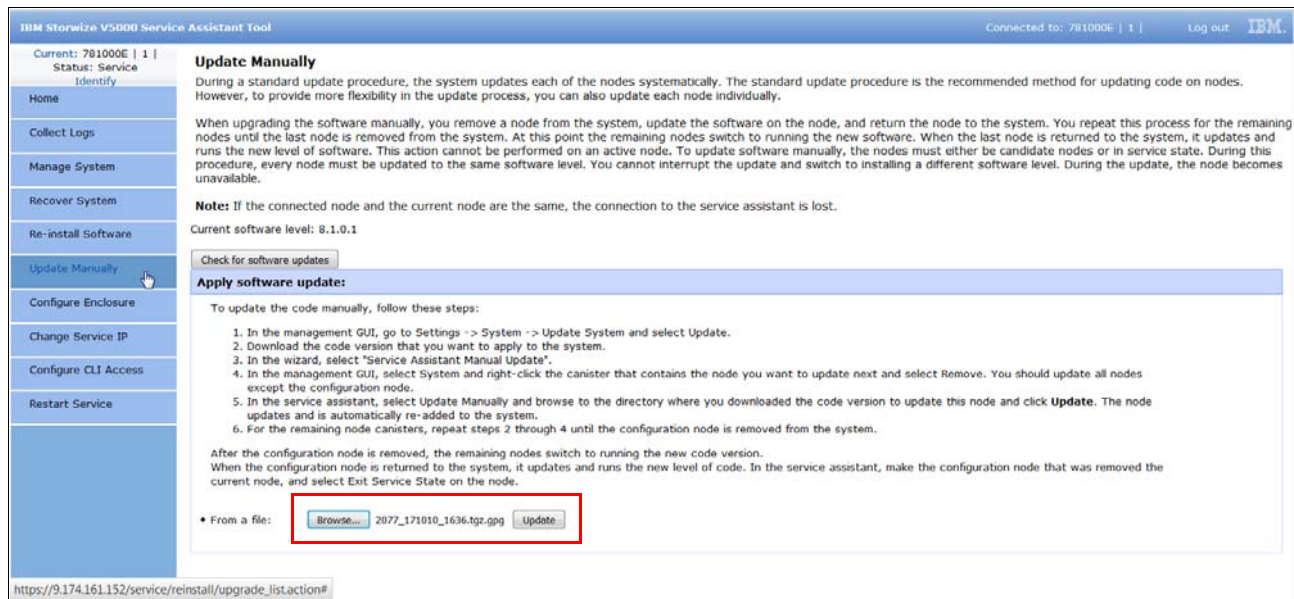


Figure 12-53 Starting the update in the Service Assistant Tool

12. Click **Update** to start the update process on the first node and wait for the node to finish updating.

Non-configuration nodes can be reintroduced automatically into the system after the update finishes. Updating and adding the node again can last 20 - 40 minutes.

The management GUI shows the progress of the update, as shown in Figure 12-54.

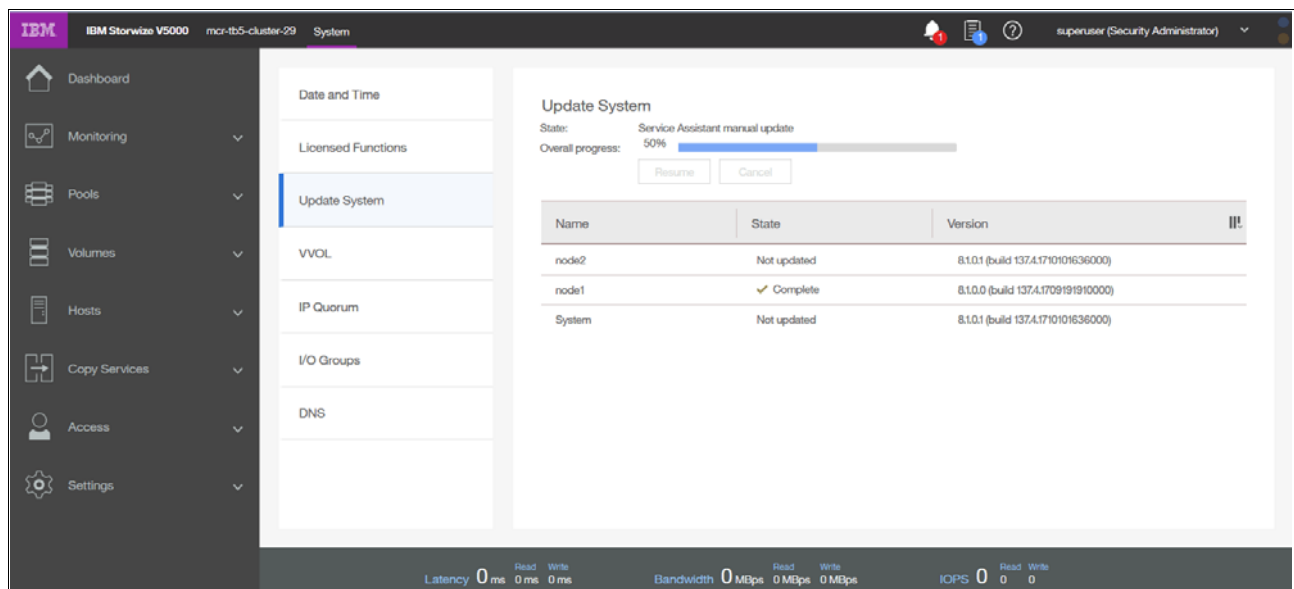


Figure 12-54 Manual update progress

13. Repeat steps 7 - 12 for the remaining nodes, leaving the configuration node until last.

14. After you remove the configuration node from the cluster, you are asked whether you want to refresh the panel, as shown in Figure 12-55 on page 663. Select **Yes**.

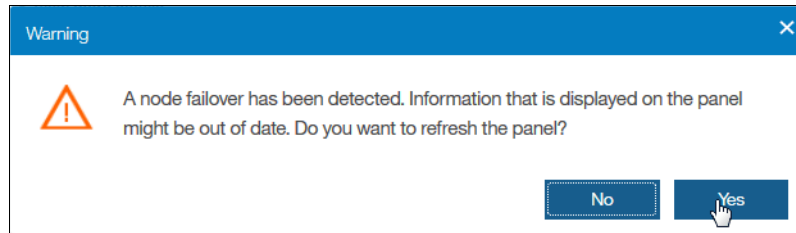


Figure 12-55 Configuration node failover warning

Important: The configuration node remains in the Service state when it is added to the cluster again. Therefore, you need to exit the Service state manually.

15. To exit the Service state, browse to the Home panel of the Service Assistant Tool and open the Actions menu. Select **Exit Service State** and click **GO**, as shown in Figure 12-56.

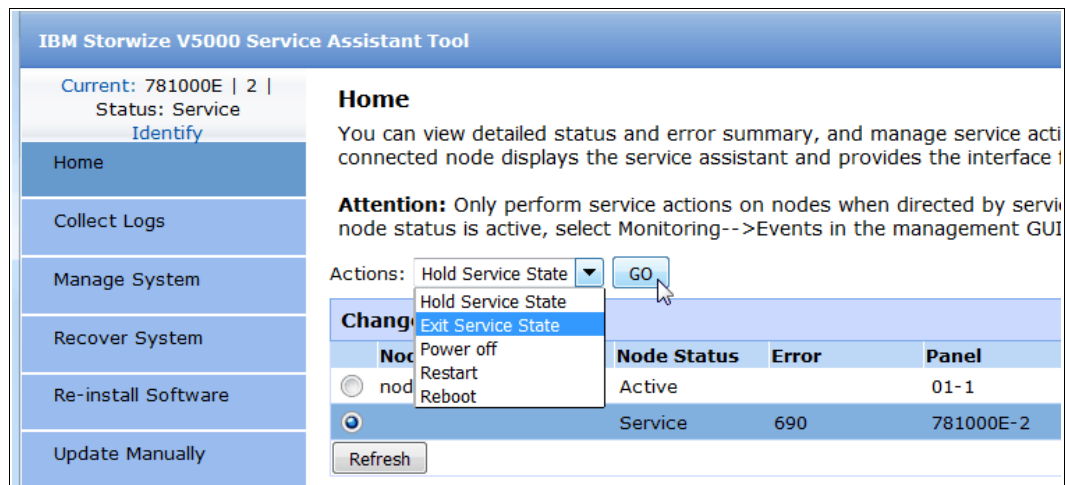


Figure 12-56 Exiting the Service state in the Service Assistant Tool

12.4.2 Updating the drive firmware

Drive firmware can be updated for all drives at the same time or individually.

To get the latest drive update package, go to the Supported Drive Types and Firmware Levels for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 web page:

<https://datacentersupport.lenovo.com/us/en/solutions/ht503947>

Note: Find the download link for the actual drive firmware at the bottom of the Web page.

Updating the firmware on individual drives

To update an individual drive, navigate to **Pools** → **Internal Storage**, right-click the drive to update, and select **Upgrade** from the Actions menu, as shown in Figure 12-57 on page 664.

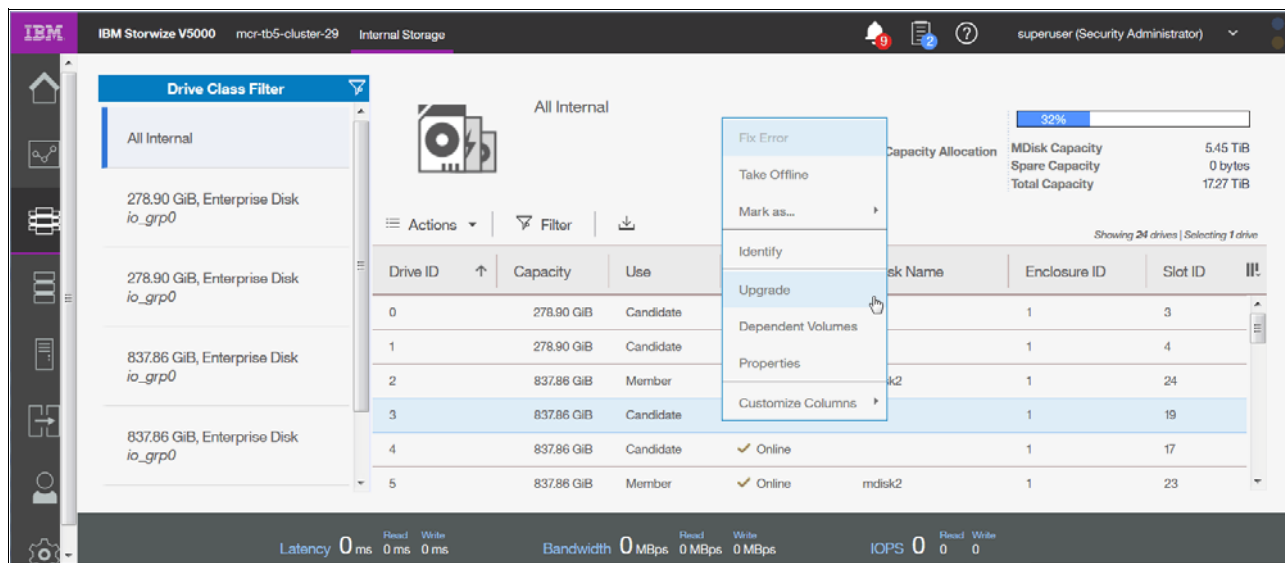


Figure 12-57 Individual drive update

Select the upgrade package, which was downloaded from the Lenovo Support site, by clicking the folder icon, and click **Upgrade**, as shown in Figure 12-58.



Figure 12-58 Individual drive update file selection

The drive firmware update takes about 2 - 3 minutes for each drive.

To verify the new firmware level, right-click the drive and select **Properties**, as shown in Figure 12-59 on page 665.

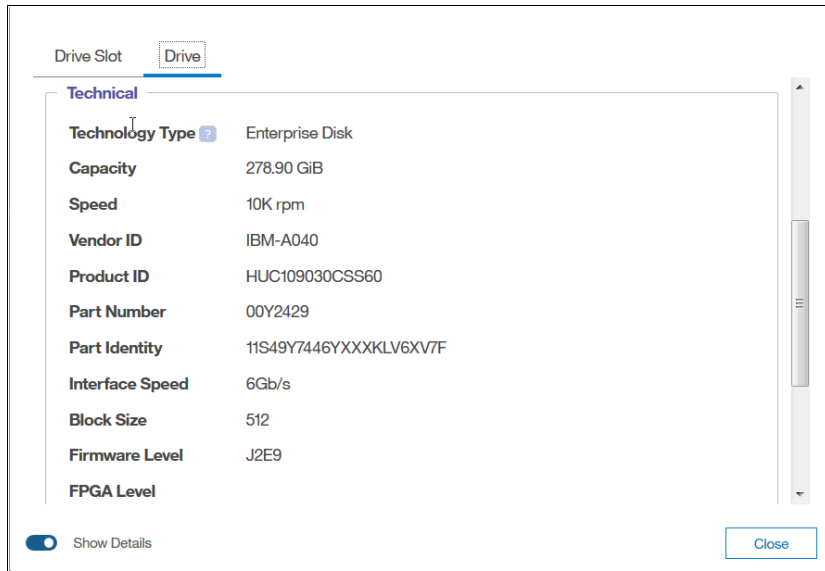


Figure 12-59 Individual drive update result

Updating the firmware on all drives

Here we show how to use the management GUI to update all of the drives in a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030:

- Go to **Pools** → **Internal Storage**

Figure 12-60 shows how to update all drives through the Actions menu in the Internal Storage panel. Under Drive Class Filter, click **All Internal**. In the Actions menu, click **Upgrade All**.

Note: If any drives are selected, the Actions menu displays actions for the selected drives and the Upgrade All option does not appear. If a drive is selected, deselect it by holding down the Ctrl key and clicking the drive.

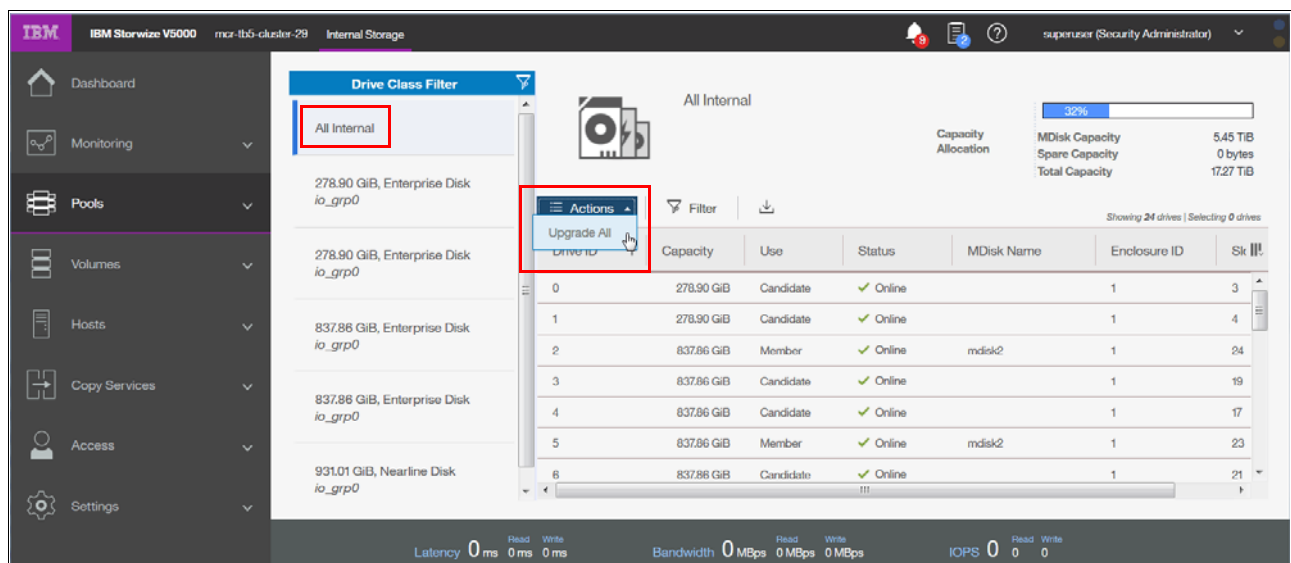


Figure 12-60 Update of multiple drives in the Internal Storage panel

After you initiate the drive upgrade process by either of the previous two options, the panel in Figure 12-61 is displayed. Select the drive upgrade package, which was downloaded from the Lenovo Support site, by clicking the folder icon, and click **Upgrade**. You can also override newer versions of firmware when you checkmark the option **Install the firmware even if the drive is running a newer version and if directed by the support center**.

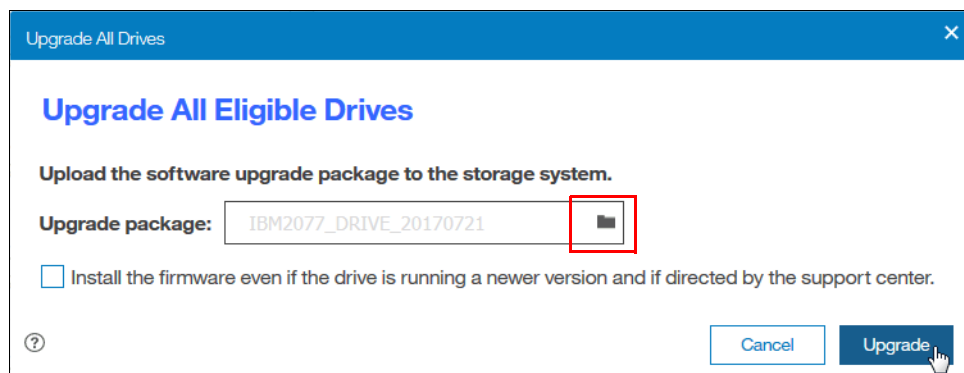


Figure 12-61 Upload the software upgrade package for multiple drives

All drives that require an update can now be updated.

12.5 Monitoring

Any issue that is reported by your Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems must be fixed as soon as possible. Therefore, it is important to configure the system to send automatic notifications when a new event is reported. You can select the type of event for which you want to be notified. For example, you can restrict notifications to only events that require immediate action.

Several event notification mechanisms are available:

- | | |
|---------------|---|
| Email | Email notifications can be configured to send emails to one or more email addresses. With this mechanism, individuals can receive notifications wherever they have email access, including mobile devices. |
| SNMP | SNMP notifications can be configured to send a Simple Network Management Protocol (SNMP) traps report to a data center management system that consolidates SNMP reports from multiple systems. With this mechanism, you can monitor your data center from a single workstation. |
| Syslog | Syslog notifications can be configured to send a syslog report to a data center management system that consolidates syslog reports from multiple systems. With this mechanism, you can monitor your data center from a single location. |

If your system is within warranty, or you have a hardware maintenance agreement, configure your Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems to send email events directly to IBM if an issue that requires hardware replacement is detected. This mechanism is known as *Call Home*. When an event is received, IBM automatically opens a problem report and, if appropriate, contacts you to verify whether replacement parts are required.

Important: If you set up Call Home to the IBM Support Center, ensure that the contact details that you configured are correct and kept up-to-date when personnel changes.

12.5.1 Email notifications and Call Home

The Call Home function of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 use the email notification mechanism to send emails to the specific IBM Support Center. You enable Call Home by configuring email notifications. Then, you can optionally add other email addresses to notify.

To configure Call Home and other optional email addresses, complete the following steps:

1. Browse to **Settings** → **Notifications** → **Email** and select **Enable Notifications**, as shown in Figure 12-62.




Figure 12-62 Enabling email notifications

For the correct functionality of email notifications, ensure that Simple Mail Transfer Protocol (SMTP) is enabled on the management network and not, for example, blocked by firewalls.

If Email Notification is not enabled, you will get a periodically warning like shown in Figure 12-63.




Figure 12-63 Configure Call Home info

2. Configure the SMTP servers. You can add several servers by clicking the plus (+) sign, as shown in Figure 12-64.




Figure 12-64 Email Servers

3. Figure 12-65 on page 668 shows the entry for Call Home. This Email Address is given, and can't be changed.

Call Home

Email Address

callhome0@de.ibm.com

☒ Error Events ☒ Inventory

Test

Figure 12-65 Call Home

4. You can add several recipients to receive notifications. Press the + sign to add a new Email Address. Figure 12-66 shows one entry.

Email Users

Email Address	Notifications				
	Error	Warning	Info	Inventory	
admin@testcenter.com	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	+ -

Figure 12-66 Email Users

5. It is very important to add an Email contact, who is responsible for this Storage System. Provide the contact information of the system owner, who can be contacted by the IBM Support Center when necessary, Figure 12-67 shows such an entry. Ensure that you always keep this information up-to-date.

Email Contact

* Contact Name

James

* Email Reply Address

jamesr@ok.ibm.com

* Telephone (Primary)

+44-181-9676899

Telephone (Alternate)

* Required

Figure 12-67 Email Contact

6. Also the System Location is important. This Information will be used by the support personnel to send the Support Representative to the failing system. If there is only a minor problem this Info will be used to send the CRU parts to the given address. Figure 12-68 on page 669 shows how the System Location panel should be filled out. Ensure that you always keep this information up-to-date.

System Location

* Company Name: IBM

* Street Address: 123 MayHouse

* City: London

* State or Province: XX

* Postal Code: 123456

* Comment: third floor

* Country or Region: United Kingdom

* Required

Figure 12-68 System Location

- You can include an inventory file into your Email to check the actual inventory of your system. Figure 12-69 shows you the location where you can set the checkmark to indicate you want to receive inventory details. The emails include an inventory report that describes the system hardware and critical configuration information. Object names and other information, such as IP addresses, are not sent. Based on the information that is received, IBM can inform you whether the hardware or software that you are using requires an upgrade because of a known issue.

Inventory Service

☒ Enable inventory reporting

Reporting Interval: 7 days

Figure 12-69 Inventory details

- Click **Save**.
- Select as shown in Figure 12-70 **Edit** → **Call Home** → **Test** to test the Call Home function.

Call Home

Send a test email to the email user. The email user must already exist and this panel must be in edit mode to enable this function.

Test

Figure 12-70 Test Call Home.

Disabling and enabling notifications

Email notifications can be temporarily or permanently disabled at any time, as shown in Figure 12-71 on page 670. Disabling email notifications is a preferred practice when you run maintenance tasks, such as upgrading code or replacing parts. After the maintenance operation, remember to re-enable the email notification function.




Figure 12-71 Disabling email notifications

The same results can be achieved by using the CLI and entering the **svctask stopmail** and **svctask startmail** commands.

12.6 Audit log

The *audit log* is useful when you analyze past configuration events, especially when you try to determine, for example, how a volume ended up being shared by two hosts, or why the volume was overwritten. The audit log is included in the support package to aid in problem determination.

The audit log tracks action commands that are issued through the CLI or the management GUI. It provides the following entries:

- ▶ Name of the user who issued the action command
- ▶ Name of the actionable command
- ▶ Time stamp of when the actionable command was issued on the configuration node
- ▶ Parameters that were issued with the actionable command

Failed commands and view commands are not logged in the audit log. Certain service commands are not logged either. The **svconfig backup**, **cpdumps**, and **ping** service commands are not logged.

To access the audit log by using the GUI, browse to **Access** → **Audit Log**, as shown in Figure 12-72.

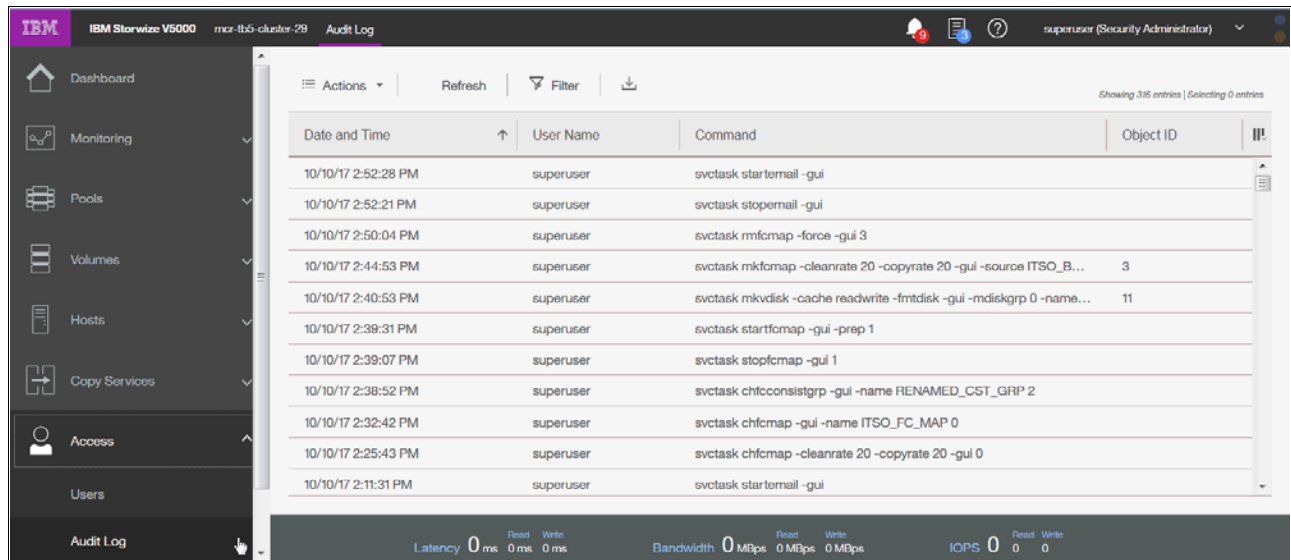


Figure 12-72 Audit log panel

Right-clicking any column header opens the option menu in which you can select columns that are shown or hidden. It is also possible to click the Column icon on the far right of the column headers to open the option menu.

Figure 12-73 shows all of the possible columns that can be displayed in the audit log view.

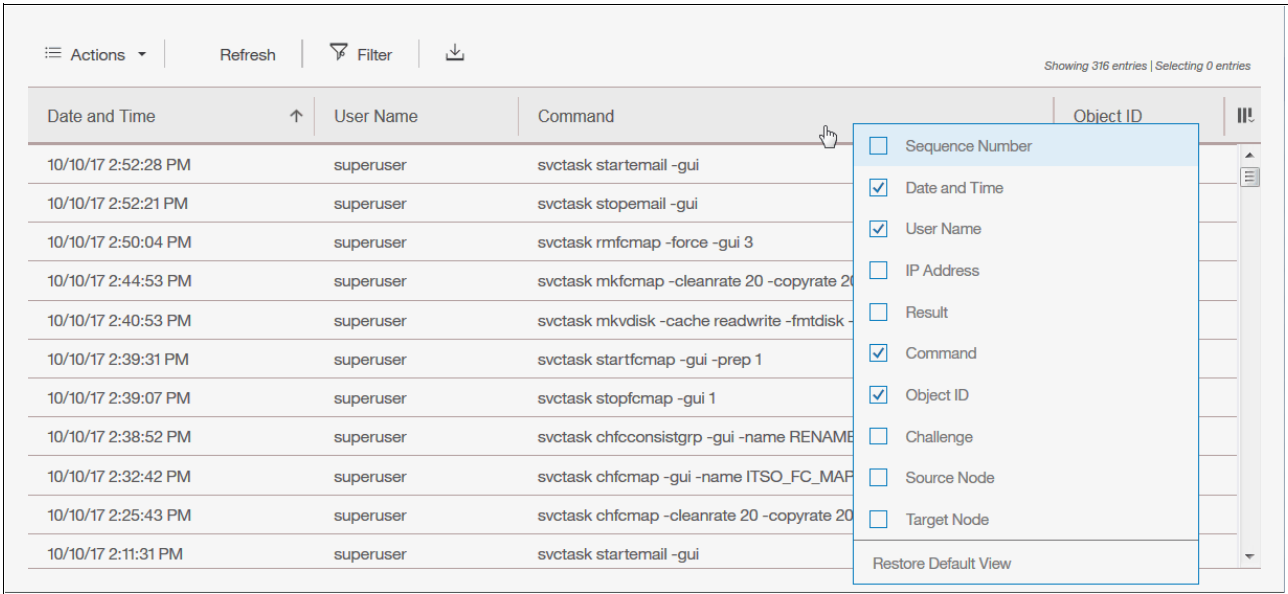


Figure 12-73 Possible audit log columns

12.7 Event log

Whenever a significant change in the status of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 is detected, an event is submitted to the *event log*. All events are classified as *alerts* or *messages*.

An alert is logged when the event requires action. Certain alerts have an associated error code that defines the service action that is required. The service actions are automated through the fix procedures. If the alert does not have an error code, the alert represents an unexpected change in the state. This situation must be investigated to see whether it is expected or represents a failure. Investigate an alert and resolve it when it is reported.

A message is logged when a change that is expected is reported, for instance, a FlashCopy operation completes.

To check the event log, browse to **Monitoring** → **Events**, as shown in Figure 12-74 on page 672.

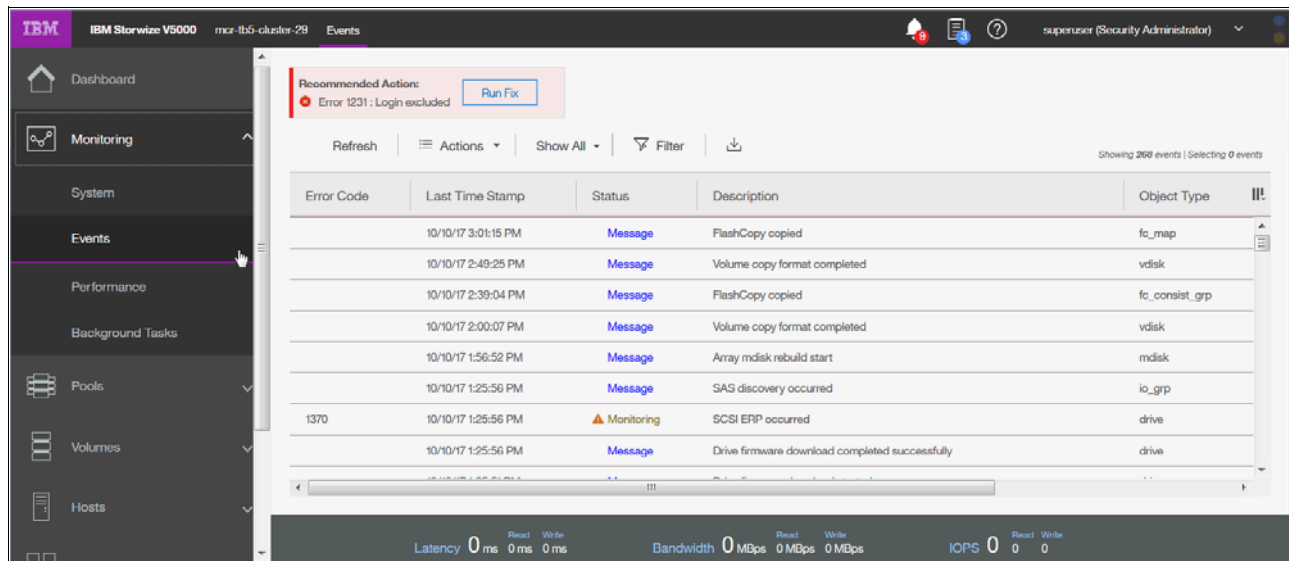


Figure 12-74 Event log

12.7.1 Managing the event log

The event log features a size limit. After the event log is full, newer entries replace the older entries, which are not required. To avoid a repeated event that fills the event log, certain records in the event log refer to multiple occurrences of the same event. When event log entries are coalesced in this way, the time stamp of the first occurrence of the problem and the time stamp the last occurrence of the problem are saved in the log entry. A count of the number of times that the error condition occurred is also saved in the log entry. Other data refers to the last occurrence of the event.

Event log panel columns

Right-clicking any column header opens the option menu in which you can select columns that are shown or hidden. It is also possible to click the Column icon on the far right of the column headers to open the option menu.

Figure 12-75 on page 673 shows all of the possible columns that can be displayed in the error log view.

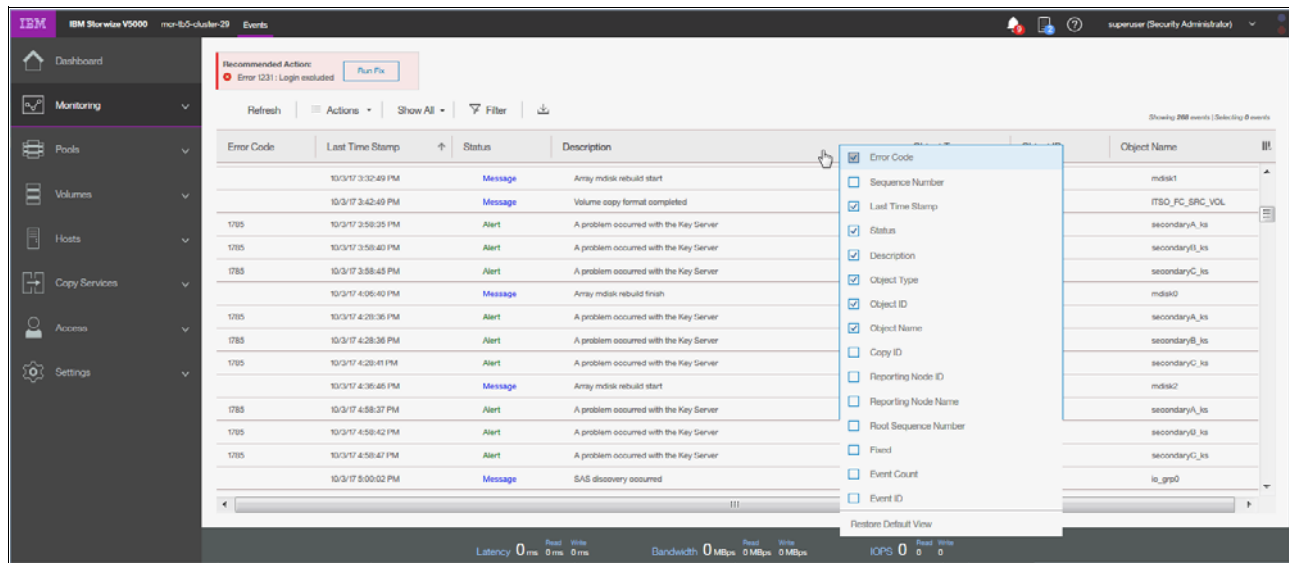


Figure 12-75 Possible event log columns

Event log filter options

The event log can be filtered by using the options that are shown in Figure 12-76.

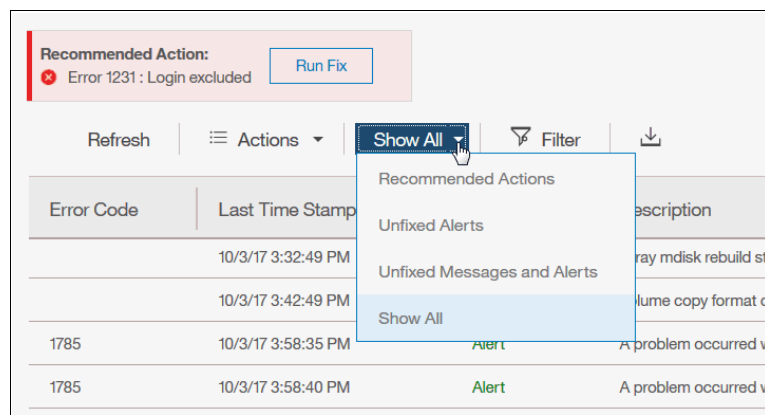


Figure 12-76 Event log filter options

Each option is described:

► Recommended Actions (default)

Only events with Recommended Actions (Status Alert) are displayed. For each problem that is selected, you can:

- Run a fix procedure
- View the properties

► Unfixed Alerts

Displays only the alerts that are not fixed. For each entry that is selected, you can:

- Run a fix procedure on any alert with an error code
- Mark an event as fixed
- Filter the entries to show them by specific minutes, hours, or dates
- Reset the date filter

- View the properties.
- **Unfixed Messages and Alerts**
This option lists unfixed events. This option is useful to find events that must be handled, but no actions are required or recommended. For each entry that is selected, you can:
 - Run a fix procedure on any alert with an error code
 - Mark an event as fixed
 - Filter the entries to show them by specific minutes, hours, or dates
 - Reset the date filter
 - View the properties
- **Show All**
This option lists all available events. For each entry that is selected, you can:
 - Run a fix procedure on any alert with an error code
 - Mark an event as fixed
 - Filter the entries to show them by specific minutes, hours, or dates
 - Reset the date filter
 - View the properties

Some events require a certain number of occurrences in 25 hours before they are displayed as unfixed. If they do not reach this threshold in 25 hours, they are flagged as expired. Monitoring events are below the coalesce threshold and are usually transient.

You can also sort events by time or error code. When you sort by error code, the most serious events (those with the lowest numbers) are displayed first. You can select any event that is listed and select **Actions** → **Properties** to view details about the event.

Important: Check for this filter option if no event is listed. Events might exist that are not associated with recommended actions.

Figure 12-77 shows an event log with no items when the Recommended Actions filter was selected, which does not necessarily mean that the event log is clear. To check whether the log is clear, click **Show All**.

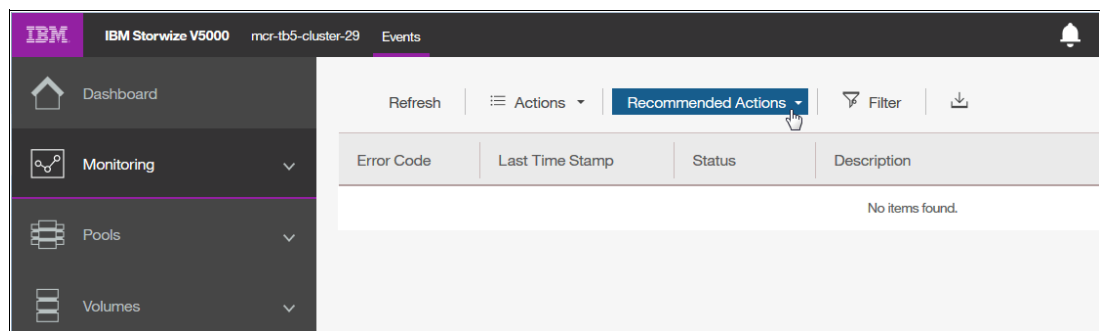


Figure 12-77 Event log with no recommended actions

Actions on a single event

Right-clicking a single event gives options that might be used for that specific event, as shown in Figure 12-78 on page 675.

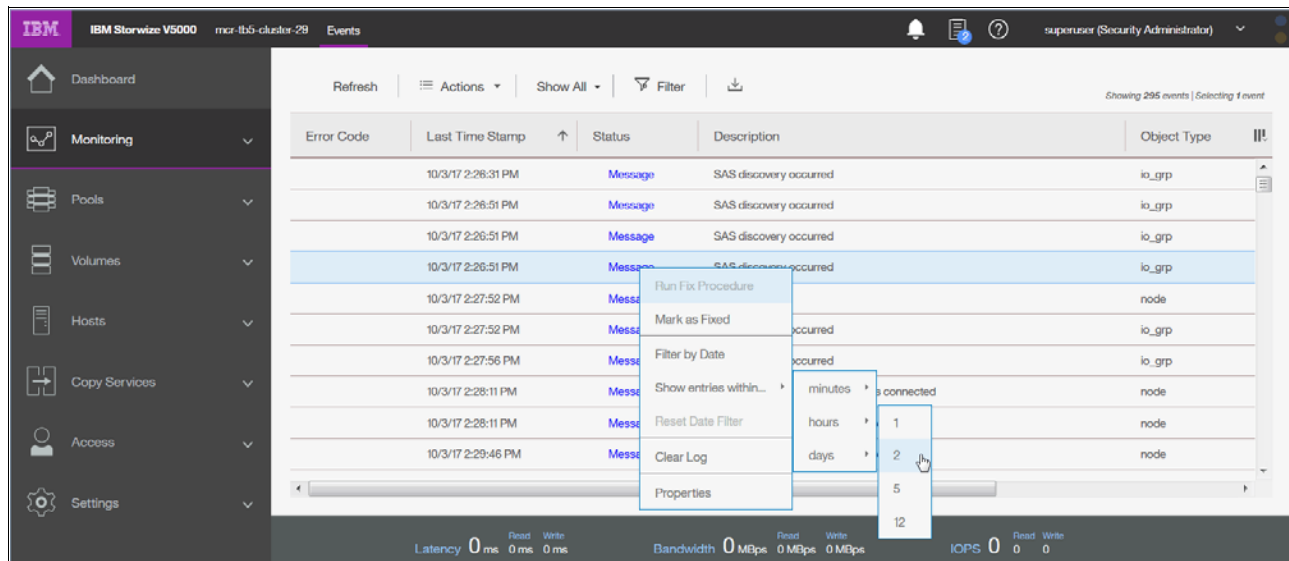


Figure 12-78 Possible actions on a single event

Each option is described:

- Run Fix Procedure

This option starts the fix procedure for this specific event. You can start a fix procedure even if the procedure is not the recommended next action. However, we advise that you fix the error with the highest priority first.

- Mark as Fixed

This option marks this specific event as fixed. Message events must be marked as fixed to stop them from showing in the event log.

- Filter by Date

This option limits the event log entries to the events that occurred between an interval that is defined by the user.

- Show entries within (minutes/hours/days)

This option limits the event log entries to the events that occurred within the last period:

- 1, 5, 10, 15, 30, or 45 minutes
- 1, 2, 5, or 12 hours
- 1, 4, 7, 15, or 30 days

- Reset Date Filter

This option clears the Filter by Date.

- Clear Log

This option clears the complete event log, even if only one event was selected.

Important: These actions cannot be undone and might prevent the system from being analyzed when severe problems occur.

- Properties

This option provides more information for the selected event that is shown in the list.

Recommended actions

A fix procedure invokes a wizard that is known as a Directed Maintenance Procedure (DMP) that helps to troubleshoot and correct the cause of an error. Certain DMPs reconfigure the system based on your responses, ensure that actions are carried out in the correct sequence, and prevent or mitigate the loss of data. For this reason, you must always run the fix procedure to fix an error, even if the fix might seem obvious.

To run the fix procedure for the error with the highest priority, go to the Recommended Action panel at the top of the Events page and click **Run Fix**, as shown in Figure 12-79. When you fix higher-priority events first, the system often can automatically mark lower-priority events as fixed.




Figure 12-79 Next recommended action

12.7.2 Alert handling and recommended actions

All events that are in Alert status require attention. Alerts are listed in priority order. Alerts must be fixed sequentially by using the available fix procedures.

Example: Array mdisk not protected by sufficient spares

For example, look at an error that was raised by taking a drive offline in an array with redundancy of one.

This example can show how faults are represented in the error log, how information about the fault can be gathered, and how the Recommended Action (DMP) can be used to fix the error:

► Detecting the alert

The Health Status indicator shows a red alert. The Status Alerts indicator (on top of the GUI) shows one alert. Click the alert to retrieve the specific information, as shown in Figure 12-80.




Figure 12-80 Status alert for an individual entry

Review the event log for more information.

- Gathering additional information

More details about the event are available by clicking the event and selecting **Details**. This information might help you fix a problem or analyze a root cause. Figure 12-81 shows the properties for the previous event.

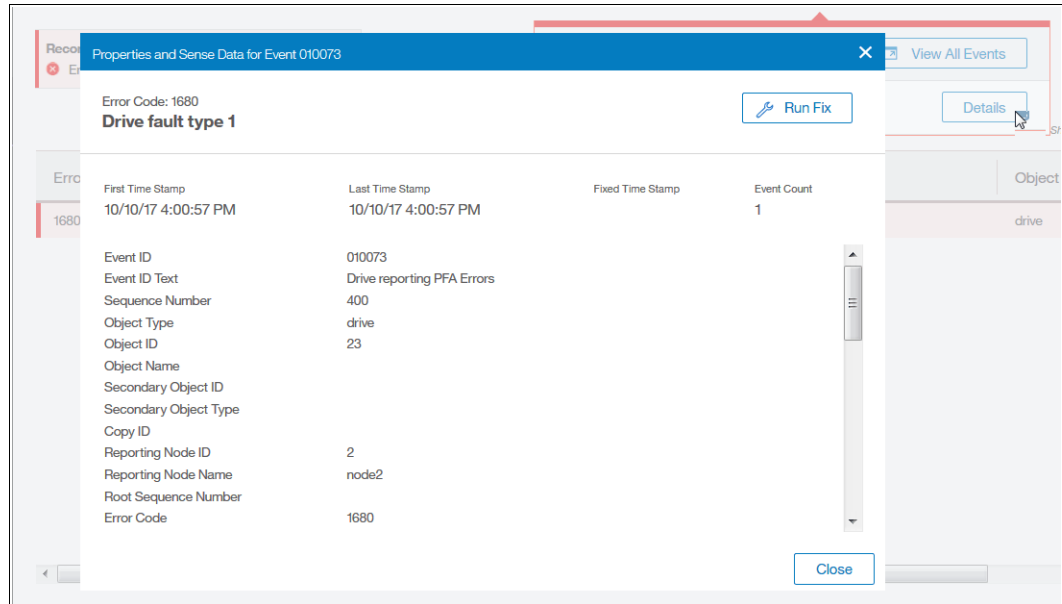


Figure 12-81 Alert properties

- Run the Recommended Action (DMP)

We highly advise that you use the DMP to fix any alerts. You can miss tasks that are running in the background when you bypass the DMP. Not all alerts have available DMPs.

Figure 12-82 shows how to start the DMP by selecting **Run Fix** at the top of the window. This option always runs the recommended action.

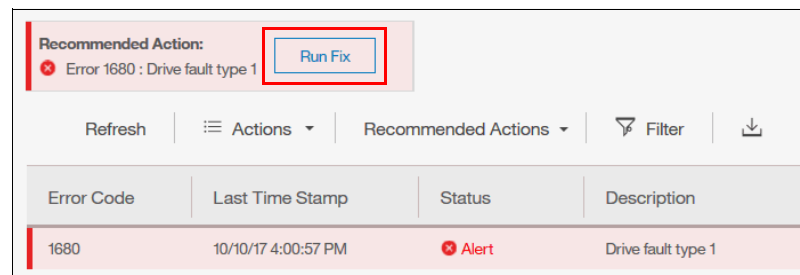


Figure 12-82 Starting the DMP (first option)

Figure 12-83 on page 678 shows how to start the DMP by right-clicking the alert record and selecting **Run Fix Procedure**. You can use this option to run a fix procedure that might not be the recommended action.

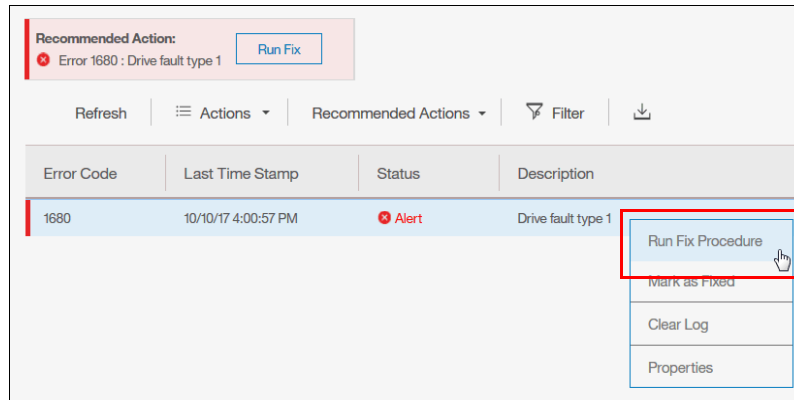


Figure 12-83 Starting the DMP (second option)

The steps and panels of a DMP are specific to the error. When all of the steps of the DMP are processed successfully, the recommended action is complete and the problem is fixed usually. Figure 12-84 shows that the Health Status changed to green and both the Status Alerts indicator and the Recommended Action box disappeared, implying that no more actions must be taken.

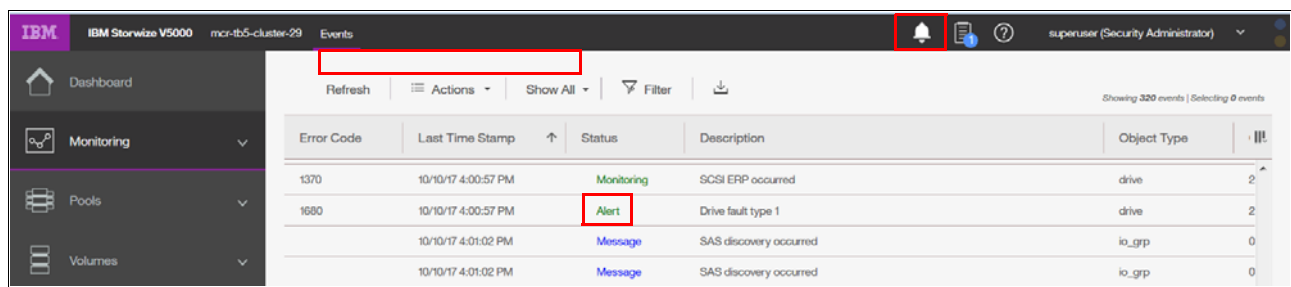


Figure 12-84 Event log with no outstanding recommended action

Handling multiple alerts

Figure 12-85 shows the event log with multiple alerts.

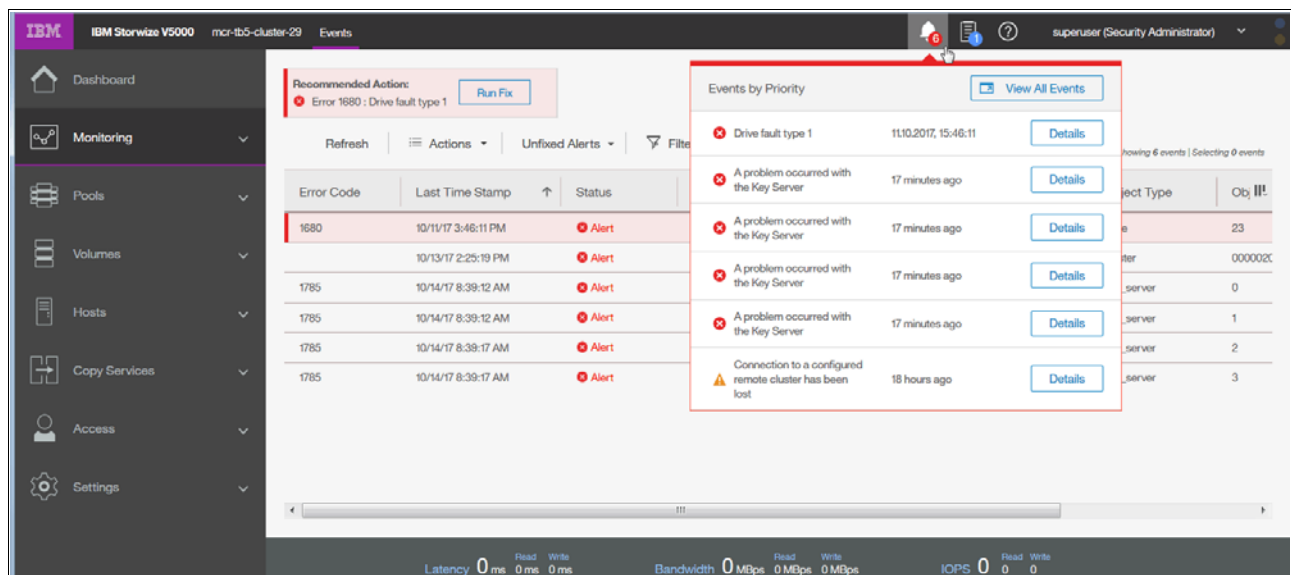


Figure 12-85 Multiple alert events that are displayed in the event log

The Recommended Action function orders the alerts by severity and displays the events with the highest severity first. If multiple events have the same severity, they are ordered by date and the oldest event is displayed first.

Events are ordered by severity. The first event is the most severe. Events are ordered by severity in the following way:

- ▶ Unfixed alerts (sorted by error code). The lowest error code has the highest severity.
- ▶ Unfixed messages.
- ▶ Monitoring events (sorted by error code). The lowest error code has the highest severity.
- ▶ Expired events.
- ▶ Fixed alerts and messages.

The less severe events are often fixed with the resolution of the most severe events.

12.8 Support Assistance

Support assistance enables support personnel to access the system to complete troubleshooting and maintenance tasks. You can configure either local support assistance, where support personnel visit your site to fix problems with the system, or remote support assistance. Both local and remote support assistance use secure connections to protect data exchange between the support center and system. More access controls can be added by the system administrator. Assistance can be provided at your location or through a remote connection to your system.

Local support assistance

Use local support assistance if you have restrictions that require on-site support only. Unlike other authentication methods, you can audit all actions that support personnel conduct on the system when local support assistance is configured. Support personnel can log on to your system by using a console or over your intranet. These users can be authenticated only by a challenge-response mechanism. Support personnel obtain the challenge-response access either through virtual private network (VPN) or over a telephone call with another support person or the administrator at the support center.

Note that if you want to enable remote support assistance or use the Assist On-Site tool, you must configure local support assistance.

Remote support assistance

With remote support assistance, support personnel can visit on site and they can also access the system remotely through a secure connection from the support center. However, before you enable remote support assistance between the system and support, you first need to configure local support assistance. You must ensure that call home is configured and a valid email server is specified. Call home automatically contacts support when critical errors occur on the system. Call home sends a return email that communicates information back to the system such as a Problem Management Report (PMR) number that tracks the problem until it is resolved.

Note that you cannot enable remote support assistance and use the Assist On-Site tool at the same time.

In addition, a service IP address must be configured before you set up remote support assistance. During system initialization, you can optionally set up a service IP address and remote support assistance. If you did not configure a service IP address, go to **Settings** →

Network → Service IPs to configure a service IP for each node on the system. Optionally, you need to configure a proxy server if you use a firewall to protect your internal network.

When you enable remote support assistance, a shared-token is also generated by the system and sent to the support center. If the system needs support services, support personnel can be authenticated onto the system with a challenge-response mechanism. Use the **chsra** command to enable remote support assistance on the system. After support personnel obtain the response code, it is entered to gain access to the system. Service personnel have three attempts to enter the correct response code. After three failed attempts, the system generates a new random challenge and support personnel must obtain a new response code.

Support roles

When you enable local support assistance, support personnel are assigned either the Monitor role or the Restricted Administrator role. The Monitor role can view, collect, and monitor logs and errors to determine the solution to problems on the system. The Restricted Administrator role gives support personnel access to administrator tasks to help solve problems on the system. However, this role restricts these users from deleting volumes or pools, unmapping hosts, or creating, deleting, or changing users. Roles limit access of the assigned user to specific tasks on the system. Users with the service role can set the time and date on the system, delete dump files, add and delete nodes, apply service, and shut down the system. They can also view objects and system configuration but cannot configure, modify, or manage the system or its resources. They also cannot read user data.

12.8.1 Configuring support assistance

You find the Support Assistance Screen under **Settings → Support → Support Assistance** as shown in Figure 12-86.

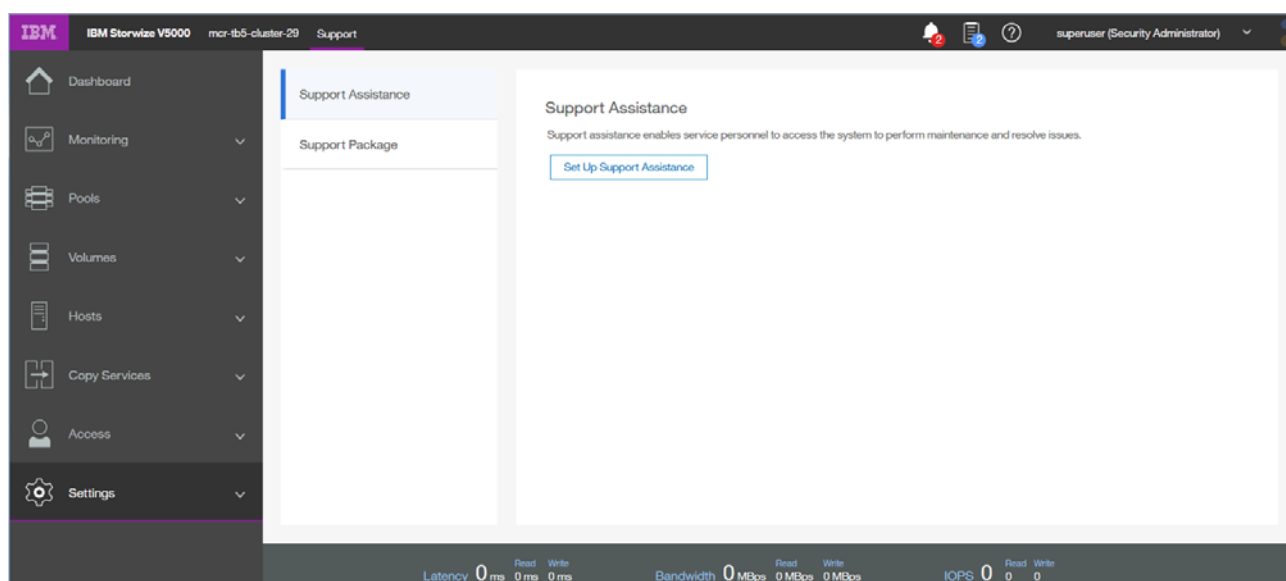


Figure 12-86 Support Assistance

12.8.2 Set up Support Assistant

To Support assistance enables support personnel to access the system to complete troubleshooting and maintenance tasks. You can configure either local support assistance, where support personnel visit your site to fix problems with the system, or remote support

assistance. Both local and remote support assistance use secure connections to protect data exchange between the support center and system. More access controls can be added by the system administrator. The system supports both local and remote support assistance. Use local support assistance if you have restrictions that require on-site support only. Unlike other authentication methods, you can audit all actions that support personnel conduct on the system when local support assistance is configured. With remote support assistance, support personnel can visit on site and they can also access the system remotely through a secure connection from the support center. However, before you enable remote support assistance between the system and support, you first need to configure local support assistance. Support personnel rely on the support package, such as snaps, dumps, and various trace files, to troubleshoot issues on the system. The management GUI and the command-line interface support sending this data to the support center securely. Additionally, support personnel can download new builds, patches, and fixes automatically to the system with your permission.

To configure support assistance, complete the following:

In the management GUI, select **Settings** → **Support** → **Support Assistance** → **Set Up Support Assistance**. See Figure 12-87.

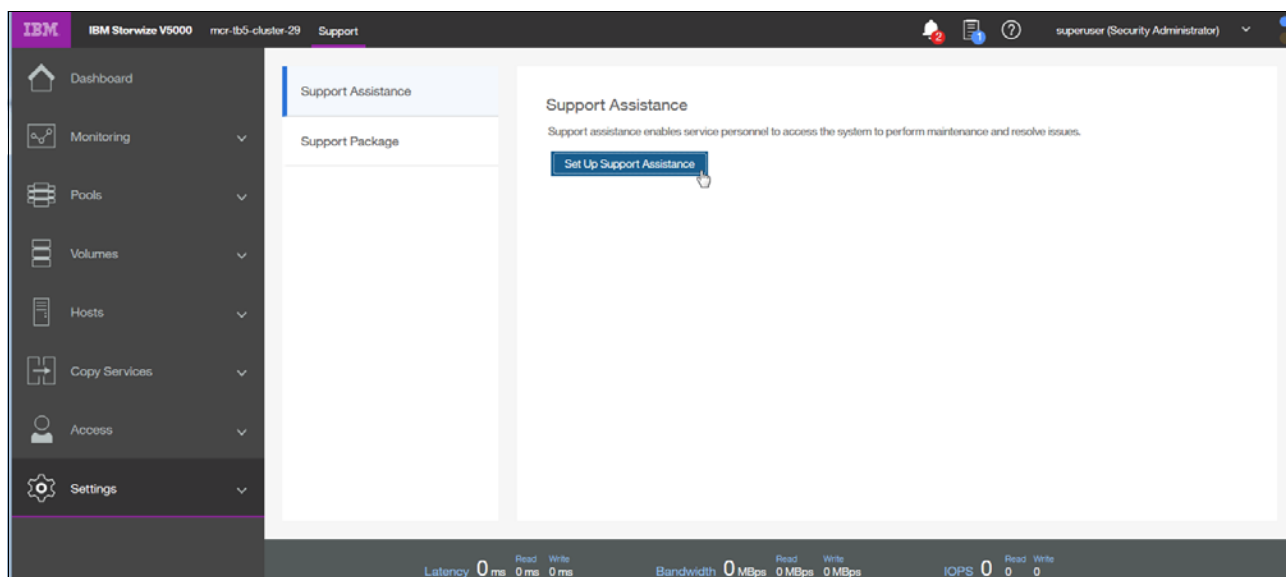


Figure 12-87 Support Assistance

If you selected to configure both local and remote support assistance, verify the pre-configured support centers. Optionally, enter the name, IP address, and port for the proxy server on the Remote Support Centers page. A proxy server is used in systems where a firewall is used to protect your internal network or if you want to route traffic from multiple storage systems to the same place

Enable local support

You have to select: **I want support personnel to work on-site only**

Figure 12-88 on page 682 shows how to enable local support.

Select this option to configure local support assistance. Use this option if your system has certain restrictions that require on-site maintenance. If you select this option, click Finish to set up local support assistance.

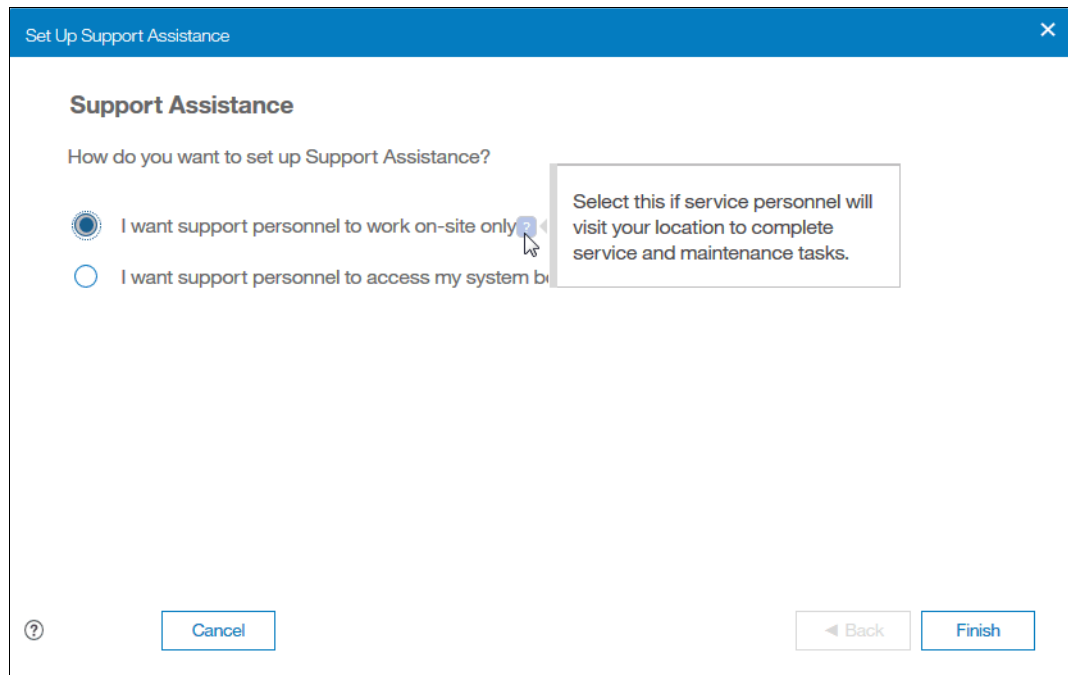


Figure 12-88 Enable local support

The screen in Figure 12-89 appears.

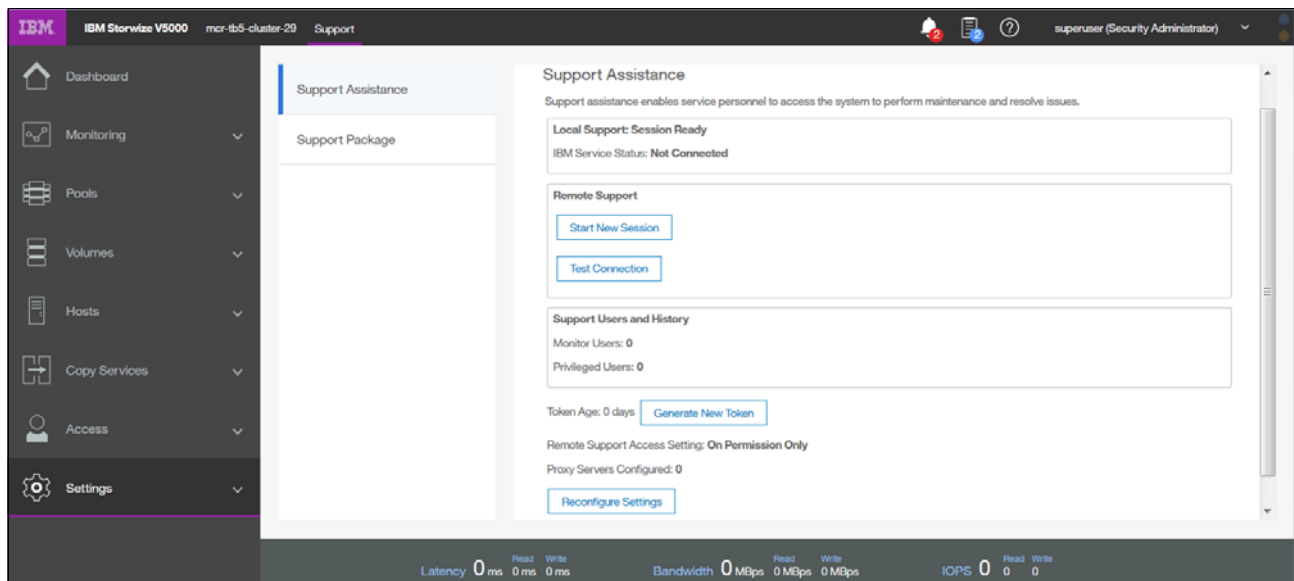


Figure 12-89 Local Support Definitions

Under **Support assistance** → **Start new Session** (marked) you can select the time which the Remote support Session can be idle before the system disconnects the line, see Figure 12-90 on page 683.

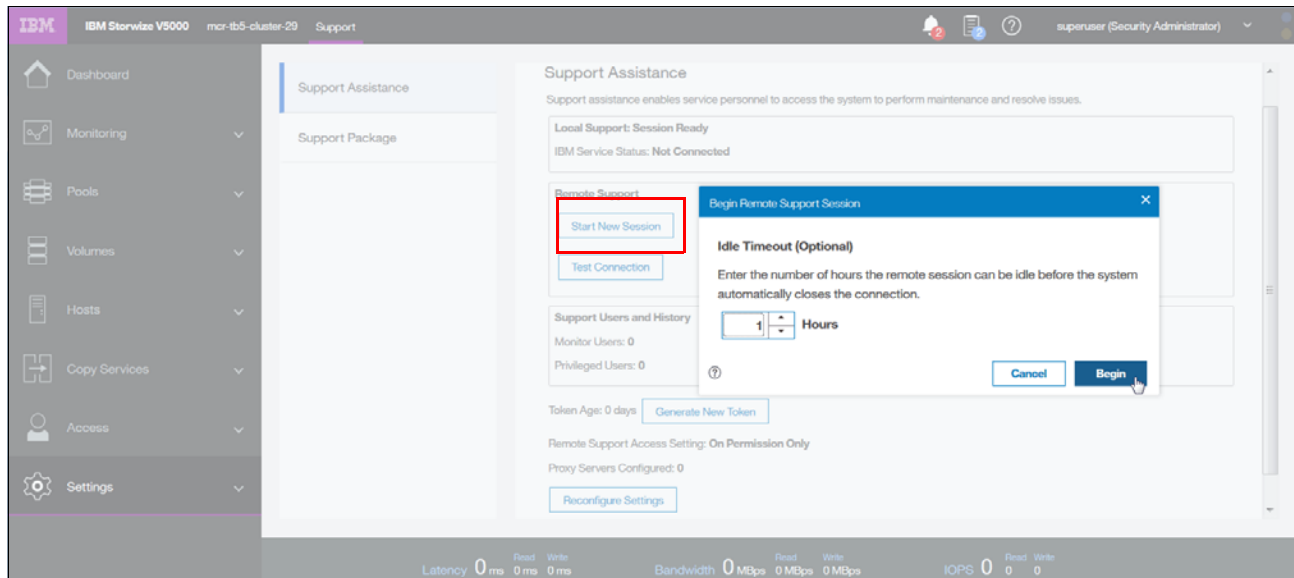


Figure 12-90 Set Idle time before the line will be disconnected

Press **Begin** to begin a new session.

Test Connection lets you test the connectivity as shown in Figure 12-91.

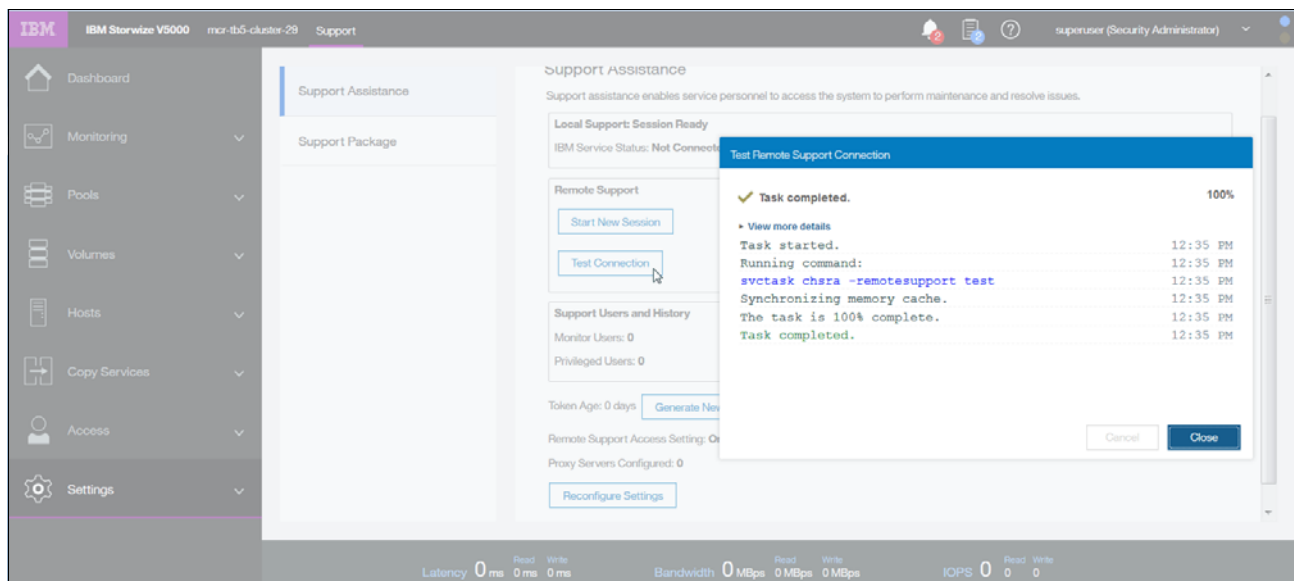


Figure 12-91 Test connection

Figure 12-92 shows the pop up testing the line to the Service Center.




Figure 12-92 Test Service Center Connection

An overview of remote users is shown in Figure 12-93.



Figure 12-93 Support Users

A new Token can be generated by pressing the button **Generate New Token** as shown in Figure 12-94.

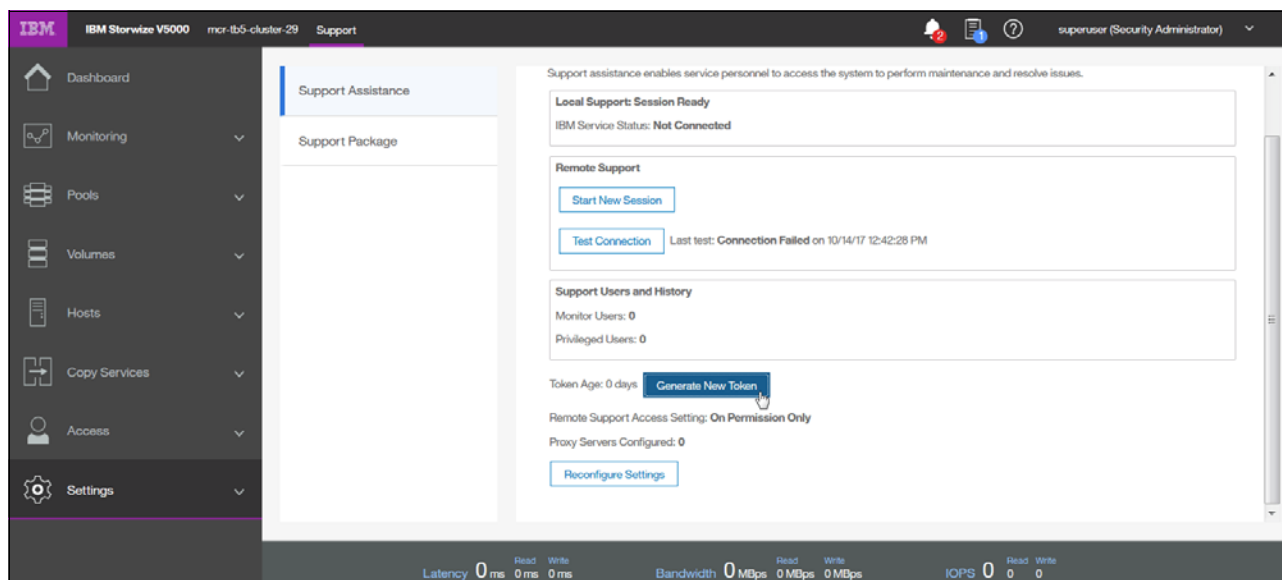


Figure 12-94 Generate New Token

When you enable remote support assistance, the system generates a support assistance token. This shared security token is sent to the support center and is used for authentication during support assistance sessions. Updating a token is essentially overwriting the existing token, then sending it securely to the support assistance administration server in an email message. You specify the email addresses of the support assistance administration servers when you configure support assistance. If the email is not received in time for a support incident or cannot be sent for some reason, a service engineer can manually add the token to the administration server. Before you can update a token, you must enable the support assistance feature. You can update the token periodically as a security practice, similar to how you update passwords.

To update a shared support assistance token, enter the following command:

```
svctask chsra -updatetoken
```

If settings change over time you can reconfigure your settings using the button **Reconfigure Settings** as shown in Figure 12-95 on page 685.

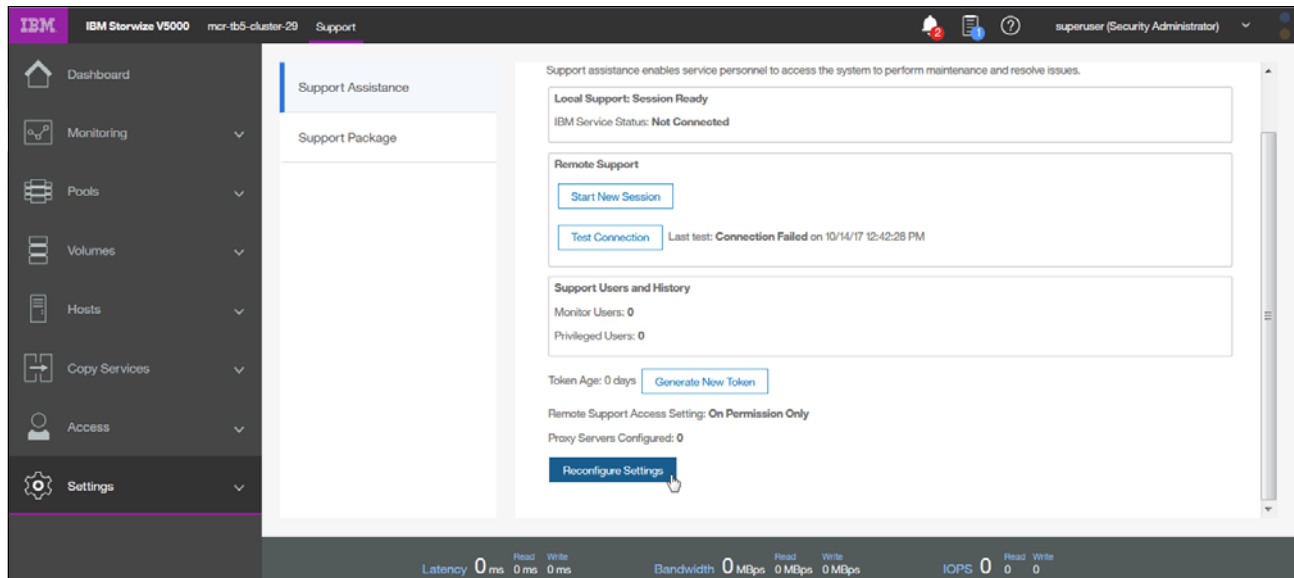


Figure 12-95 Reconfigure Settings

Enable remote support

If you are configuring remote support assistance ensure that the following prerequisites are met:

- ▶ Ensure that call home is configured with a valid email server
- ▶ Ensure that a valid service IP address is configured on each node on the system
- ▶ If your system is behind a firewall or if you want to route traffic from multiple storage systems to the same place, you must configure a Remote Support Proxy server. Before you configure remote support assistance, the proxy server must be installed and configured separately. During the set-up for support assistance, specify the IP address and the port number for the proxy server on the Remote Support Centers page
- ▶ If you do not have firewall restrictions and the storage nodes are directly connected to the Internet, request your network administrator to allow connections to 129.33.206.139 and 204.146.30.139 on Port 22
- ▶ Both uploading support packages and downloading software require direct connections to the Internet. A DNS server must be defined on your system for both of these functions to work
- ▶ To ensure that support packages are uploaded correctly, configure the firewall to allow connections to the following IP addresses on port 443: 129.42.56.189, 129.42.54.189, and 129.42.60.189
- ▶ To ensure that software is downloaded correctly, configure the firewall to allow connections to the following IP addresses on port 22: 170.225.15.105, 170.225.15.104, 170.225.15.107, 129.35.224.105, 129.35.224.104, and 129.35.224.107

Using the management GUI

To configure remote support assistance, complete the following.

Select → **I want support personnel to access my system both on-site and remotely**, see Figure 12-96 on page 686.

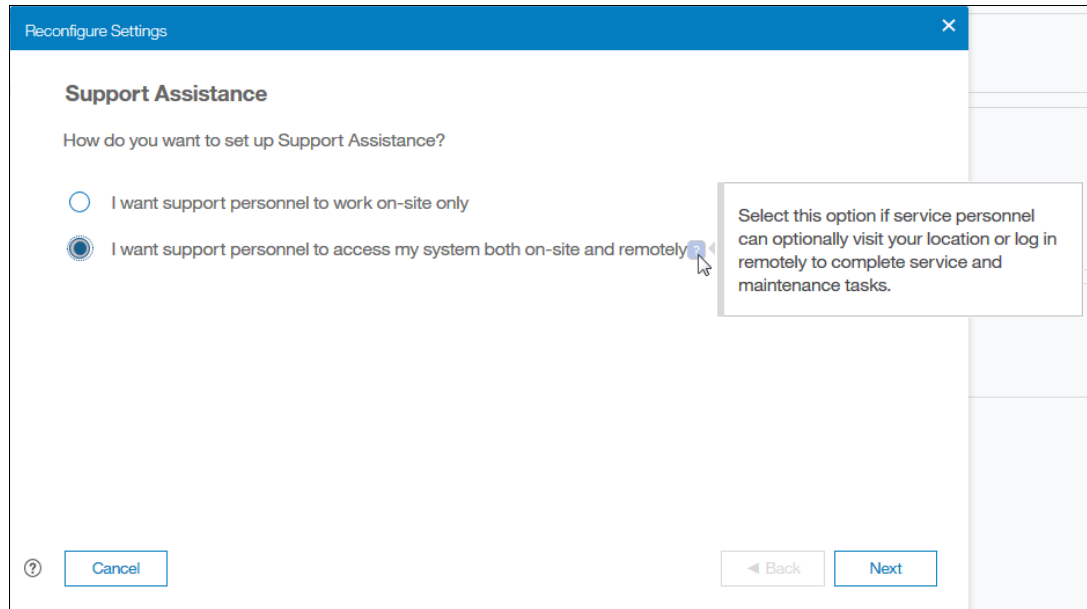


Figure 12-96 Enable remote support

Select this option to configure remote support assistance. Use this option to allow support personnel to access your system through a secure connection from the support center. Secure remote assistance requires a valid service IP address, call home, and an optional proxy server if a firewall is used to protect your internal network. If you select this option, click Next to specify IP addresses for the support center and optional proxy server. See Figure 12-97 on page 687.

Reconfigure Settings

Support Centers

Support centers respond to manual and automatic service requests from the system. The following support centers are configured on the system:

Name	IP Address	Port
default_support_center0	129.33.206.139	22
default_support_center1	204.146.30.139	22

Remote Support Proxy (Optional)

Required for network configurations using a firewall, or for systems without direct connection to the network.

Name

Customer Proxy 2

IP

173.63.152.234

Port

8080

To remove an existing proxy server, select from the list below and click Next.

Name	IP Address	Port
<input type="checkbox"/> Customer Proxy	173.63.152.233	8080

Cancel

Back

Next

Figure 12-97 Support Centers

Click **Next**. On the Remote Support Access Settings page, select one of these options to control when support personnel can access your system to conduct maintenance and fix problems:

At Any Time — Support personnel can access the system at any time. For this option, remote support session does not need to be started manually and sessions remain open continuously.

On Permission Only — The system administrator must grant permission to support personnel before they can access the system. See Figure 12-98 on page 688.

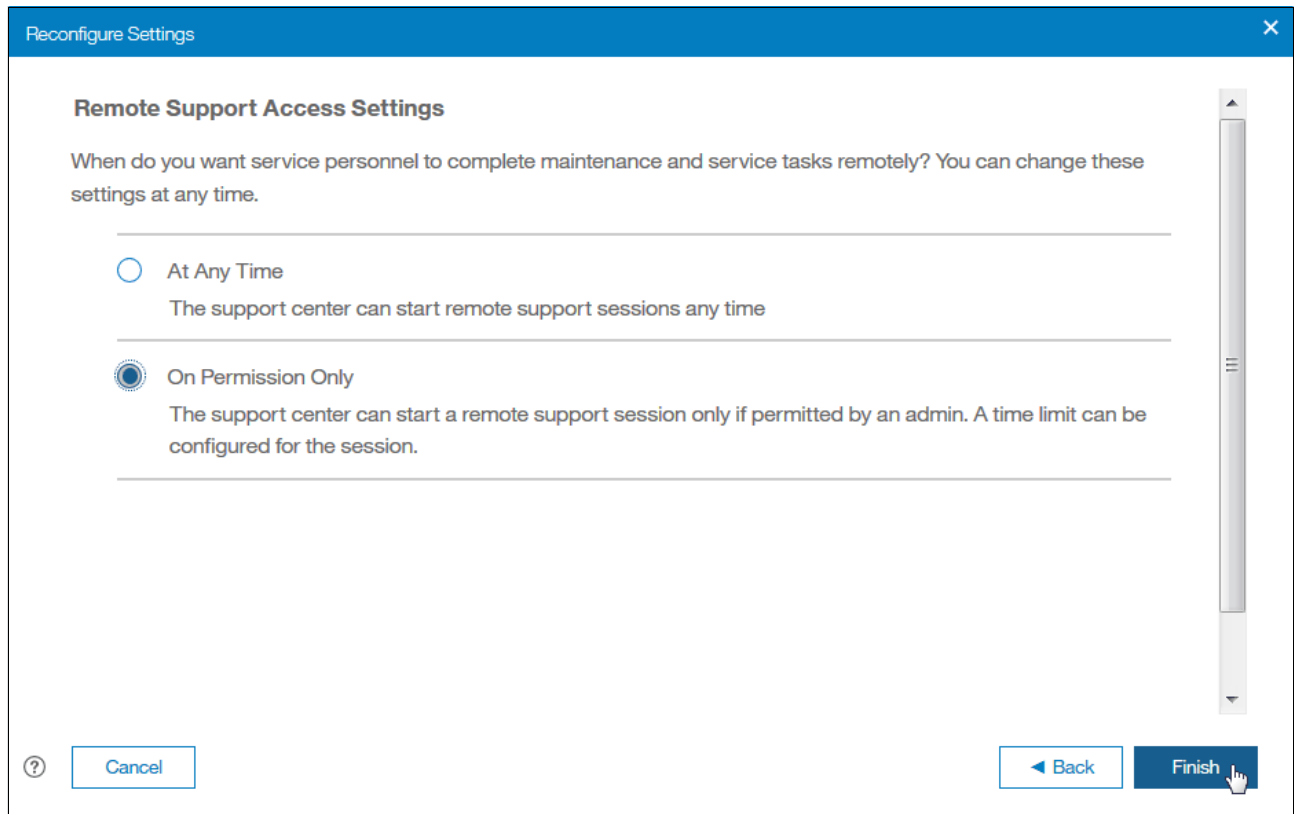


Figure 12-98 Remote Support Access Settings

Click **Finish**. After you configure remote support assistance with permission only, you can start sessions between the support center and the system. On the Support Assistance page, select Start New Session and specify the number of hours the session can be idle before the support user is logged off from the system. See Figure 12-99.

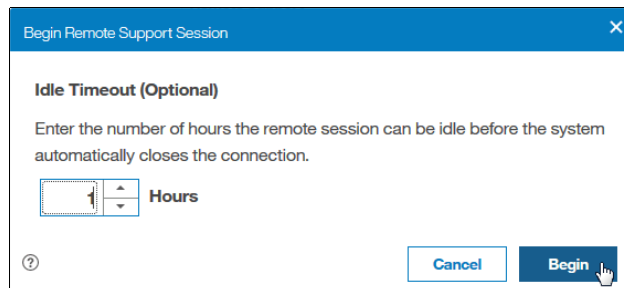


Figure 12-99 Idle Timeout Setting

If you plan to use the command-line interface to configure local support assistance, enter the following command:

chsra -enable

To configure remote support assistance, enter the following command:

chsra -remotesupport enable

12.8.3 Disable Support Assistance

You can disable support assistance by using the command-line interface (CLI). When you disable support assistance, the support assistance token is deleted. All active secure remote access user sessions are closed immediately and a secure email message is sent to the administration server to indicate that secure remote access is disabled on the system.

To disable support assistance completely, enter the following command:

```
svctask chsra -disable
```

To disable remote support assistance only, enter the following command:

```
svctask chsra -remotesupport disable
```

12.9 Collecting support information

If you have an issue with a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and call the IBM Support Center, you might be asked to provide support data as described in the next section.

12.9.1 Collecting support information by using the GUI

The following information describes how to collect supporting data for the IBM Support Center.

To reach the Support Package screen go to **Settings** → **Support** → **Support Package** following screens opens. See Figure 12-100.

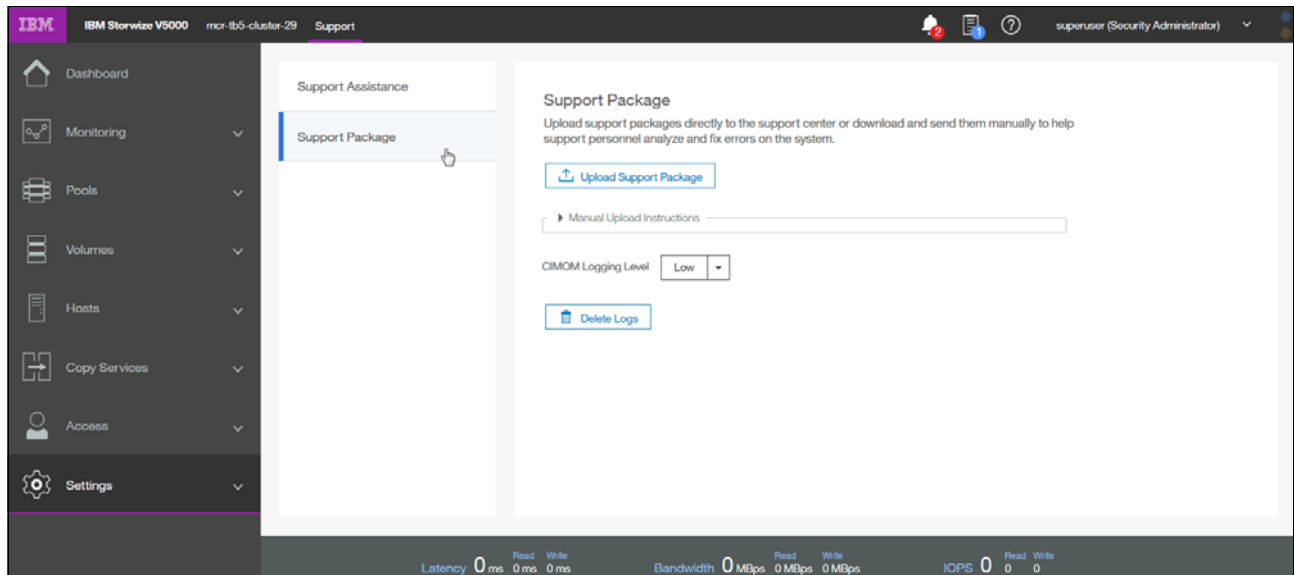


Figure 12-100 Support Package

12.9.2 Automatic upload of Support Packages

You can use the management GUI or the command-line interface to upload support packages to the support center. If support assistance is configured on your systems, you can either

automatically or manually upload new support packages to the support center to help analyze and resolve errors on the system. You can select individual logs to either download to review or send directly to the support center for analysis.

Before automatically uploading a support package, ensure that the following prerequisites are configured on the system:

- ▶ Ensure that all of the nodes on the system have internet access.
- ▶ Ensure that a valid service IP address is configured on each node on the system.
- ▶ Configure at least one valid DNS server for domain name resolution. To configure a DNS server on the system, select **Settings** → **System** → **DNS** and specify valid IP addresses and names for one or more DNS servers. You can also use the `mkdnsserver` command to configure DNS servers.
- ▶ Configure the firewall to allow connections to the following IP addresses on port 443: 129.42.56.189, 129.42.54.189, and 129.42.60.189. To test connections to the support center, select **Settings** → **Support** → **Support Assistance**. On the Support Assistance page, select Test Connection to verify connectivity between the system and the support center.

The management GUI supports uploading new or existing support packages to support automatically. See Figure 12-101.

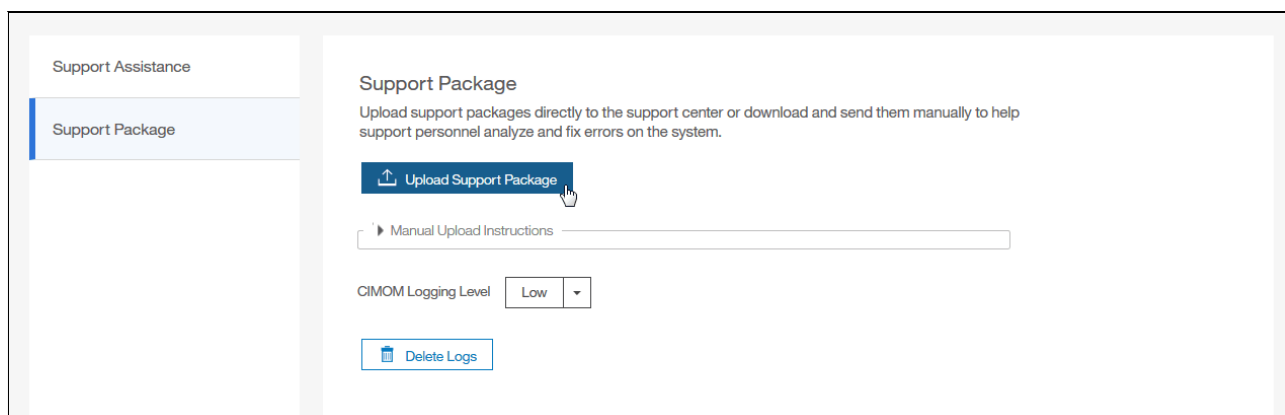


Figure 12-101 Upload Support Package

When you press → **Upload Support Package** the selection screen opens as shown in Figure 12-102 on page 691.

Figure 12-102 Upload Support package

On the Upload Support Package page, enter the Problem Management Report (PMR) number that is associated with the support package that you are uploading. If you do not have a PMR number, click **Don't have a PMR?** to open the Service Request (SR) tool to generate a PMR. You need a IBM Partner ID to register.

Note: If you are not sure if a PMR exists or do not want to create a new PMR, the package can still be sent to the support center. The machine serial number and type are used to route the package to the support center. However, specifying a PMR number can decrease response time for support personnel. You can call the Lenovo Support Line or use the Lenovo Support Portal to open a call. Go to the following address:

<https://datacentersupport.lenovo.com/us/en/>

Specify the type of package that you want to generate and upload to the support center by selecting:

- Standard logs
This support package contains the most recent logs that were collected from the system. These logs are most commonly used by the IBM Support Center to diagnose and solve problems.
- Standard logs plus one existing statesave
This support package contains the standard logs from the system and the most recent statesave from any of the nodes in the system. Statesaves are also known as *memory dumps* or *live memory dumps*.
- Standard logs plus the most recent statesave from each node

This option is used most often by the support team for problem analysis. They contain the standard logs from the system and the most recent statesave from each node in the system.

► Standard logs plus new statesave

This option might be requested by the IBM Support Center team for problem determination. It generates a new statesave (livedump) for all of the nodes and packages them with the most recent logs.

The support center will let you know which package they need.

Click Upload. After the new support package is generated, a summary panel displays the progress of the upload. If the upload is unsuccessful or encounters errors, verify the connection between the system and the support center and retry the upload.

If you decide that you want to upload the support package later you can use the function shown in Figure 12-103.




Figure 12-103 Upload Existing Package

If you press **Upload Existing Package** a screen opens as shown in Figure 12-104.

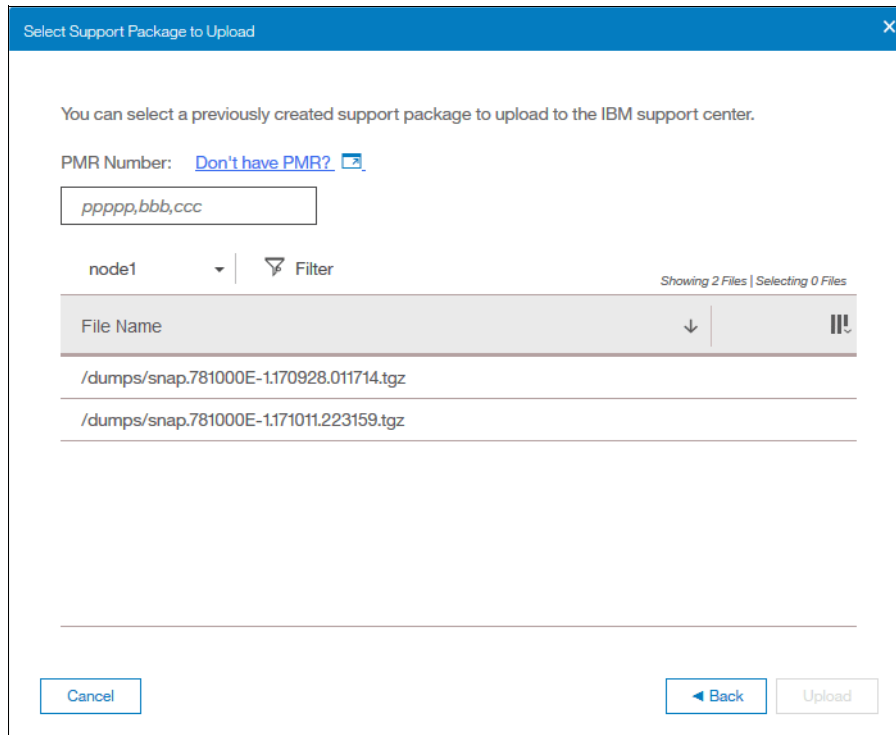


Figure 12-104 Select Support Package to Upload

Using the command-line interface

To upload a support package or other file with the command-line interface, complete these steps:

Enter the following command:

```
satask supportupload -pmr pmr_number -filename fullpath/filename
```

where the `pmr_number` is the number of an existing PMR and `fullpath/filename` is the full path and the name of the file that you are uploading. The `-pmr` and `-filename` parameters are not required. If you do not specify a PMR number, the file is uploaded by using the machine serial and type to route the file to the support center. If you do not specify a file name, the latest support package is uploaded.

To verify the progress of the upload to the support center, enter the following command:

```
!scmdstatus
```

In the results of this command, verify that the `supportupload_status` is `Complete`, which indicates that the upload is successfully completed. Other possible values for this parameter include `Active`, `Wait`, `Abort`, and `Failed`. If the upload is `Active`, you can use the `supportupload_progress_percent` parameter to view the progress for the upload.

If you want to generate a new support package, complete these steps:

Enter the following command in the command-line interface:

```
satask snap -upload -pmr pmr_number
```

where the `pmr_number` is the number of an existing PMR. The command generates a new support package and uploads it to the support center with the identifying PMR number. If you do not have a PMR number that corresponds with support package, then you can use the following command:

```
satask snap -upload
```

The command generates a new support package and uploads it to the support center by using the machine type and serial to route the package.

To verify the progress of the upload to the support center, enter the following command:

```
!scmdstatus
```

In the results of this command, verify that the `supportupload_status` is `Complete`, which indicates that the upload is successfully completed. Other possible values for this parameter include `Active`, `Wait`, `Abort`, and `Failed`. If the upload is `Active`, you can use the `supportupload_progress_percent` parameter to view the progress for the upload.

12.9.3 Manual upload of Support Packages

You can use the management GUI or the command-line interface to upload manually support packages to the support center. If support assistance is configured on your systems, you can manually upload new support packages to the support center to help analyze and resolve errors on the system. You can select individual logs to either download to review or send directly to the support center for analysis.

Using the management GUI

The management GUI supports manually uploading support packages. Manually uploading support packages require that you download either a new support package or an existing support package to your system and then upload the file to support directly.

To manually upload a new support package to the support center, complete these steps:

In the management GUI, select **Settings** → **Support** → **Support Package**.

On the Support Package page, expand **Manual Upload Instructions**. Figure 12-105.

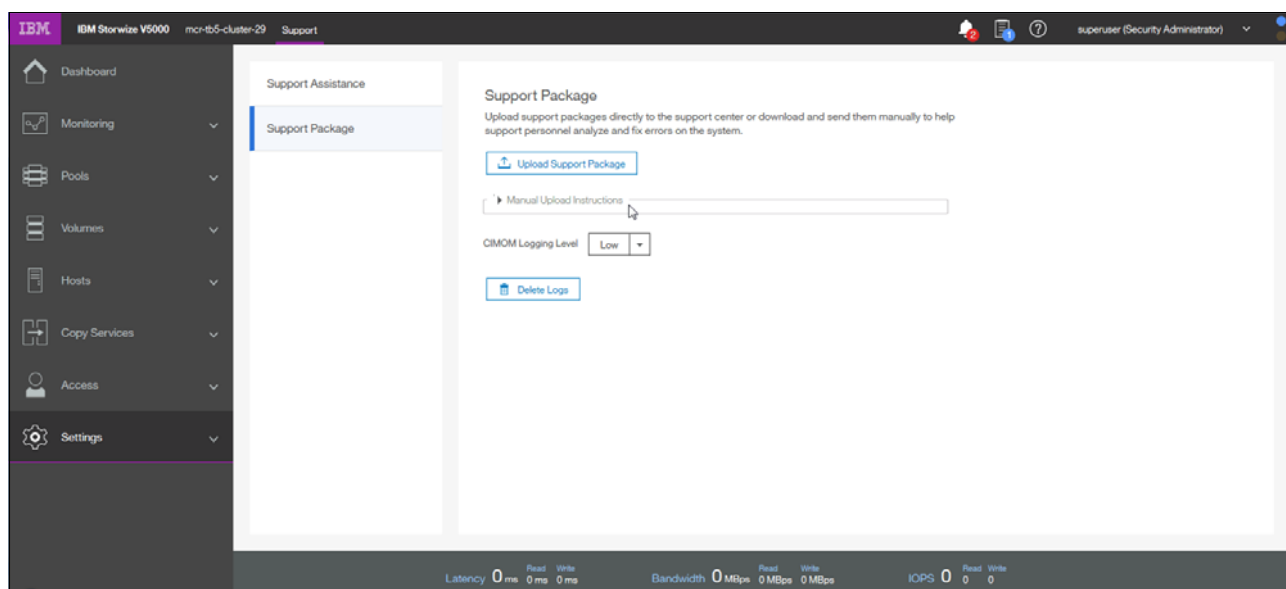


Figure 12-105 Manual Upload Instructions

In the Manual Upload Instructions section, click Download Support Package. See Figure 12-106.

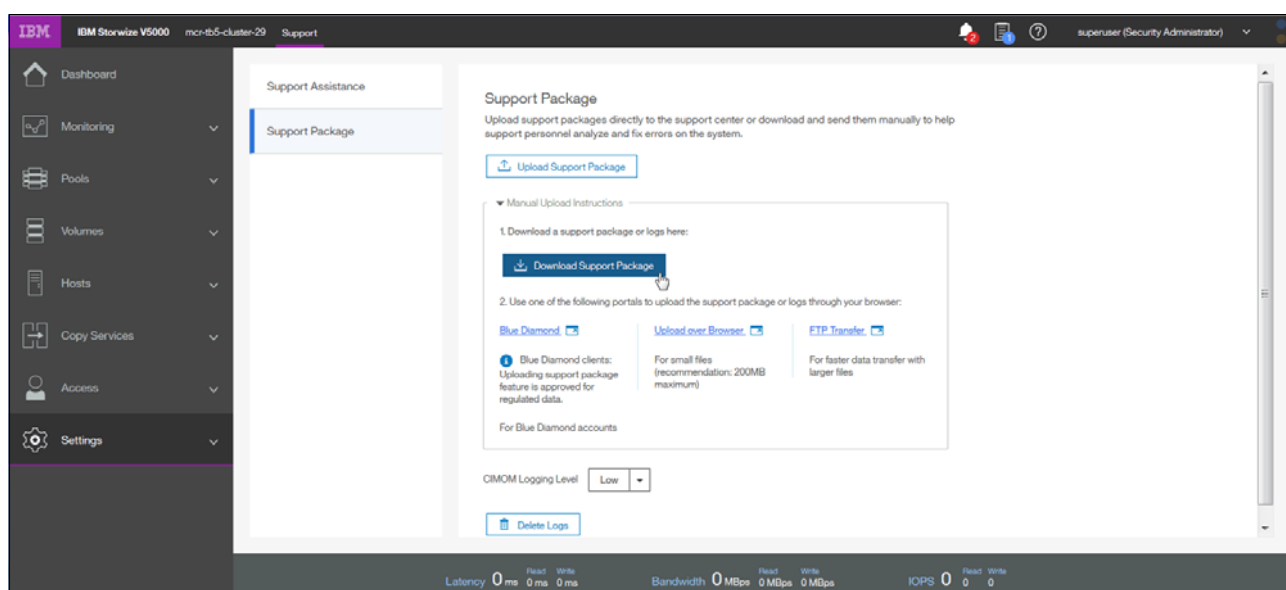


Figure 12-106 Download Support Package

On the Download New Support Package or Log File panel, select one of these types of support packages to download which are shown in Figure 12-107.

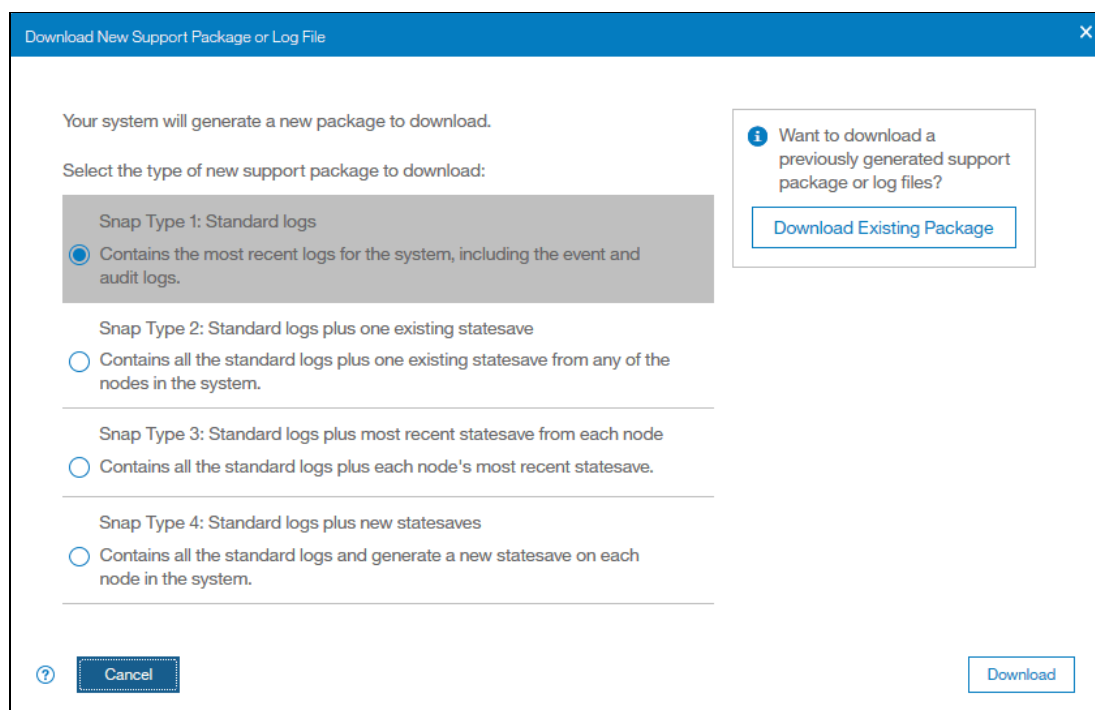


Figure 12-107 Download Support Package or Log file

The type to select depends on the event that is being investigated. For example, if you notice that a node is restarted, capture the snap file with the latest existing statesave. If needed, the IBM Support Center can notify you of the package that is required.

The following components are included in each type of support package:

- ▶ Standard logs
This support package contains the most recent logs that were collected from the system. These logs are most commonly used by the IBM Support Center to diagnose and solve problems.
- ▶ Standard logs plus one existing statesave
This support package contains the standard logs from the system and the most recent statesave from any of the nodes in the system. Statesaves are also known as *memory dumps* or *live memory dumps*.
- ▶ Standard logs plus the most recent statesave from each node
This option is used most often by the support team for problem analysis. They contain the standard logs from the system and the most recent statesave from each node in the system.
- ▶ Standard logs plus new statesave
This option might be requested by the IBM Support Center team for problem determination. It generates a new statesave (livedump) for all of the nodes and packages them with the most recent logs.

Click **Download** to download the support package to your local computer.

After the download completes to your local computer, you can upload the package to the support center with one of the following methods shown in Figure 12-108.




Figure 12-108 Download portals

This explains the different portals and in each case a Webpage opens in your browser.

► **Blue Diamond**

Select the link to log in to the BlueDiamond portal. BlueDiamond provides enhanced security and support for healthcare clients. You must be a registered BlueDiamond client to use this option. After you accept the terms of service for the upload, log into the BlueDiamond portal with your user name and password.

► **Upload over Browser**

Use this option for small files under 200 MB. Select the link to upload the support package to the support web page through the web browser. On the support web page, complete the following steps:

Enter a valid PMR number that is associated with this support package. In the Upload is for field, select Other. Enter a valid email address for the contact for this package.

► **FTP Transfer**

Use this option for larger files. Select the link to send the package to support with file transfer protocol (FTP). You can send packages to support with standard FTP (non-secure), secure FTP, or with SFTP, which is FTP over secure shell protocol (SSH). On the support port for FTP transfers, select the type of FTP you want to use and follow the instructions for that method.

Using the command-line interface

To upload a support package or other file with the command-line interface, complete these steps:

Enter the following command:

```
satask supportupload -pmr pmr_number -filename fullpath/filename
```

where the `pmr_number` is the number of an existing PMR and `fullpath/filename` is the full path and the name of the file that you are uploading. The `-pmr` and `-filename` parameters are not required. If you do not specify a PMR number, the file is uploaded by using the machine serial and type to route the file to the support center. If you do not specify a file name, the latest support package is uploaded.

To verify the progress of the upload to the support center, enter the following command:

```
lscmdstatus
```

In the results of this command, verify that the `supportupload_status` is Complete, which indicates that the upload is successfully completed. Other possible values for this parameter include Active, Wait, Abort, and Failed. If the upload is Active, you can use the `supportupload_progress_percent` parameter to view the progress for the upload.

If you want to generate a new support package, complete these steps:

Enter the following command in the command-line interface:

```
satask snap -upload -pmr pmr_number
```

where the `pmr_number` is the number of an existing PMR. The command generates a new support package and uploads it to the support center with the identifying PMR number. If you do not have a PMR number that corresponds with support package, then you can use the following command:

```
satask snap -upload
```

The command generates a new support package and uploads it to the support center by using the machine type and serial to route the package.

To verify the progress of the upload to the support center, enter the following command:

```
lscmdstatus
```

In the results of this command, verify that the `supportupload_status` is `Complete`, which indicates that the upload is successfully completed. Other possible values for this parameter include `Active`, `Wait`, `Abort`, and `Failed`. If the upload is `Active`, you can use the `supportupload_progress_percent` parameter to view the progress for the upload.

12.9.4 Collecting support information by using the SAT

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 management GUI collects information from all of the components in the system. The Service Assistant Tool (SAT) collects information from all node canisters. The *snap file* is the information that is collected and packaged in a single file.

If the package is collected by using the Service Assistant Tool, ensure that the node from which the logs are collected is the current node, as shown in Figure 12-109.

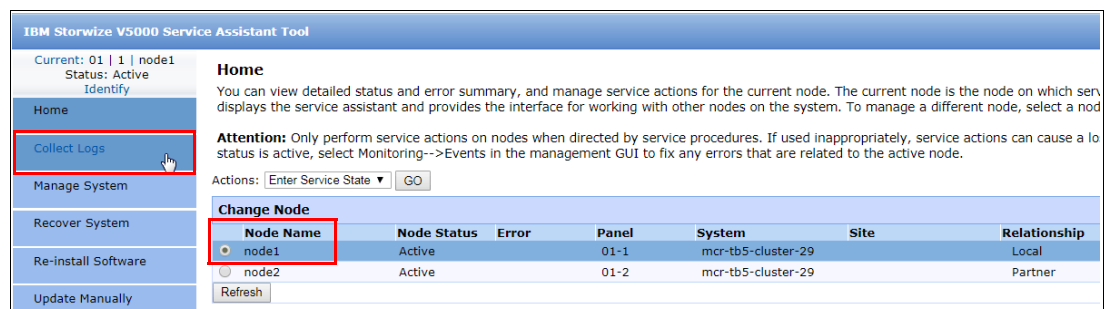


Figure 12-109 Accessing the Collect Logs panel in the Service Assistance Tool

Support information can be downloaded with or without the latest statesave, as shown in Figure 12-110 on page 698.

Figure 12-110 Collect Logs panel in the Service Assistance Tool

Accessing the SAT by using the technician port

If your system or one of your node canisters is inaccessible through the administrative network, you can connect a personal computer directly to the technician port on the node canister to access the Service Assistant Tool.

Note: This procedure starts the initialization tool if the node canister that is being serviced is in the candidate state, if no system details are configured, and if the partner node is not in the active state.

Complete the following steps:

1. Configure Dynamic Host Configuration Protocol (DHCP) on the Ethernet port of the personal computer to connect to the node canister.
Alternatively, if the personal computer does not support DHCP, configure the static IPv4 address 192.168.0.2 on the port.
2. On the Lenovo Storage V3700 V2 system or Lenovo Storage V3700 V2 XP system, reenabling the technician port by completing the following steps:
 - a. Create a text file with the **satask chserviceip -techport enable -force** command.
 - b. Save the file as **satask.txt** in the root directory of the Universal Serial Bus (USB) stick.
 - c. Insert the USB stick in the USB port of the node that you want to service.
 - d. Wait until no write activity is recognized and remove the USB stick.

Note: The Lenovo Storage V5030 system has a dedicated technician port that is always enabled so this step is unnecessary.

3. Connect an Ethernet cable between the port on the personal computer and the technician port. The technician port is labeled with a **T** on the rear of the node canister.
4. Open a supported web browser on the personal computer and browse to the **http://192.168.0.1** URL.

Note: If the cluster is active and you connect to the configuration node, this URL opens the management GUI. If you want to access the SAT in this case, browse to

<http://192.168.0.1/service>

5. Complete the correct procedure to service the canister.
6. Log out of the Service Assistant Tool and disconnect the Ethernet cable from the technician port.
7. On the Lenovo Storage V3700 V2 system or Lenovo Storage V3700 V2 XP system, disable the technician port by running the command that is shown in Example 12-3.

Example 12-3 Disabling the technician port

```
>satask chserviceip -techport disable
```

SAS port 2 can then be used again to provide extra Ethernet connectivity for system management, iSCSI, and IP replication.

12.10 Powering off the system and shutting down the infrastructure

The following sections describe the process to power off the system and to shut down and start an entire infrastructure that contains a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

12.10.1 Powering off

Important: Never power off your Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems by powering off the power supply units (PSUs), removing both PSUs, or removing both power cables from a running system. It can lead to inconsistency or loss of the data that is staged in the cache.

You can power off a node canister or the entire system. When you power off only one node canister for each I/O group, all of the running tasks remain active while the remaining node takes over.

Powering off the system is typically planned in site maintenance (power outage, building construction, and so on) because all components of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 are redundant and replaceable while the system is running.

Important: If you are powering off the entire system, you lose access to all volumes that are provided by this system. Powering off the system also powers off all Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 nodes. All data is flushed to disk before the power is removed.

Before you power off the system, stop all hosts with volumes that are allocated to this system. This step can be skipped for hosts with volumes that are provisioned with mirroring (host-based mirror) from different storage systems. However, skipping this step means that errors that relate to lost storage paths and disks can be logged on the host error log.

Note: If a canister or the system is powered off, a local visit can be required to either reseal the canister or power cycle the enclosures.

Powering off a node canister

To power off a canister by using the GUI, complete the following steps:

1. Browse to **Monitoring** → **System** and rotate the enclosure to the rear view, as shown in Figure 12-111.




Figure 12-111 Rotating the system image

2. Right-click the required canister and select **Power Off Canister**, as shown in Figure 12-112.

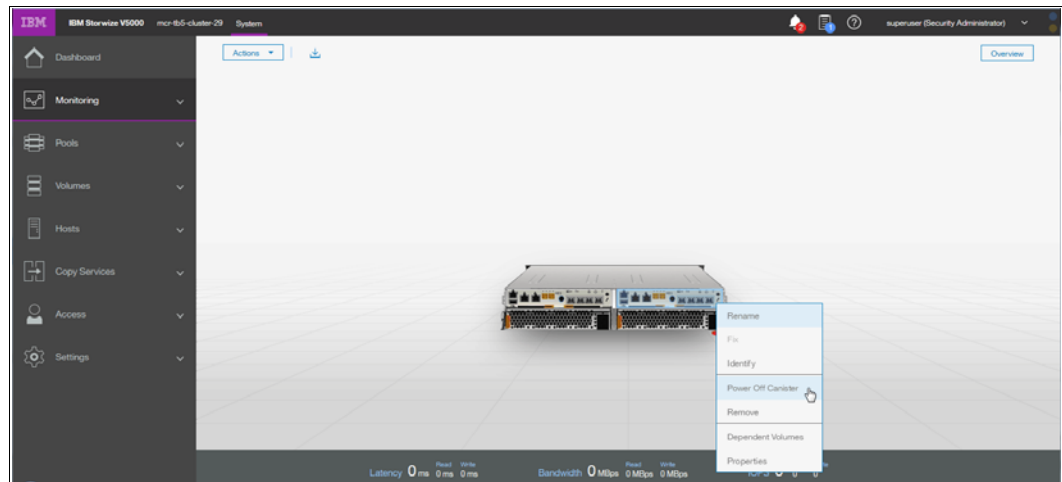


Figure 12-112 Powering off the canister

3. Confirm that you want to power off the canister by entering the confirmation code and clicking **OK**, as shown in Figure 12-113 on page 701.

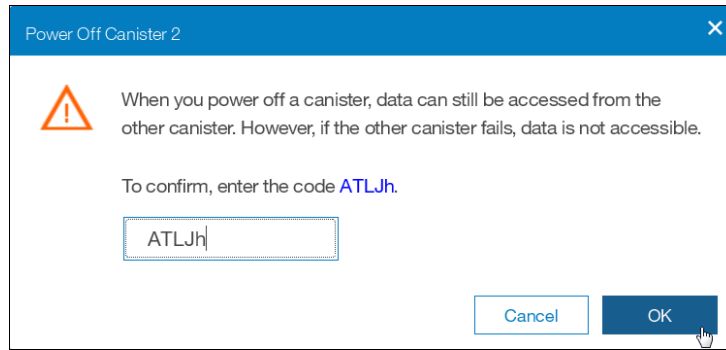


Figure 12-113 Canister power off confirmation window

4. After the node canister is powered off, you can confirm that it is offline in the System panel, as shown in Figure 12-114.




Figure 12-114 Checking the canister state

To power off a node canister by using the CLI, use the command that is shown in Example 12-4.

Example 12-4 Powering off a canister by using the CLI

```
>svctask stopsystem -node 2
```

```
Are you sure that you want to continue with the shut down? (y/yes to confirm)
```

Powering off the system

To power off the entire system by using the GUI, complete the following steps:

1. Browse to **Monitoring** → **System**, click **Actions** → **Power Off System**, as shown in Figure 12-115 on page 702.

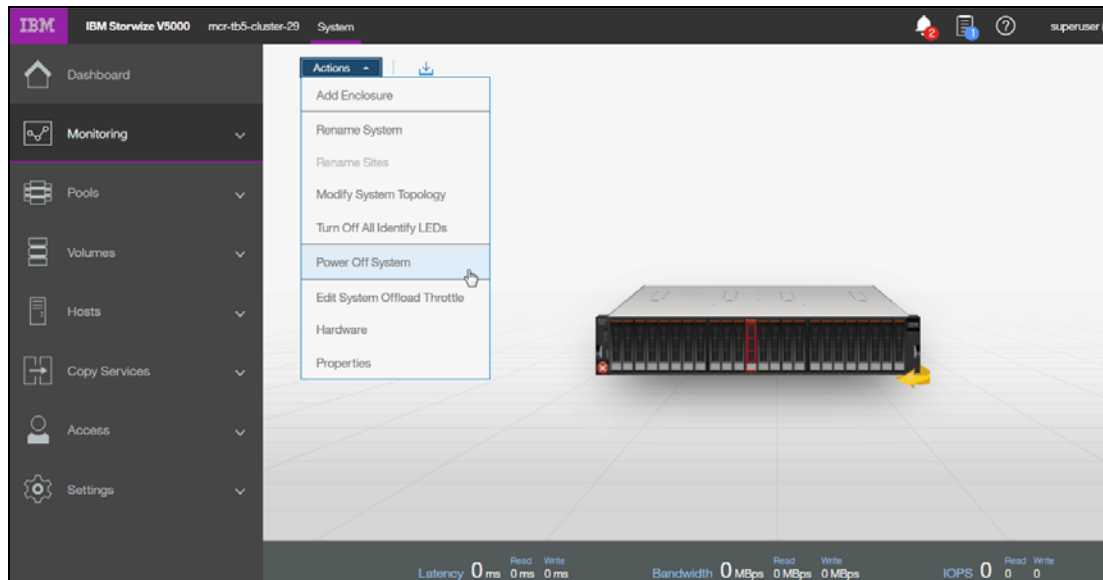


Figure 12-115 Powering off the system

2. Confirm that you want to power off the system by entering the confirmation code and clicking **OK**, as shown in Figure 12-116. Ensure that all FlashCopy, Metro Mirror, Global Mirror, data migration operations, and forced deletions are stopped or allowed to complete before you continue.

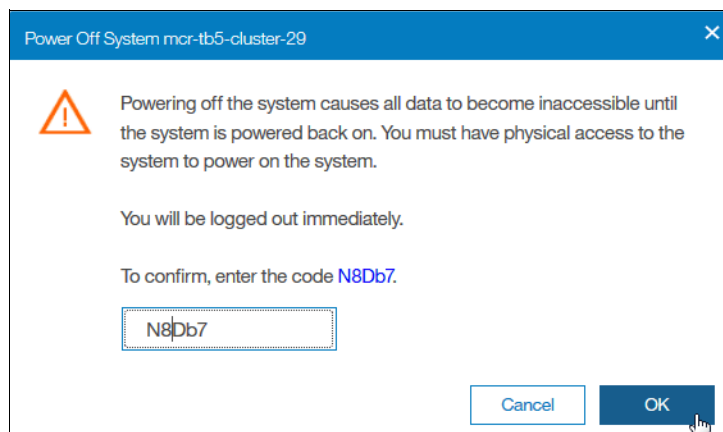


Figure 12-116 Power Off System confirmation window

To power off the system by using the CLI, use the command that is shown in Example 12-5. Ensure that all FlashCopy, Metro Mirror, Global Mirror, data migration operations, and forced deletions are stopped or allowed to complete before you continue.

Example 12-5 Powering off the system by using the CLI

```
>svctask stopsystem
```

```
Are you sure that you want to continue with the shut down? (y/yes to confirm)
```

Wait for the power LED on the node canisters to blink slowly, which indicates that the power off operation completed.

Note: When you power off a Lenovo Storage V3700 V2, V3700 V2 XP, or V5030, it does not automatically restart. You must manually restart the system by removing and reapplying the power / powercords.

12.10.2 Shutting down and starting up the infrastructure

To shut down an entire infrastructure (storage, servers, and applications), complete the following steps:

1. Power off your servers and all applications.
2. Power off your Lenovo Storage V3700 V2, V3700 V2 XP, or V5030 system by using either the GUI or the CLI.
3. Remove the power cords that are connected to both power supplies in the rear of the enclosure on every control and expansion enclosure.
4. Power off your storage area network (SAN) switches.

To start an entire infrastructure, complete the following steps:

1. Power on your SAN switches and wait until the boot completes.
2. Power on any expansion enclosures by connecting the power cord to both power supplies in the rear of the enclosure or by turning on the power circuit.
3. Power on the control enclosures by connecting the power cords to both power supplies in the rear of the enclosure and by turning on the power circuits.

The system starts. The system starts successfully when the status LEDs of all node canisters in the control enclosure are permanently on, which takes no longer than 10 minutes.

Power on your servers and start all applications.

Encryption

Encryption protects against the potential exposure of sensitive user data that is stored on discarded, lost, or stolen storage devices. Lenovo Storage V3700 V2 XP, and Lenovo Storage V5030 support optional encryption of data at-rest. Lenovo Storage V3700 V2 does not support encryption.

Specifically, this chapter provides information about the following topics:

- ▶ 13.1, “Planning for encryption” on page 706
- ▶ 13.2, “Defining encryption of data at-rest” on page 706
- ▶ 13.3, “Activating encryption” on page 711
- ▶ 13.4, “Enabling encryption” on page 719
- ▶ 13.5, “Configuring additional providers” on page 739
- ▶ 13.6, “Migrating between providers” on page 743
- ▶ 13.7, “Recovering from a provider loss” on page 744
- ▶ 13.8, “Using encryption” on page 745
- ▶ 13.9, “Rekeying an encryption-enabled system” on page 753
- ▶ 13.10, “Migrating between key providers” on page 758
- ▶ 13.11, “Disabling encryption” on page 759

13.1 Planning for encryption

Data at-rest encryption is a powerful tool that can help organizations protect confidentiality of sensitive information. However encryption, like any other tool, needs to be used correctly to fulfill its purpose.

There are multiple drivers for an organization to implement data at-rest encryption. These can be internal, such as protection of confidential company data, and ease of storage sanitization, or external, like compliance with legal requirements or contractual obligations.

Therefore, before configuring encryption on the storage, the organization should define its needs and, if it is decided that data at-rest encryption is a required measure, include it in the security policy. Without defining the purpose of the particular implementation of data at-rest encryption, it would be difficult or impossible to choose the best approach to implementing encryption and verifying if the implementation meets the set goals.

Below is a list of items which may be worth considering during the design of a solution including data at-rest encryption:

- ▶ Legal requirements
- ▶ Contractual obligations
- ▶ Organization's security policy
- ▶ Attack vectors
- ▶ Expected resources of an attacker
- ▶ Encryption key management
- ▶ Physical security

There are multiple regulations that mandate data at-rest encryption, from processing of Sensitive Personal Information to guidelines of the Payment Card Industry. If there are any regulatory or contractual obligations that govern the data which will be held on the storage system, they often provide a wide and detailed range of requirements and characteristics that need to be realized by that system. Apart from mandating data at-rest encryption, these documents may contain requirements concerning encryption key management.

Another document which should be consulted when planning data at-rest encryption is the organization's security policy

The final outcome of a data at-rest encryption planning session should be replies to three questions:

1. What are the goals that the organization wants to realize using data at-rest encryption?
2. How will data at-rest encryption be implemented?
3. How can it be demonstrated that the proposed solution realizes the set goals?

13.2 Defining encryption of data at-rest

Encryption is the process of encoding data so that only authorized parties can read it. Secret keys are used to encode the data according to well-known algorithms.

Encryption of data at-rest as implemented in Lenovo Storage V series system is defined by the following characteristics:

- ▶ *Data at-rest* means that the data is encrypted on the end device (drives).
- ▶ The algorithm that is used is the Advanced Encryption Standard (AES) US government standard from 2001.

- Encryption of data at-rest complies with the Federal Information Processing Standard 140 (FIPS-140) standard, but is not certified.
- Ciphertext stealing XTS-AES-256 is used for data encryption.
- AES 256 is used for master access keys.
- The algorithm is public. The only secrets are the keys.
- A symmetric key algorithm is used. The same key is used to encrypt and decrypt data.

The encryption of system data and metadata is not required, so they are not encrypted.

Encryption is enabled at a system level and all of the following prerequisites must be met *before* you can use encryption:

- You must purchase an encryption license before you activate the function.
If you did not purchase a license, contact a Lenovo sales representative or Lenovo Business Partner to purchase an encryption license.
- At least three USB flash drives are required if you plan not to use a key management server. They are available as a feature code from Lenovo (see the note on 722).
- You must activate the license that you purchased.
- Encryption must be enabled.

Note: Only data at-rest is encrypted. Host to storage communication and data sent over links used for Remote Mirroring are not encrypted.

Figure 13-1 shows an encryption example. Encrypted disks and encrypted data paths are marked in blue. Unencrypted disks and data paths are marked in red. In this example the server sends unencrypted data to a SAN Volume Controller 2145-DH8 system, which stores hardware-encrypted data on internal disks. The data is mirrored to a remote IBM Storwize V5000 for Lenovo system using Remote Copy. The data flowing through the Remote Copy link is not encrypted. Because the IBM Storwize V5000 for Lenovo is unable to perform any encryption activities, data on the IBM Storwize V5000 for Lenovo is not encrypted.




Figure 13-1 Encryption on single site

To enable encryption of both data copies, the IBM Storwize V5000 for Lenovo must be replaced by an encryption capable system, as shown in Figure 13-2 on page 708. After such replacement both copies of data are encrypted, but the Remote Copy communication between both sites remains unencrypted.




Figure 13-2 Encryption on both sites

Figure 13-3 shows an example configuration that uses both software and hardware encryption. Software encryption is used to encrypt an external virtualized storage system. Hardware encryption is used for internal, SAS-attached disk drives.




Figure 13-3 Example of software encryption and hardware encryption

Placement of hardware encryption and software encryption in the Lenovo Storage code stack are shown in Figure 13-4 on page 709. The functions that are implemented in software are shown in blue. The external storage system is shown in yellow. The hardware encryption on the SAS chip is marked in pink. Compression is performed before encryption. Therefore, it is possible to realize benefits of compression for the encrypted data.

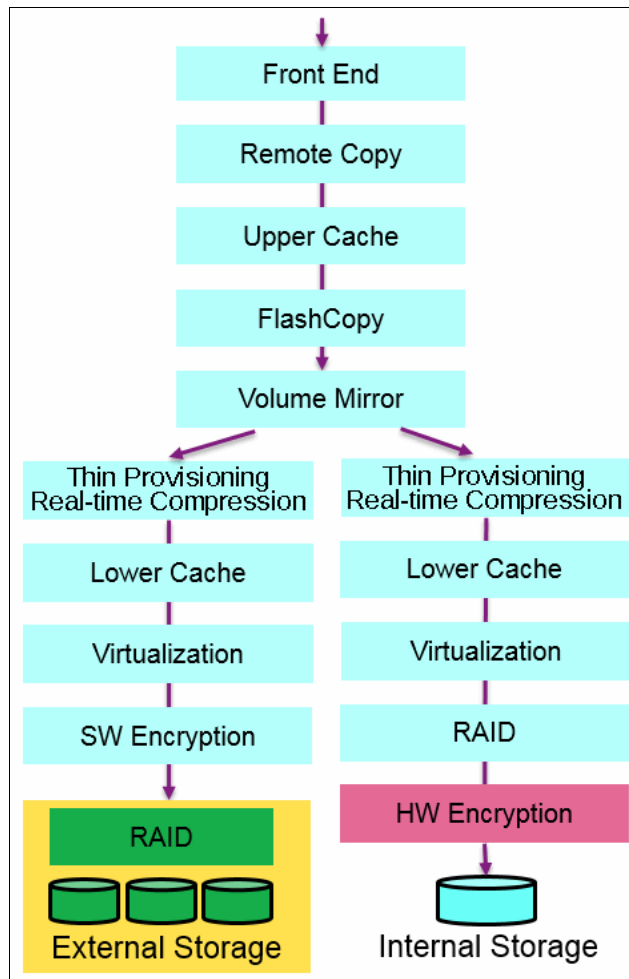


Figure 13-4 Encryption software stack

Each volume copy can use different encryption methods (hardware, software). It is also allowed to have volume copies with different encryption status (encrypted versus unencrypted). The encryption method depends only on the pool that is used for the specific copy. You can migrate data between different encryption methods by using volume migration or volume mirroring.

13.2.1 Encryption methods

There are two ways to perform encryption Lenovo storage V series systems: hardware encryption and software encryption. Both methods of encryption protect against the potential exposure of sensitive user data that are stored on discarded, lost, or stolen media. Both can also facilitate the warranty return or disposal of hardware. Which method is used for encryption is chosen automatically by the system based on the placement of the data:

- ▶ Hardware encryption: Data is encrypted by using Serial Attached SCSI (SAS) hardware. Used only for internal storage.
- ▶ Software encryption: Data is encrypted by using nodes' CPU (encryption code leverages AES-NI CPU instruction set). Used only for external storage.

Note: Software encryption is available in code V7.6 and later.

Both methods of encryption use the same encryption algorithm, the same key management infrastructure, and the same license.

Note: The design for encryption is based on the concept that a system should either be encrypted or not encrypted. Encryption implementation is intended to encourage solutions that contain only encrypted volumes or only unencrypted volumes. For example, once encryption is enabled on the system, all new objects (e.g. pools) are by default created as encrypted.

13.2.2 Encryption keys

Hardware and software encryption use the same encryption key infrastructure. The only difference is the object that is encrypted by using the keys. The following objects can be encrypted:

- ▶ Pools (software encryption)
- ▶ Child pools (software encryption)
- ▶ Arrays (hardware encryption)

Encryption keys can be described as follows:

- ▶ Keys are unique for each object, and they are created when the object is created.
- ▶ Two types of keys are defined in the system:
 - Master access key:
 - The master access key is created when encryption is enabled.
 - The master access key can be stored on USB flash drives and/or a key server(s). One master access key is created for each enabled encryption key provider.
 - It can be copied or backed up as necessary.
 - It is *not* permanently stored anywhere in the system.
 - It is required at boot time to unlock access to encrypted data.
 - Data encryption keys (one for each encrypted object):
 - Data encryption keys are used to encrypt data. When an encrypted object (such as an array, a pool, or a child pool) is created, a new data encryption key is generated for this object.
 - Managed disks (MDisk) that are not self-encrypting are automatically encrypted using the data encryption key of the pool or child pool that they belong to.
 - Managed Disks (MDisks) that are self-encrypting will not be re-encrypted using the data encryption key of the pool or child pool they belong to by default. You can override this default by manually configuring the MDisk as not self-encrypting.
 - Data encryption keys are stored in secure memory.
 - During cluster internal communication data encryption keys are encrypted with the master access key.
 - Data encryption keys cannot be viewed.
 - Data encryption keys cannot be changed.
 - When an encrypted object is deleted, its data encryption key is discarded (*secure erase*).

Important: If all master access key copies are lost and the system must cold reboot, all encrypted data is gone. No method exists, even for Lenovo, to decrypt the data without the keys. If encryption is enabled and the system cannot access the master access key, all SAS hardware is offline, including unencrypted arrays.

13.2.3 Encryption licenses

Encryption is a licensed feature that uses key-based licensing.

No trial licenses for encryption exist on the basis that when the trial runs out, the access to the data would be lost. Therefore, you must purchase an encryption license before you activate encryption. Licenses are generated by IBM Data storage feature activation (DSFA) based on the serial number (S/N) and the machine type and model number (MTM) of the nodes.

You can activate an encryption license during the initial system setup (on the **Encryption** screen of the initial setup wizard) or later on, in the running environment.

Contact your Lenovo sales representative or Lenovo Business Partner to purchase an encryption license.

13.3 Activating encryption

The first step to use encryption is to activate your encryption license.

Activation of the license can be performed in one of two ways: Automatically or manually. Both methods are available during the initial system setup and on the running system.

13.3.1 Obtaining an encryption license

You must purchase an encryption license before you activate encryption. If you did not purchase a license, contact a Lenovo sales representative or Lenovo Business Partner to purchase an encryption license.

When you purchase a license, you should receive a function authorization document with an authorization code printed on it. This code allows you to proceed using the automatic activation process.

If the automatic activation process fails or if you prefer using the manual activation process, use this page to retrieve your license keys:

<s://www.ibm.com/storage/dsfa/storwize/selectMachine.wss>

Ensure that you have the following information:

- ▶ Machine type (MT)
- ▶ Serial number (S/N)
- ▶ Machine signature
- ▶ Authorization code

See 13.3.5, “Activate the license manually” on page 717 for instructions about how to retrieve the machine signature of a node.

13.3.2 Start activation process during initial system setup

One of the steps in the initial setup enables encryption license activation. The system asks “Was the encryption feature purchased for this system?”. To activate encryption at this stage, follow these steps:

1. Select **Yes**, as shown in Figure 13-5.




Figure 13-5 Encryption activation during initial system setup

2. The Encryption window displays information about your storage system, as shown in Figure 13-6 on page 713.

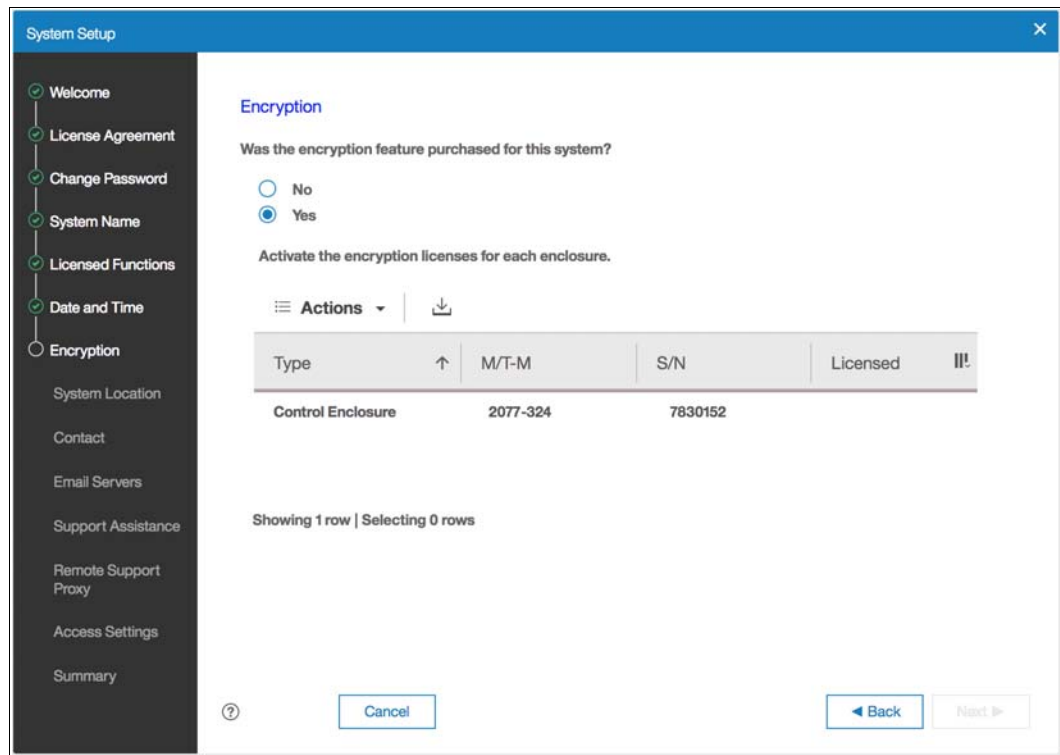


Figure 13-6 Information storage system during initial system setup

3. Right clicking on the node opens a context menu with two license activation options (**Activate License Automatically** and **Activate License Manually**), as shown in Figure 13-7. Use either option to activate encryption. See 13.3.4, “Activate the license automatically” on page 715 for instructions about how to complete an automatic activation process. See “Activate the license manually” on page 717 for instructions on how to complete a manual activation process.

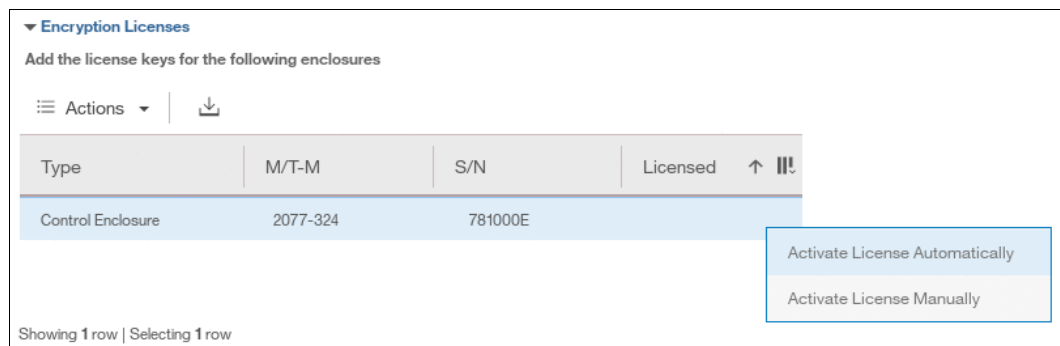


Figure 13-7 Selecting license activation method

4. After either activation process is complete, you can see a green check mark in the column labeled **Licensed** next to a node for which the license was enabled and you can proceed with the initial system setup by clicking **Next**, as shown in Figure 13-8 on page 714.

Note: Every enclosure needs an active encryption license before you can enable encryption on the system.

Attempting to add a non-licensed enclosure to an encryption-enabled system will fail.




Figure 13-8 Successful encryption license activation during initial system setup

13.3.3 Start activation process on a running system

To activate encryption on a running system, follow these steps:

1. Click **Settings** → **System** → **Licensed Functions** and click **Encryption Licenses**, as shown in Figure 13-9.

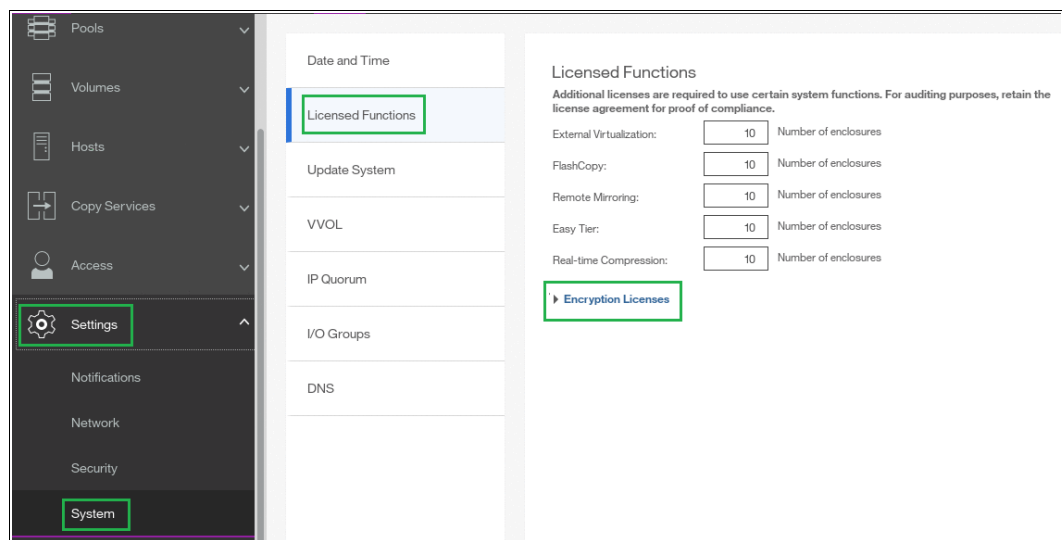


Figure 13-9 Expanding Encryption Licenses section on the Licensed Functions view

2. The Encryption Licenses window displays information about your nodes. Right click on the node on which you want to install an encryption license. This will open a context menu with two license activation options (**Activate License Automatically** and **Activate License Manually**), as shown in Figure 13-10 on page 715. Use either option to activate

encryption. See 13.3.4, “Activate the license automatically” on page 715 for instructions on how to complete an automatic activation process. See 13.3.5, “Activate the license manually” on page 717 for instructions on how to complete a manual activation process.

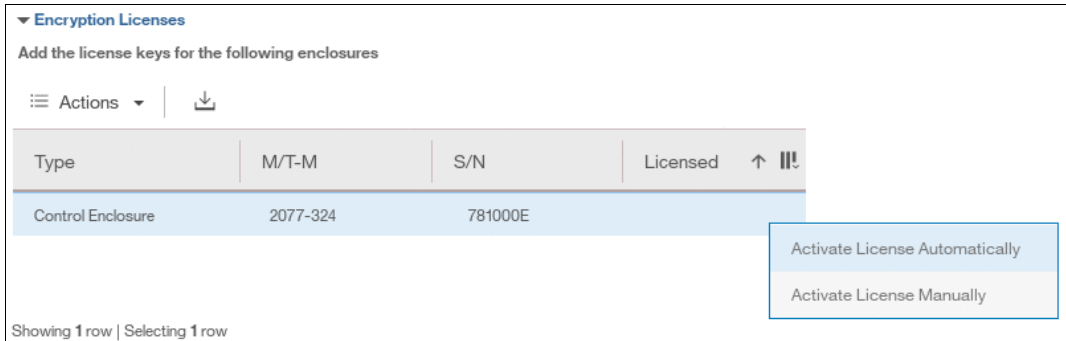


Figure 13-10 Select the node on which you want to enable the encryption

- After either activation process is complete, you can see a green check mark in the column labeled **Licensed** for the node, as shown in Figure 13-11.




Figure 13-11 Successful encryption license activation on a running system

13.3.4 Activate the license automatically

Important: To perform this operation, the personal computer that is used to connect to the GUI and activate the license must be able to connect to the Internet.

To activate the encryption license for a node automatically, follow this procedure:

- Select **Activate License Automatically**, the Activate License Automatically window opens, as shown in Figure 13-12 on page 716.

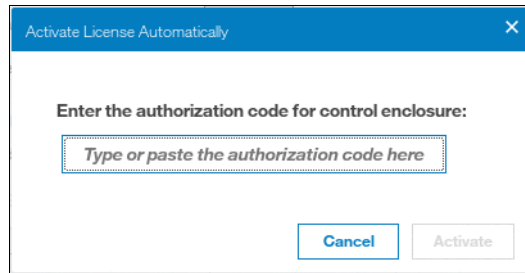


Figure 13-12 Encryption license Activate License Automatically window

2. Enter the authorization code that is specific to the node that you selected, as shown in Figure 13-13. You can now click **Activate**.

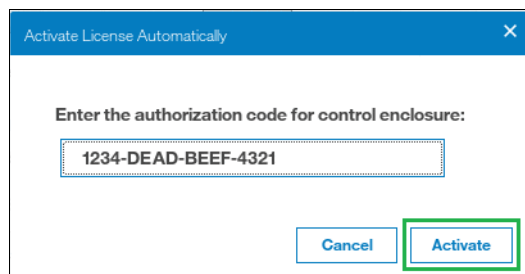


Figure 13-13 Entering an authorization code

3. The system connects to Lenovo to verify the authorization code and retrieve the license key. Figure 13-14 shows a window which is displayed during this connection. If everything works correctly, the procedure takes less than a minute.

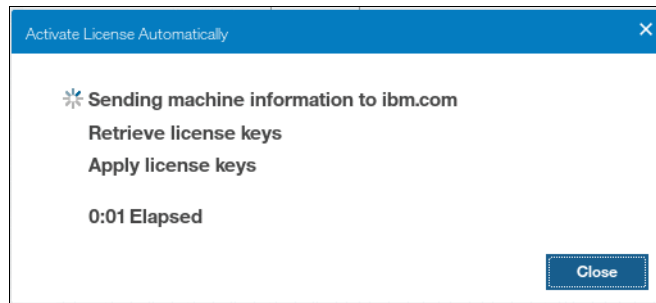


Figure 13-14 Activating encryption

4. After the license key has been retrieved, it is automatically applied as shown in Figure 13-15 on page 717.

▼ Encryption Licenses

Add the license keys for the following enclosures

☰ Actions ▾ | ⬇

Type	M/T-M	S/N	Licensed	↑	⌵
Control Enclosure	2077-324	781000E	✓		

Showing 1 row | Selecting 1 row

Figure 13-15 Successful encryption license activation

Problems with automatic license activation

If connections problems occur with the automatic license activation procedure, the system times out after 3 minutes with an error.

Check whether the personal computer that is used to connect to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI and activate the license can access the internet. If you are unable to complete the automatic activation procedure, try to use the manual activation procedure that is described in 13.3.5, “Activate the license manually” on page 717.

Although authorization codes and encryption license keys use the same format (four groups of four hexadecimal digits), you can only use each of them in the appropriate activation process. If you use a license key when the system expects an authorization code, the system will display an error message, as shown in Figure 13-16.

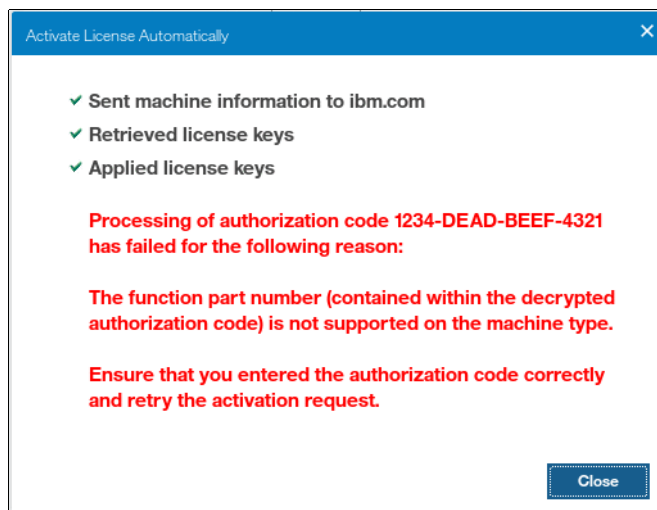


Figure 13-16 Authorization code failure

13.3.5 Activate the license manually

To manually activate the encryption license for a node, follow this procedure:

1. Select **Activate License Manually**, the Manual Activation window opens, as shown in Figure 13-17 on page 718.

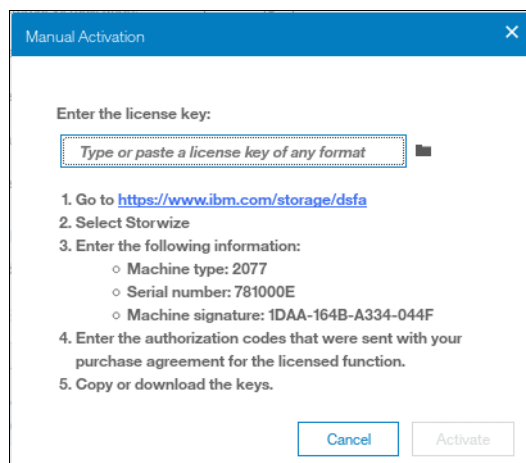


Figure 13-17 Manual encryption license activation window

2. If you have not done so already, you need to obtain the encryption license for the node. The information required to obtain the encryption license is displayed in the *Manual Activation* window. Use this data to follow the instructions in 13.3.1, “Obtaining an encryption license” on page 711.
3. You can enter the license key either by typing it, by using cut or copy and paste, or by clicking the folder icon and uploading to the storage system the license key file downloaded from DSFA. In Figure 13-18, the sample key is already entered. You can now click **Activate**.

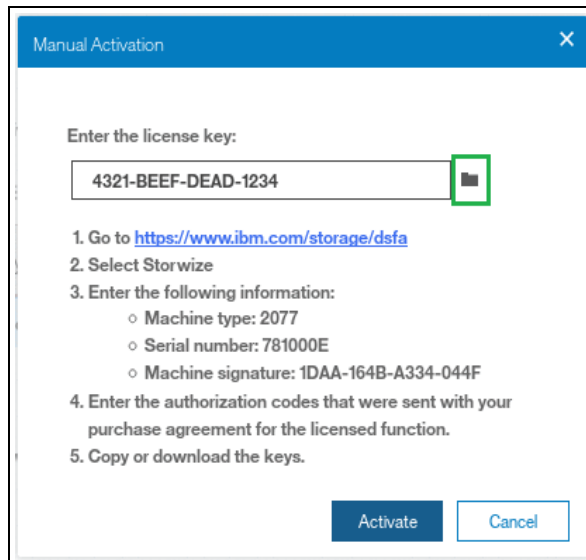


Figure 13-18 Entering an encryption license key

4. Once the task completes successfully, the GUI shows that encryption is licensed for the given node, as shown in Figure 13-19 on page 719.

▼ Encryption Licenses

Add the license keys for the following enclosures

⋮ Actions ▾ | ⬇

Type	M/T-M	S/N	Licensed	⬆ ⚠
Control Enclosure	2077-324	781000E	✓	

Showing 1 row | Selecting 1 row

Figure 13-19 Successful encryption license activation

Problems with manual license activation

Although authorization codes and encryption license keys use the same format (four groups of four hexadecimal digits), you can only use each of them in the appropriate activation process. If you use an authorization code when the system expects a license key, the system will display an error message, as shown in Figure 13-20.

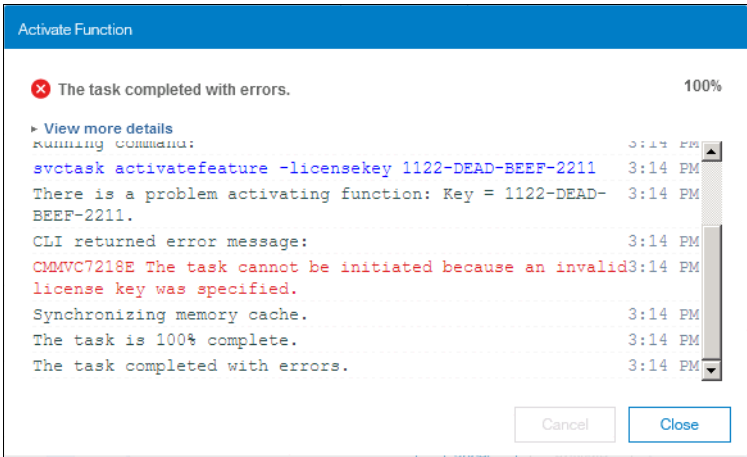


Figure 13-20 License key failure

13.4 Enabling encryption

This section describes the process to create and store system master access key copies, also referred to as encryption keys. These keys can be stored on any or both of two key providers: USB flash drives or a key server.

Key server support is available in controller firmware code V7.8 and later. Additionally controller firmware V8.1 introduces the ability to define up to four encryption key servers, which is a recommended configuration, as it increases key provider availability.

Support for simultaneous use of both USB flash drives and a key server is available in controller firmware V8.1 and later. Organizations that use encryption key management servers might consider parallel use of USB flash drives as a backup solution. During normal operation such drives could be disconnected and stored in a secure location. However, in the event of a catastrophic loss of encryption servers, the USB drives could still be used to unlock the encrypted storage.

The following list of key server and USB flash drive characteristics might help you to choose the type of encryption key provider that you want to use.

Key servers can have the following characteristics:

- ▶ Physical access to the system is not required to perform a rekey operation.
- ▶ Support for businesses that have security requirements precluding use of USB ports.
- ▶ Possibility to use Hardware Security Modules (HSMs) for encryption key generation.
- ▶ Ability to replicate keys between servers and perform automatic backups.
- ▶ Implementations follow an open standard (KMIP) that aids in interoperability.
- ▶ Ability to audit operations related to key management.
- ▶ Ability to separately manage encryption keys and physical access to storage systems.

USB flash drives have the following characteristics:

- ▶ Physical access to the system might be required to process a rekey operation.
- ▶ No moving parts with almost no read or write operations to the USB flash drive.
- ▶ Inexpensive to maintain and use.
- ▶ Convenient and easy to have multiple identical USB flash drives available as backups. You can just copy them if needed.

Important: Maintaining confidentiality of the encrypted data hinges on security of the encryption keys. Pay special attention to ensuring secure creation, management and storage of the encryption keys.

13.4.1 Starting the Enable Encryption wizard

After the license activation step is successfully completed, you can now enable encryption. You can enable encryption after completion of the initial system setup using either GUI or CLI. There are two ways in the GUI to start the **Enable Encryption** wizard. It can be started by clicking *Run Task* button next to **Enable Encryption** on the *Suggested Tasks* window as shown in Figure 13-21.

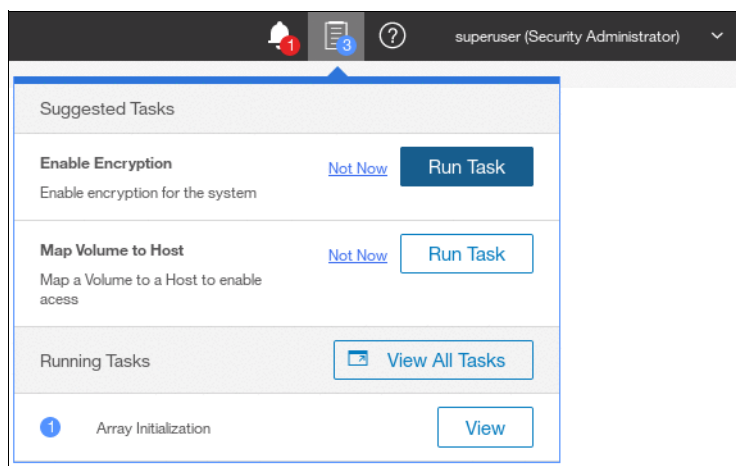


Figure 13-21 Enable Encryption from the Suggested Tasks window

You can also click **Settings** → **Security** → **Encryption** and click **Enable Encryption**, as shown in Figure 13-22 on page 721.

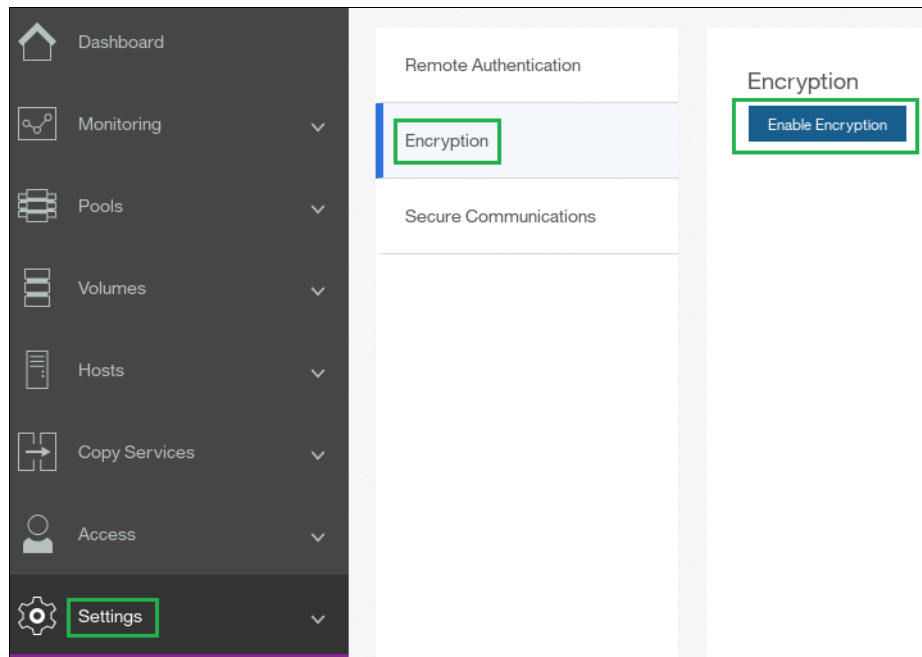


Figure 13-22 Enable Encryption from the Security panel

The **Enable Encryption** wizard starts by asking which encryption key provider to use for storing the encryption keys, as shown in Figure 13-23. You can enable either or both providers.

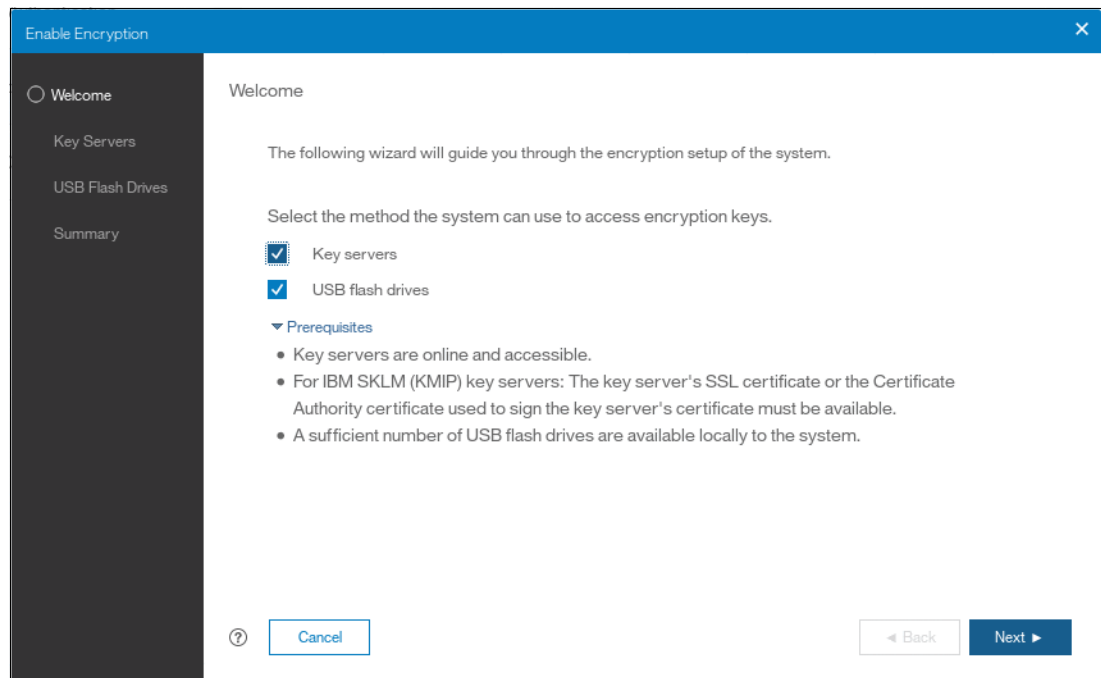


Figure 13-23 Enable Encryption wizard Welcome screen

The next section will present a scenario in which both encryption key providers are enabled at the same time. See 13.4.2, “Enabling encryption using USB flash drives” on page 722 for instructions on how to enable encryption using only USB flash drives provider. See 13.4.3,

“Enabling encryption using key servers” on page 726 for instructions on how to enable encryption using key server(s) as the sole encryption key provider.

13.4.2 Enabling encryption using USB flash drives

Note: The system needs at least three USB flash drives to be present before you can enable encryption using this encryption key provider. Lenovo USB flash drives are recommended, although other flash drives might work. You can use any USB ports in any node of the cluster. After creating the USB flash drives you can copy them if you need more than four.

Using USB flash drives as the encryption key provider requires a minimum of three USB flash drives to store the generated encryption keys. Because the system will attempt to write the encryption keys to any USB key inserted into a node port, it is critical to maintain physical security of the system during this procedure.

While the system enables encryption, you are prompted to insert USB flash drives into the system. The system generates and copies the encryption keys to all available USB flash drives.

Ensure that each copy of the encryption key is valid before you write any user data to the system. The system validates any key material on a USB flash drive when it is inserted into the canister. If the key material is not valid, the system logs an error. If the USB flash drive is unusable or fails, the system does not display it as output. Figure 13-79 on page 757 shows an example where the system detected and validated three USB flash drives.

If your system is in a secure location with controlled access, one USB flash drive for each canister may remain inserted in the system. If there is a risk of unauthorized access, then all USB flash drives with the master access keys must be removed from the system and stored in a secure place.

Securely store all copies of the encryption key. For example, any USB flash drives holding an encryption key copy, that are not left plugged into the system, can be locked in a safe. Similar precautions must be taken to protect any other copies of the encryption key that are stored on other media.

Notes: Generally, create at least one additional copy on another USB flash drive for storage in a secure location. You can also copy the encryption key from the USB drive and store the data on other media, which may provide additional resilience and mitigate risk that the USB drives used to store the encryption key come from a faulty batch.

Every encryption key copy must be stored securely to maintain confidentiality of the encrypted data.

A minimum of one USB flash drive with the correct master access key is required to unlock access to encrypted data after a system restart such as a system-wide reboot or power loss. No USB flash drive is required during a warm reboot, such as a node exiting service mode or a single node reboot. The data center power-on procedure needs to ensure that USB flash drives containing encryption keys are plugged into the storage system before it is powered on.

During power-on, insert USB flash drives into the USB ports on two supported canisters to safeguard against failure of a node, node's USB port, or USB flash drive during the power-on procedure.

To enable encryption using USB flash drives as the only encryption key provider follow these steps:

1. In the Enable Encryption wizard Welcome tab, select **USB flash drives** and click **Next**, as shown in Figure 13-24.

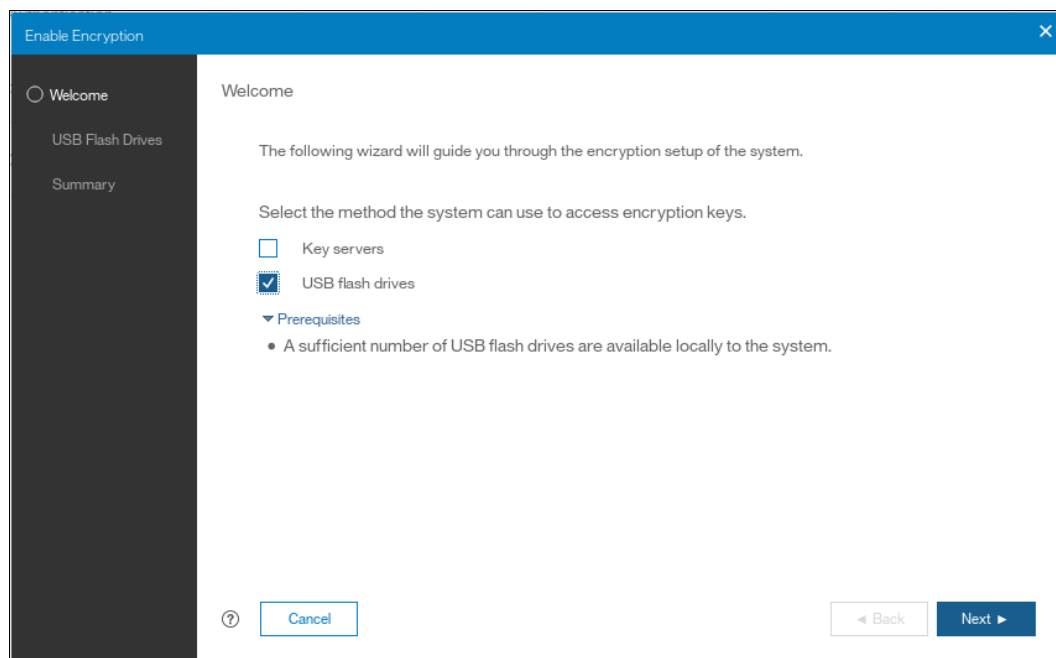


Figure 13-24 Selecting USB flash drives in the Enable Encryption wizard

2. If there are fewer than 3 USB flash drives inserted into the system, you will be prompted to insert additional drives, as shown in Figure 13-25 on page 724. The system will report how many additional drives need to be inserted.




Figure 13-25 Waiting for USB flash drives to be inserted

Note: The **Next** option remains disabled and the status at the bottom is kept at 0 until at least three USB flash drives are detected.

3. Insert the USB flash drives into the USB ports as requested.
4. After the minimum required number of drives is detected, the encryption keys are automatically copied on the USB flash drives, as shown in Figure 13-26.




Figure 13-26 Writing the master access key to USB flash drives

You can keep adding USB flash drives or replacing the ones already plugged in to create new copies. When done, click **Next**.

5. The number of keys that were created is shown in the Summary tab, as shown in Figure 13-27. Click **Finish** to finalize the encryption enablement.

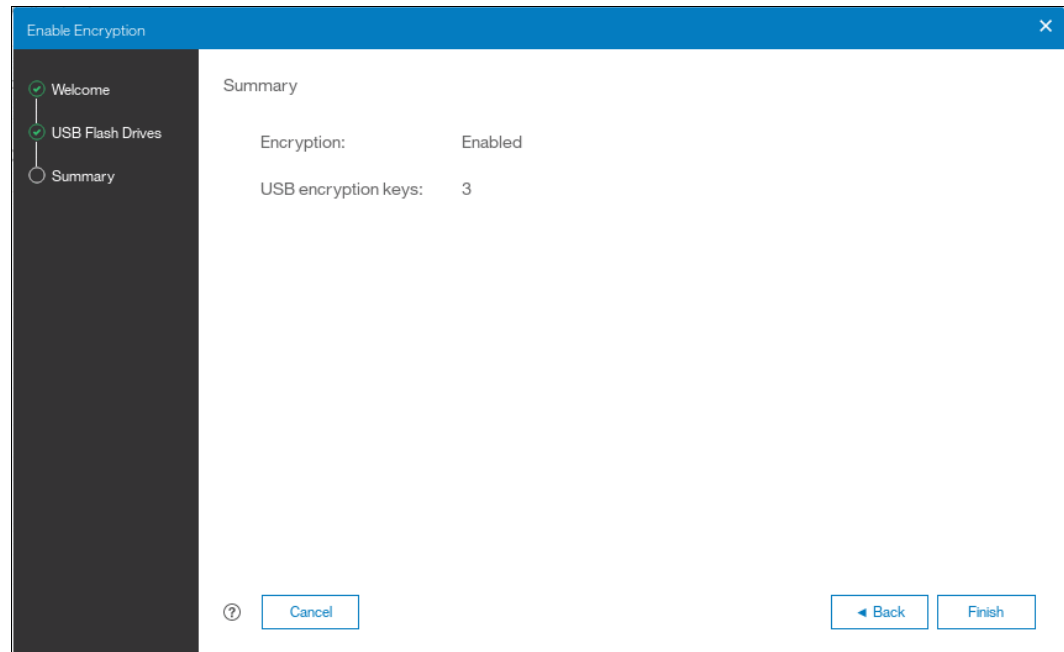


Figure 13-27 Commit the encryption enablement

6. You receive a message confirming that the encryption is now enabled on the system, as shown in Figure 13-28.

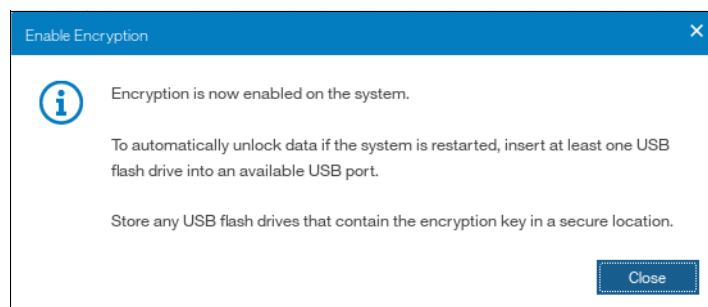


Figure 13-28 Encryption enabled message using USB flash drives

7. You can confirm that encryption is enabled, as well as verify which key providers are in use, by going to **Settings** → **Security** → **Encryption**, as shown in Figure 13-29 on page 726.




Figure 13-29 Encryption view showing using USB flash drives as the enabled provider

13.4.3 Enabling encryption using key servers

A key server is a centralized system that receives and then distributes encryption keys to its clients including Lenovo Storage V series systems.

Lenovo Storage V series system supports use of an IBM Security Key Lifecycle Manager key server as an encryption key provider. SKLM supports Key Management Interoperability Protocol (KMIP), which is a standard for management of cryptographic keys.

Note: Make sure, that the key management server functionality is fully independent from storage provided by systems using a key server for encryption key management. Failure to observe this requirement may create an encryption deadlock. An encryption deadlock is a situation in which none of key servers in the given environment can become operational because some critical part of the data in each server is stored on a storage system that depends on one of the key servers to unlock access to the data.

Controller firmware code V8.1 and later supports up to 4 key server objects defined in parallel.

Before you can create a key server object in the storage system, the key server must be configured. Ensure that you complete the following tasks on the SKLM server before you enable encryption on the storage system:

- ▶ Configure the SKLM server to use Transport Layer Security version 2 (TLSv2). The default setting is TLSv1, but controller firmware supports only version 2.
- ▶ Ensure that the database service is started automatically on startup.
- ▶ Ensure that there is at least one Secure Sockets Layer (SSL) certificate for browser access.
- ▶ Create a SPECTRUM_VIRT device group for Spectrum Virtualize systems.

For more information about completing these tasks, see SKLM documentation at IBM Knowledge Center at:

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

Access to the key server storing the correct master access key is required to enable encryption for the cluster after a system restart such as a system-wide reboot or power loss. Access to the key server is not required during a warm reboot, such as a node exiting service

mode or a single node reboot. The data center power-on procedure must ensure key server availability before storage system using encryption is booted.

To enable encryption using a key server follow these steps:

1. Ensure that you have **service IPs** configured on all your nodes.
2. In the Enable Encryption wizard Welcome tab, select **Key servers** and click **Next**, as shown in Figure 13-40 on page 733.

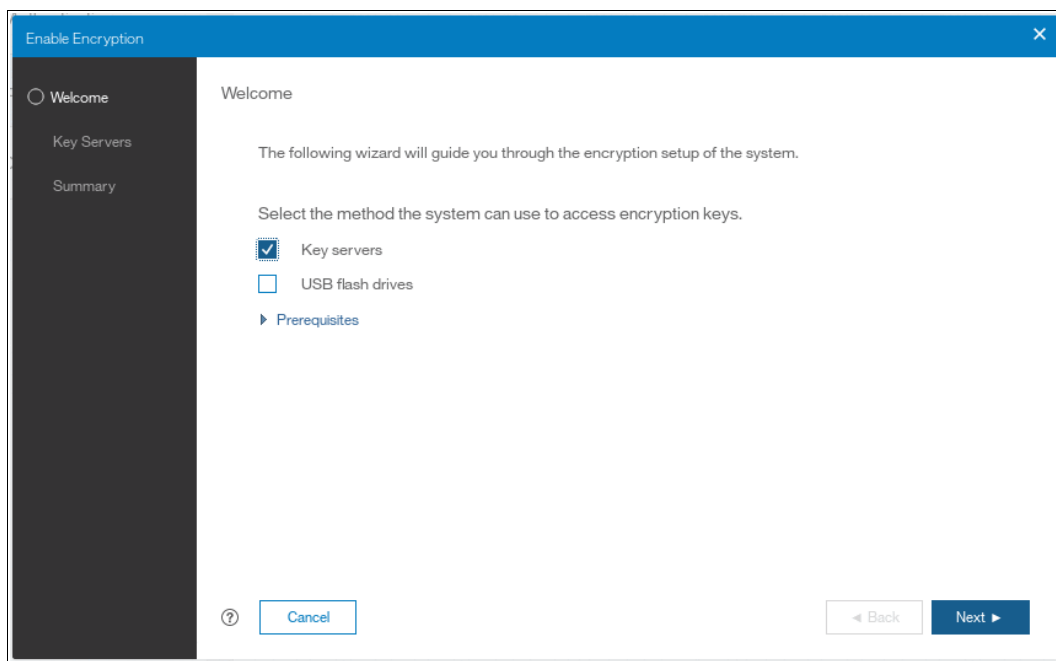


Figure 13-30 Selecting Key server as the only provider in the Enable Encryption wizard

3. The wizard moves to the Key Servers tab, as shown in Figure 13-42 on page 735. Enter the name and **IP address** of the key servers. Note that the first key server specified must be the primary SKLM key server.

Note: The supported versions of Security Key Lifecycle Manager (up to V2.7, which was the latest code version available at the time of writing) differentiate between the primary and secondary key server role. The Primary SKLM server as defined on *Key Servers* screen of Enable Encryption wizard must be the server defined as the primary by SKLM administrators.

The key server name serves just as a label, only the provided IP address will be used to actually contact the server. If the key server's TCP port number differs from the default value for the KMIP protocol (i.e. 5696), then enter the port number. An example of a complete primary SKLM configuration is shown in Figure 13-31 on page 728.

Name	IP Address	Port
primary_ks	9.174.157.2	5696

Figure 13-31 Configuration of the primary SKLM server

- If you want to add additional, secondary SKLM servers, then click on “+” and fill the data for secondary SKLM servers, as shown on Figure 13-32. You can define up to four SKLM servers. Click **Next** when you are done.

Name	IP Address	Port
primary_ks	9.174.157.2	5696
secondaryA_ks	9.174.157.3	5696
secondaryB_ks	9.174.157.4	5696
secondaryC_ks	9.174.157.5	5696

Figure 13-32 Configuring multiple SKLM servers

- The next page in the wizard is a reminder that SPECTRUM_VIRT device group dedicated for controller firmware systems must exist on the SKLM key servers. Make sure that this device group exists and click **Next** to continue, as shown in Figure 13-33 on page 729.

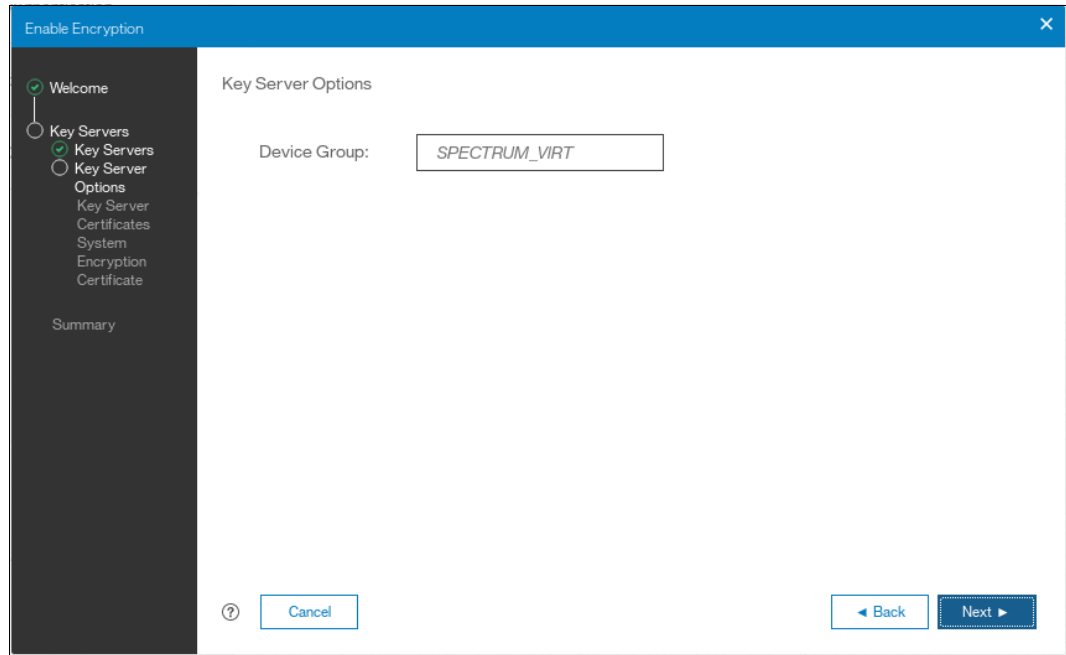


Figure 13-33 Checking key server device group

6. The next step is to enable secure communication between the controller firmware system and the SKLM key servers. This can be done by either uploading the public certificate of the certificate authority (CA) used to sign all the SKLM key server certificates, or by uploading the public SSL certificate of each key server directly. Figure 13-34 shows the case when an organization's CA certificate is used. Click **Next** to proceed to the next step.

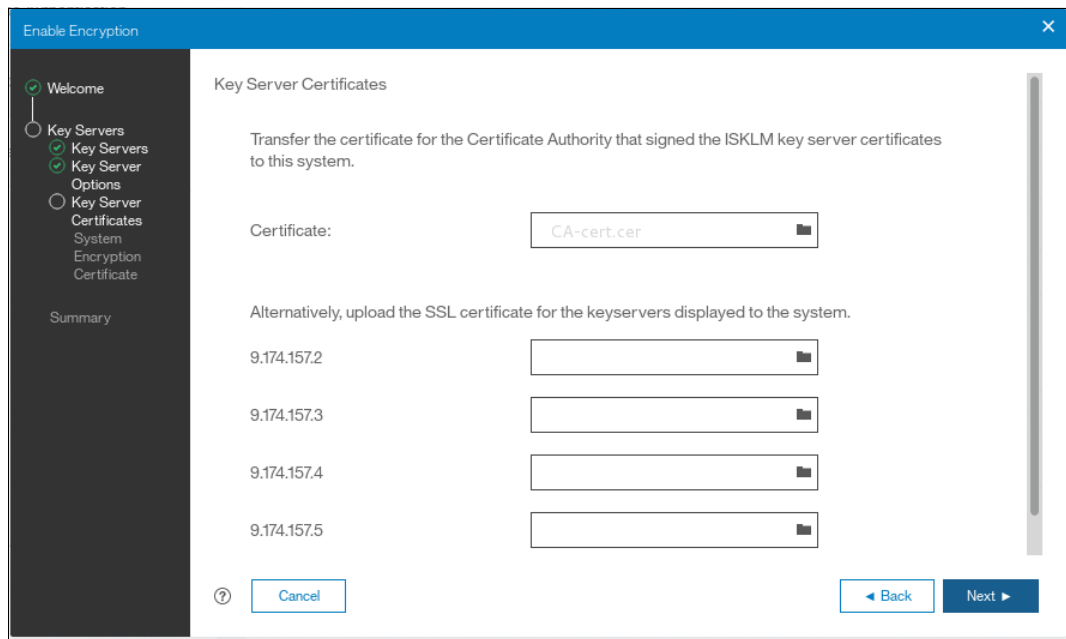


Figure 13-34 Uploading the key server or certification authority SSL certificate

7. Subsequently, configure the SKLM key server to trust the SSL certificate of the controller firmware system. You can download the controller firmware system public SSL certificate

by clicking **Export Public Key**, as shown in Figure 13-35. You should install this certificate in the SKLM key server in the SPECTRUM_VIRT device group.

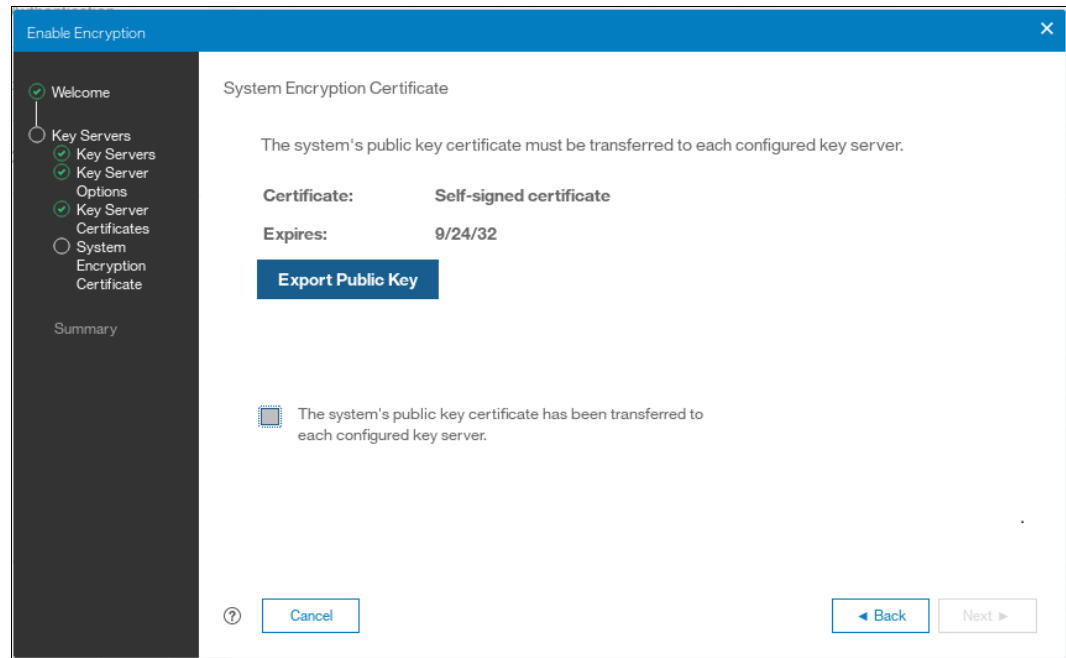


Figure 13-35 Downloading the Spectrum Virtualize SSL certificate

8. When the Spectrum Virtualize system SSL certificate has been installed on the SKLM key server, acknowledge this by selecting the box indicated in Figure 13-36 and click **Next** to proceed to the next step.

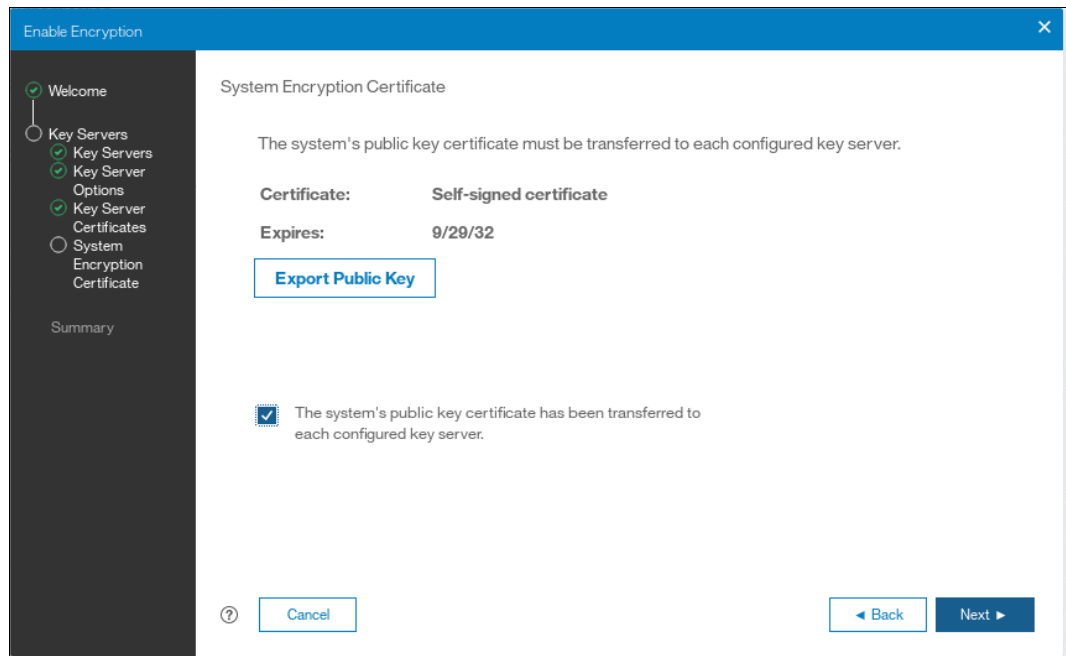


Figure 13-36 Acknowledge Spectrum Virtualize SSL certificate transfer

9. The key server configuration is shown in the Summary tab, as shown in Figure 13-37. Click **Finish** to create the key server object and finalize the encryption enablement.

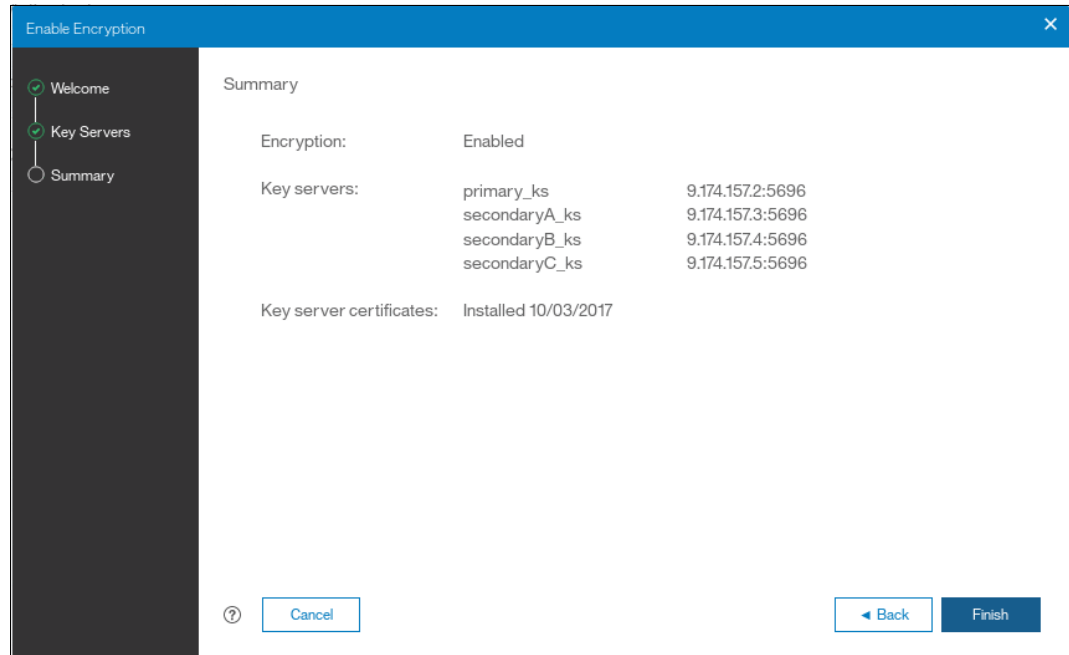


Figure 13-37 Finish enabling encryption using a key server

10. If there are no errors while creating the key server object, you receive a message that confirms that the encryption is now enabled on the system, as shown in Figure 13-38.

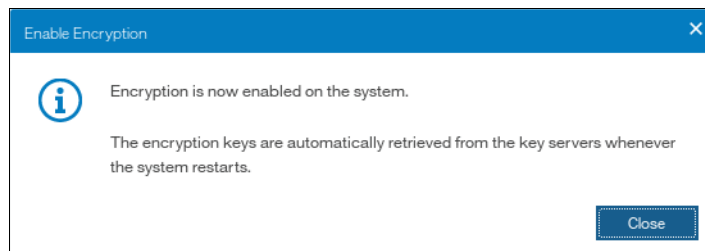


Figure 13-38 Encryption enabled message using a key server

11. Confirm that encryption is enabled in **Settings** → **Security** → **Encryption**, as shown in Figure 13-39 on page 731. Note the four green checks, which indicate, that all four SKLM servers are detected as available by the system.




Figure 13-39 Encryption enabled with only key servers as encryption key providers

13.4.4 Enabling encryption using both providers

Controller firmware code V8.1 and later allows parallel use of both key server and USB flash drive encryption key providers. It is possible to configure both providers in a single run of encryption enable wizard. To perform such configuration, the system must meet requirements of both SKLM and USB flash drive encryption key providers.

Controller firmware supports enabling encryption using a Security Key Lifecycle Manager (SKLM) key server. SKLM supports Key Management Interoperability Protocol (KMIP), which is a standard for encryption of stored data and management of cryptographic keys.

Note: Make sure, that the key management server functionality is fully independent from storage provided by systems using a key server for encryption key management. Failure to observe this requirement may create an encryption deadlock. An encryption deadlock is a situation in which none of key servers in the given environment can become operational because some critical part of the data in each server is stored on an encrypted storage system that depends on one of the key servers to unlock access to the data.

Controller firmware code V8.1 and later supports up to four key server objects defined in parallel.

Before you can create the key server object in a storage system, the key server must be configured. Ensure that you complete the following tasks on the SKLM server before you enable encryption on the storage system:

- ▶ Configure the SKLM server to use Transport Layer Security version 2 (TLSv2). The default setting is TLSv1, but controller firmware supports only version 2.
- ▶ Ensure that the database service is started automatically on startup.
- ▶ Ensure that there is at least one Secure Sockets Layer (SSL) certificate for browser access.
- ▶ Create a SPECTRUM_VIRT device group for controller firmware systems. A device group allows for restricted management of subsets of devices within a larger pool.

For more information about completing these tasks, see SKLM at IBM Knowledge Center at:

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

Access to the key server storing the correct master access key is required to enable encryption for the cluster after a system restart such as a system-wide reboot or power loss. Access to the key server is not required during a warm reboot, such as a node exiting service mode or a single node reboot. The data center power-on procedure must ensure key server availability before storage system using encryption is booted.

To enable encryption using a key server follow these steps:

1. Ensure that you have **service IPs** configured on all your nodes.
2. In the Enable Encryption wizard Welcome tab, select **Key servers** and click **Next**, as shown in Figure 13-40.

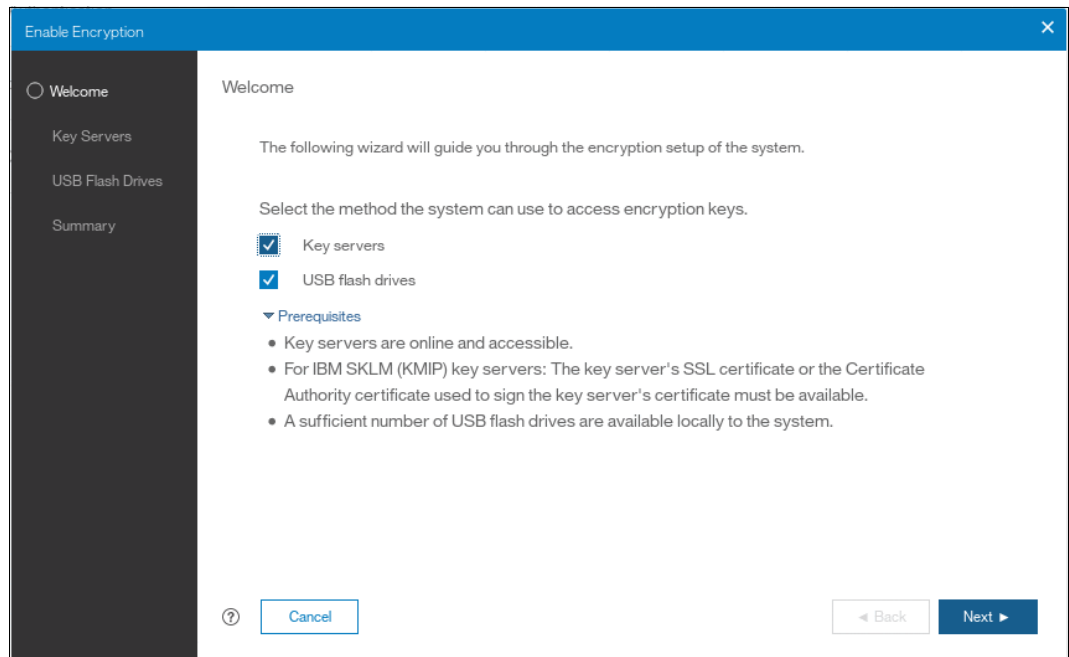


Figure 13-40 Selecting Key servers in the Enable Encryption wizard

3. The wizard moves to the Key Servers tab, as shown in Figure 13-42 on page 735. Enter the name and **IP address** of the key servers. Note that the first key server specified must be the primary SKLM key server.

Note: The supported versions of Security Key Lifecycle Manager (up to V2.7, which was the latest code version when this book was written) differentiate between primary and secondary key server role. Primary SKLM server as defined on *Key Servers* screen of Enable Encryption wizard must be the server defined as primary by SKLM administrators.

Name	IP Address	Port	
primary_ks	9.174.157.2	5696	+ - (Primary)

Figure 13-41 Configuration of the primary SKLM server

4. If you want to add additional, secondary SKLM servers, then click on “+” and fill the data for subsequent SKLM servers, as shown in Figure 13-42 on page 735. You can define up to four SKLM servers. Click **Next** when you are done.

Enable Encryption

Key Servers

Ensure that replication is enabled between each configured key server to allow for redundant copies of the encryption key.

Name	IP Address	Port	
primary_ks	9.174.157.2	5696	(Primary)
secondaryA_ks	9.174.157.3	5696	
secondaryB_ks	9.174.157.4	5696	
secondaryC_ks	9.174.157.5	5696	

Buttons: Cancel, Back, Next

Figure 13-42 Configuring multiple SKLM servers

- The next page in the wizard is a reminder that SPECTRUM_VIRT device group dedicated for Spectrum Virtualize systems must exist on the SKLM key servers. Make sure that this device group exists and click **Next** to continue, as shown in Figure 13-43.

Enable Encryption

Key Server Options

Device Group: SPECTRUM_VIRT

Buttons: Cancel, Back, Next

Figure 13-43 Checking key server device group

- The next step is to enable secure communication between the Spectrum Virtualize system to and the SKLM key servers. This can be done by either uploading the public certificate of the certificate authority used to sign all the SKLM key server certificates, or by uploading the public SSL certificate of each key server directly. Figure 13-44 on page 736 shows the

case when an organization's CA certificate is used. When either file has been selected, you can click **Next**.

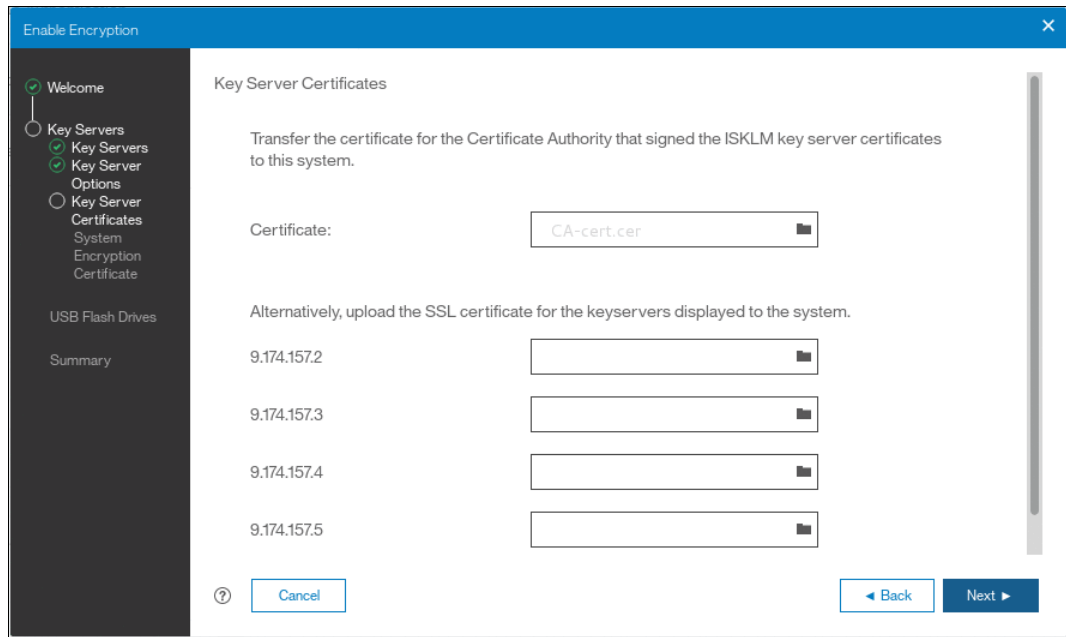


Figure 13-44 Uploading the key server or certification authority SSL certificate

7. Subsequently, configure the SKLM key server to trust the SSL certificate of the controller firmware system. You can download the controller firmware system public SSL certificate by clicking **Export Public Key**, as shown in Figure 13-45. You should install this certificate in the SKLM key servers in the SPECTRUM_VIRT device group.




Figure 13-45 Downloading the controller firmware SSL certificate

8. When the controller firmware system SSL certificate has been installed on the SKLM key server, acknowledge this by selecting the box indicated in Figure 13-46 and click **Next** to proceed to the next step.

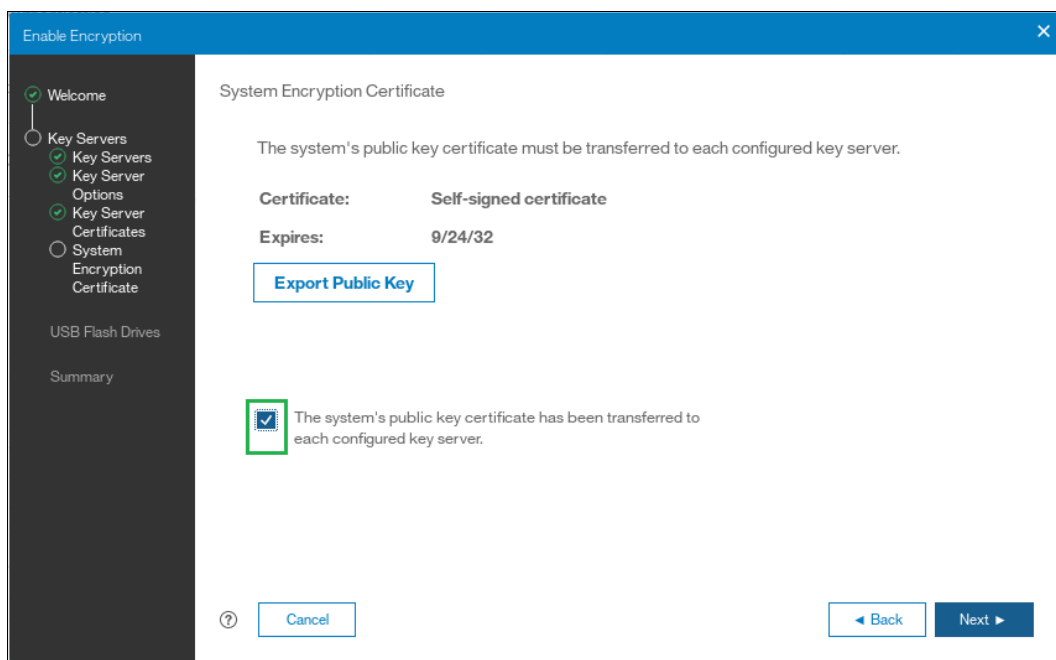


Figure 13-46 Acknowledge controller firmware SSL certificate transfer

9. The next step in the wizard is to store the master encryption key copies on USB flash drives. If there are fewer than three drives detected, the system will request plugging additional USB flash drives as shown on Figure 13-47. You cannot proceed until the required minimum number of USB flash drives is detected by the system.

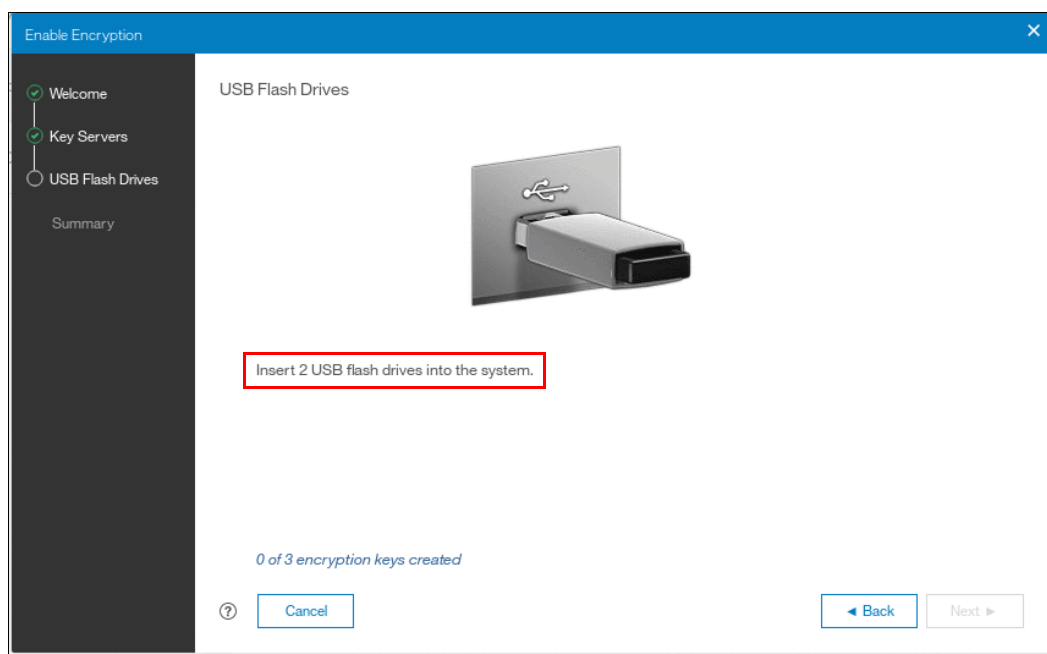


Figure 13-47 At least 3 USB flash drives are required to configure USB flash drive key provider

10. Once at least three USB flash drives are detected, the system will write master access key to each of the drives. Note that the system will attempt to write the encryption key to any flash drive it detects. Therefore, it is crucial to maintain physical security of the system during this procedure. Once the keys are successfully copied to at least three USB flash drives, the system will display a screen as shown in Figure 13-48.




Figure 13-48 Master Access Key successfully copied to USB flash drives

11. The next screen presents you with the summary of the configuration that will be implemented on the system, see Figure 13-49. Click **Finish** to create the key server object and finalize the encryption enablement.

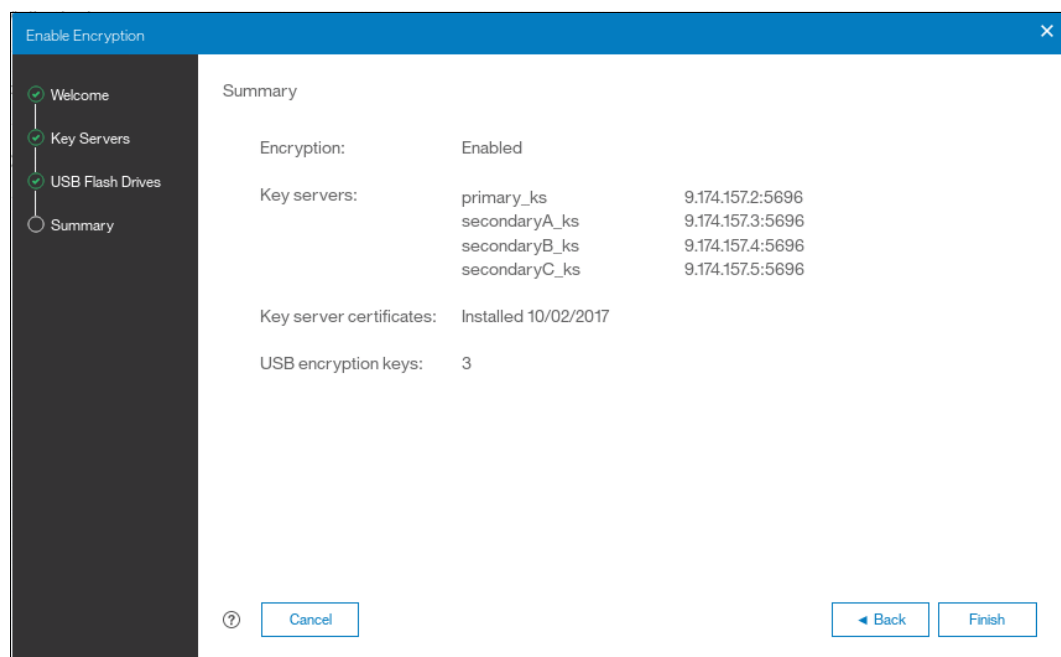


Figure 13-49 Encryption configuration summary in two providers scenario

12. If there are no errors while creating the key server object, the system displays a screen that confirms that the encryption is now enabled on the system, and that both encryption key providers are enabled (see Figure 13-50).

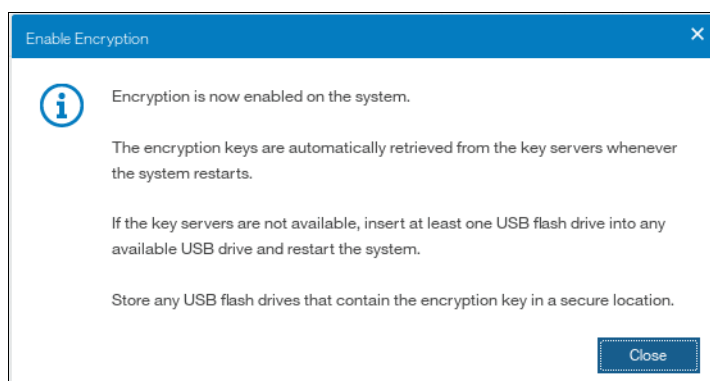


Figure 13-50 Encryption enabled message using both encryption key providers

13. You can confirm that encryption is enabled, as well as verify which key providers are in use, by going to **Settings** → **Security** → **Encryption**, as shown in Figure 13-51. Note four green check marks confirming, that the master access key is available on all four SKLM servers.




Figure 13-51 Encryption enabled with two key providers available

13.5 Configuring additional providers

Once the system is configured with a single encryption key provider, it is possible to add a second provider.

Note: If you set up encryption of your storage system when it was running controller firmware code version earlier than V7.8.0, then when you upgrade to code version V8.1 you have to rekey the master encryption key before you can enable second encryption provider.

13.5.1 Adding SKLM as a second provider

If the storage system is configured with the USB flash drive provider, it is possible to configure SKLM server(s) as a second provider. To enable SKLM server(s) as a second provider follow these steps:

1. Go to **Settings** → **Security** → **Encryption**, expand the *Key Servers* section and click on **Enable**, as shown in Figure 13-52 on page 740. Note that to enable key server as a

second provider, the system has to detect at least one USB flash drive with a current copy of the master access key.

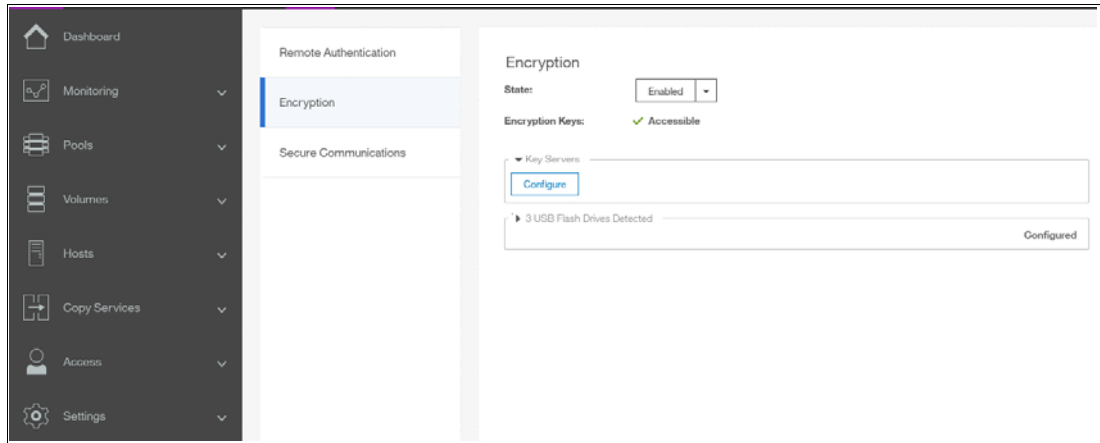


Figure 13-52 Enable SKLM server as a second provider

2. Subsequently, follow the steps required to configure the key server provider, as described in 13.4.3, “Enabling encryption using key servers” on page 726. One difference to the process described in that section is that the wizard will give you an option to migrate from the USB flash drive provider to key server provider. Select **No** to enable both encryption key providers, as shown in Figure 13-53.

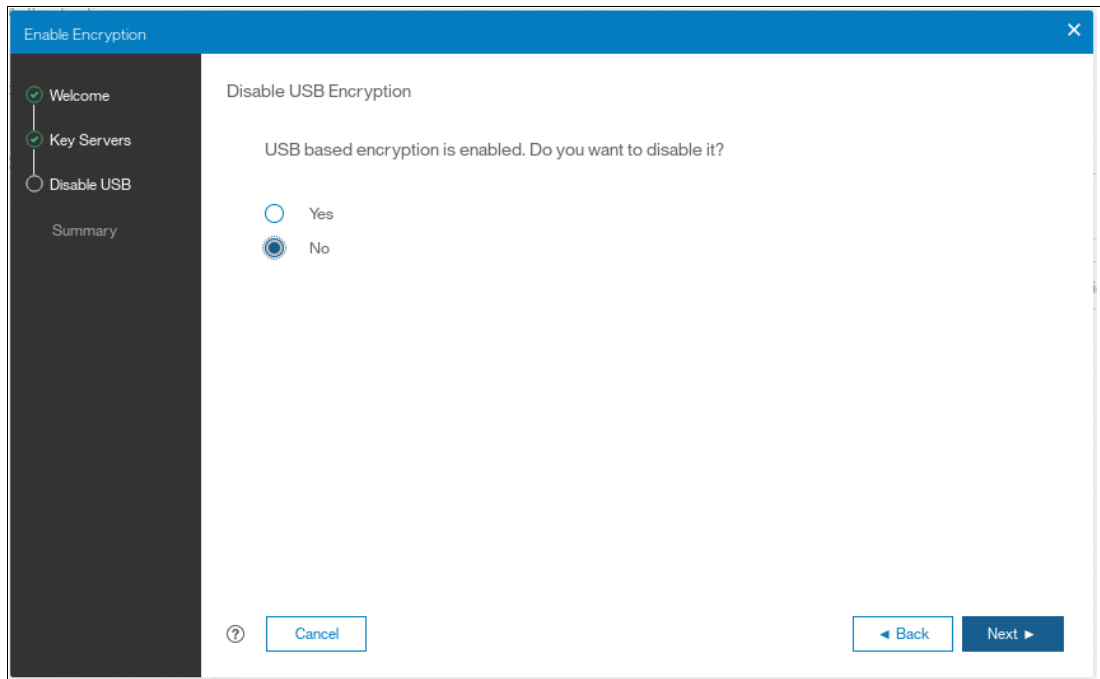


Figure 13-53 Do not disable USB flash drive encryption key provider

3. This choice is confirmed on the summary screen before the configuration is committed, as shown in Figure 13-54 on page 741.




Figure 13-54 Configuration summary before committing

4. After you click finish, the system will configure SKLM servers as a second encryption key provider. Successful completion of the task will be confirmed by a message as in Figure 13-55.

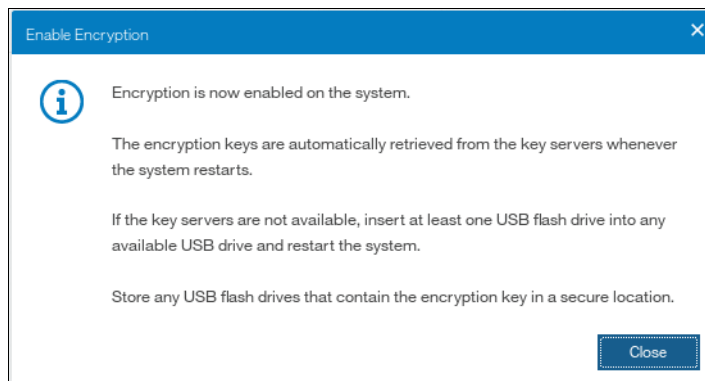


Figure 13-55 Confirmation of successful configuration of two encryption key providers

5. You can confirm that encryption is enabled, as well as verify which key providers are in use, by going to **Settings** → **Security** → **Encryption**, as shown in Figure 13-56 on page 742. Note four green check marks confirming that the master access key is available on all four SKLM servers.




Figure 13-56 Encryption enabled with two key providers available

13.5.2 Adding USB flash drives as a second provider

If the storage system is already configured with an SKLM encryption key provider, it is possible to configure USB flash drives as a second provider. To enable USB flash drives as a second provider follow these steps:

1. Go to **Settings** → **Security** → **Encryption**, expand the *USB Flash Drives* section and click on **Configure**, as shown in Figure 13-57. Please note, that to enable USB flash drives as a second provider, the system has to be able to access key servers with the current master access key.

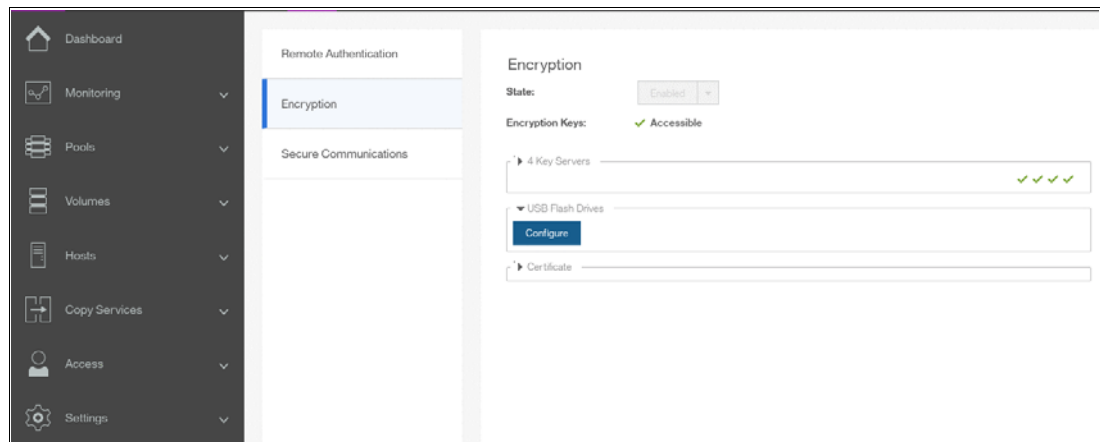


Figure 13-57 Enable USB flash drives as a second encryption key provider

2. After you click on **Configure**, you will be presented with a wizard similar to described in 13.4.2, “Enabling encryption using USB flash drives” on page 722. Note that you will not be given an option to disable SKLM provider during this process. After successful completion of the process you will be presented with a message confirming that both encryption key providers are enabled, as shown in Figure 13-58 on page 743.

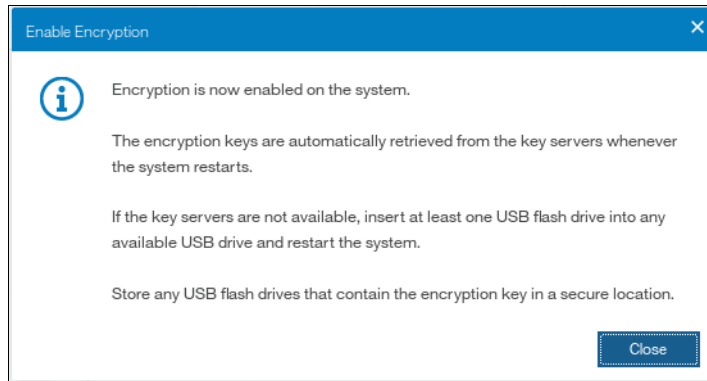


Figure 13-58 Confirmation of successful configuration of two encryption key providers

3. You can confirm that encryption is enabled, as well as verify which key providers are in use, by going to **Settings** → **Security** → **Encryption**, as shown in Figure 13-59. Note four green check marks indicating that the master access key is available on all four SKLM servers.




Figure 13-59 Encryption enabled with two key providers available

13.6 Migrating between providers

Controller firmware V8.1 introduced support for simultaneous use of both USB flash drives and a key server as encryption key providers. The system also allows migration from configuration using only USB flash drive provider to key servers provider and vice versa.

13.6.1 Migration from USB flash drive provider to encryption key server

The system is designed to facilitate migration from USB flash drives encryption key provider to encryption key server provider. If you follow the steps described in 13.5.1, “Adding SKLM as a second provider” on page 739, but when executing procedure step 2 on page 740 select **Yes** instead of **No** (see Figure 13-60 on page 744). This will cause de-activation of the USB Flash drives provider, and the procedure will complete with a single active encryption keys provider — SKLM server(s).

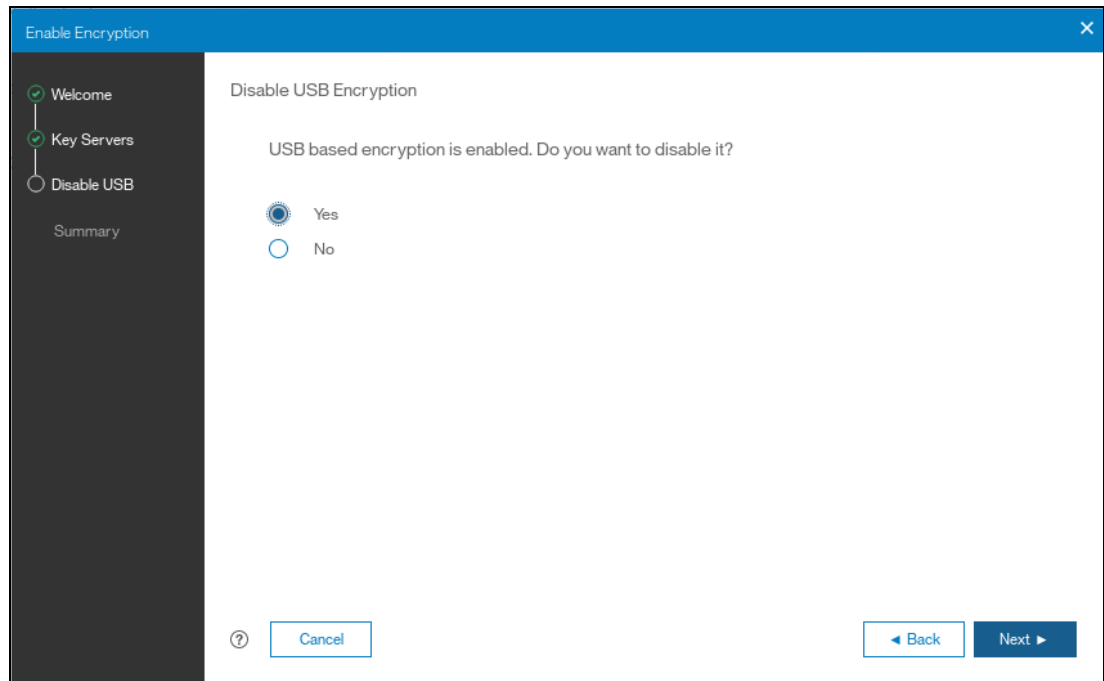


Figure 13-60 Disable USB Flash drive provider while migrating to SKLM provider

13.6.2 Migration from encryption key server to USB flash drive provider

Migration in the other direction, that is to say from using encryption key servers provider to USB flash drives provider, is not possible using only the GUI.

To perform the migration, add USB flash drives as a second provider. You can do this by following steps described in 13.5.2, “Adding USB flash drives as a second provider” on page 742. Subsequently in the CLI issue the following command:

chencryption -usb validate

to make sure that USB drives contain the correct master access key. Subsequently, disable the encryption key server provider by running the following command:

chencryption -keyserver disable

This will disable the encryption key server provider, effectively migrating your system from encryption key server to USB flash drive provider.

13.7 Recovering from a provider loss

If you have both encryption key providers enabled, and you lose one of them (by losing all copies of the encryption key kept on the USB flash drives or by losing all SKLM servers), you can recover from this situation by disabling the provider to which you lost the access. In order to disable the unavailable provider you must have access to a valid master access key on the remaining provider.

If you have lost access to the encryption key server provider, then run the command:

chencryption -keyserver disable

If you have lost access to the USB flash drives provider, then run the command

chencryption -usb disable

If you want to restore the configuration with both encryption key providers, then follow the instructions in 13.5, “Configuring additional providers” on page 739.

Note: If you lose access to all encryption key providers defined in the system, then there is no method to recover access to the data protected by the master access key.

13.8 Using encryption

The design for encryption is based on the concept that a system should either be fully encrypted or not encrypted. Encryption implementation is intended to encourage solutions that contain only encrypted volumes or only unencrypted volumes. For example, once encryption is enabled on the system, all new objects (e.g. pools) are by default created as encrypted. Some unsupported configurations are actively policed in code. For example, no support exists for creating unencrypted child pools from encrypted parent pools. However, exceptions exist:

- ▶ During the migration of volumes from unencrypted to encrypted volumes, a system might report both encrypted and unencrypted volumes.
- ▶ It is possible to create unencrypted arrays from CLI by manually overriding the default encryption setting.

Notes: Encryption support for Distributed RAID is available in controller firmware code V7.7 and later.

You must decide whether to encrypt or not encrypt an object when it is created. You cannot change this setting at a later time. To change the encryption state of stored data you have to migrate it from an encrypted object (e.g. pool) to unencrypted one, or vice versa. Volume migration is the only way to encrypt any volumes that were created before enabling encryption on the system.

13.8.1 Encrypted pools

See Chapter 4, “Storage pools” on page 139 for generic instructions on how to open the Create Pool window. After encryption is enabled, any new pool will by default be created as encrypted, as shown in Figure 13-61.

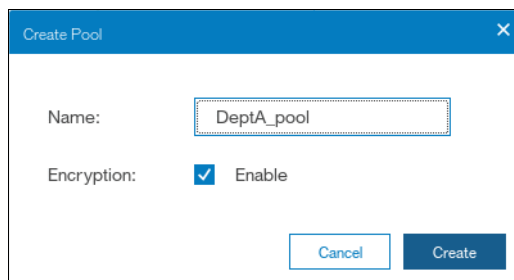


Figure 13-61 Create Pool window basic

You can click **Create** to create an encrypted pool. All storage that is added to this pool will be encrypted.

You can customize Pools view in the management GUI to show pool encryption status. Click **Pools** → and again **Pools**, and then click on the **Actions** → **Customize Columns** → **Encryption**, as shown in Figure 13-62.

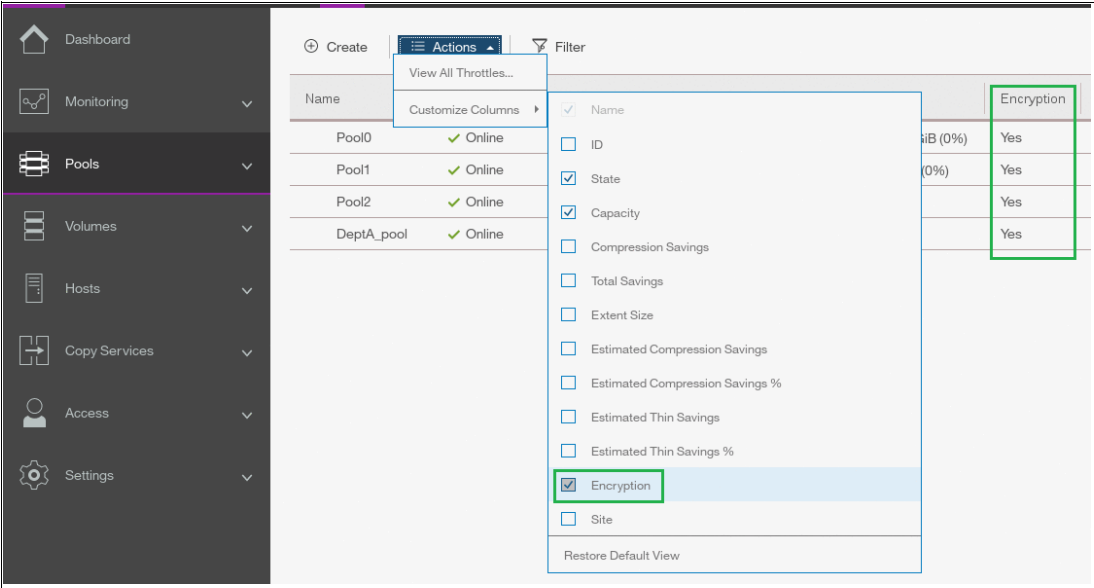


Figure 13-62 Pool encryption state

If you create an unencrypted pool, but you add only encrypted arrays or self-encrypting MDisk to the pool, then the pool will be reported as **encrypted**, because all extents in the pool are encrypted. The pool reverts back to the **unencrypted** state if you add an unencrypted array or MDisk.

Further information about how to add encrypted storage to encrypted pools is in the following sections. You can mix and match storage encryption types in a pool. Figure 13-63 on page 747 shows an example of an encrypted pool containing storage using different encryption methods.

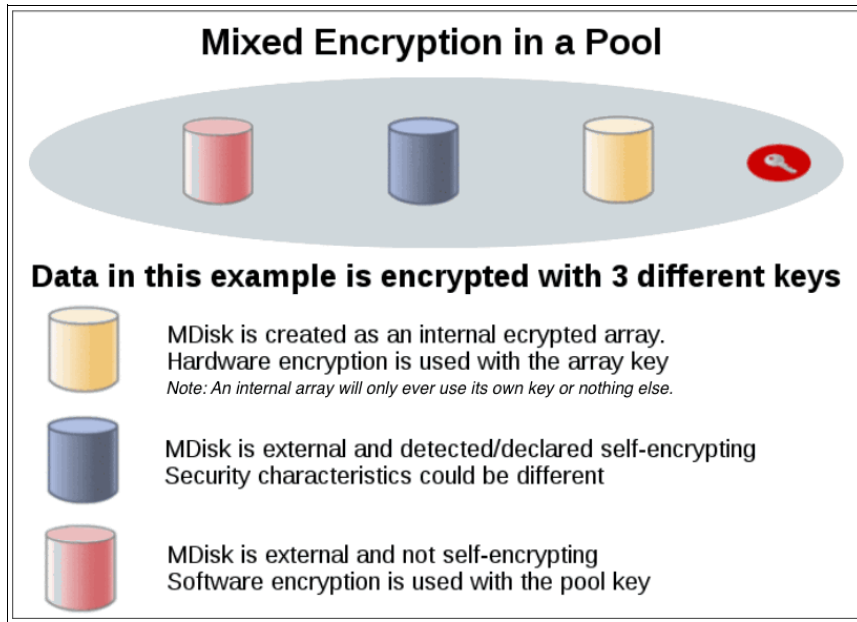


Figure 13-63 Mix and match encryption in a pool

13.8.2 Encrypted child pools

See Chapter 4, “Storage pools” on page 139 for instructions on how to open the Create Child Pool window. If the parent pool is encrypted, every child pool must be encrypted too. The GUI enforces this requirement by automatically selecting **Encryption Enabled** in the Create Child Pool window and preventing changes to this setting, as shown in Figure 13-64.

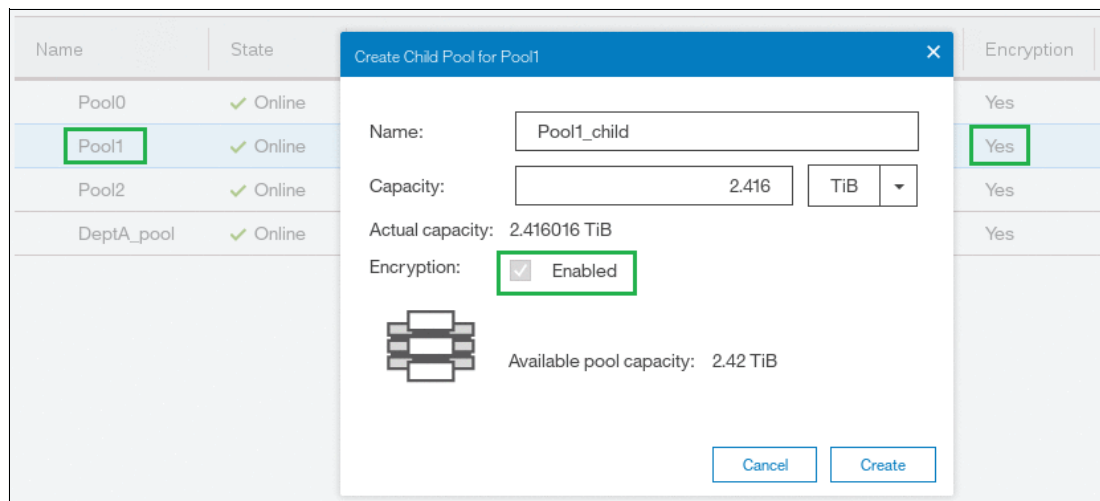


Figure 13-64 Create a child pool of an encrypted parent pool

However, if you want to create encrypted child pools from an unencrypted storage pool containing a mix of internal arrays and external MDisks, the following restrictions apply:

- The parent pool must not contain any unencrypted internal arrays
- All Lenovo Storage canisters in the system must support software encryption and have the encryption license activated

Note: An encrypted child pool created from an unencrypted parent storage pool reports as unencrypted if the parent pool contains any unencrypted internal arrays. Remove these arrays to ensure that the child pool is fully encrypted.

If you modify Pools view as described earlier in this section, you will see the encryption status of child pools, as shown in Figure 13-65. The example shows an encrypted child pool with non-encrypted parent pool.

Name	State	Capacity	Encryption
MigrationPool_1024	✓ Online	74.00 GiB / 96.00 GiB (77%)	No
MigrationPool_8192	✓ Online	10.00 GiB / 10.00 GiB (100%)	No
> Pool0	✓ Online	...	No
▼ Pool1	✓ Online	0 bytes / 64.00 GiB (0%)	No
Pool1_child	✓ Online	0 bytes / 32.00 GiB (0%)	Yes

Figure 13-65 Child pool encryption state

13.8.3 Encrypted arrays

See Chapter 4, “Storage pools” on page 139 for instructions on how to add internal storage to a pool. After encryption is enabled, all newly built arrays are hardware encrypted by default. The graphical user interface (GUI) supports only this default option.

Note: To create an unencrypted array when encryption is enabled use the command-line interface (CLI) to run the `mkarray -encrypt no` command. However, you cannot add unencrypted arrays to an encrypted pool.

You can customize MDisks by Pools view to show array encryption status. Click **Pools** → **Mdisk by Pools**, and then click on **Actions** → **Customize Columns** → **Encryption** as shown in Figure 13-66.

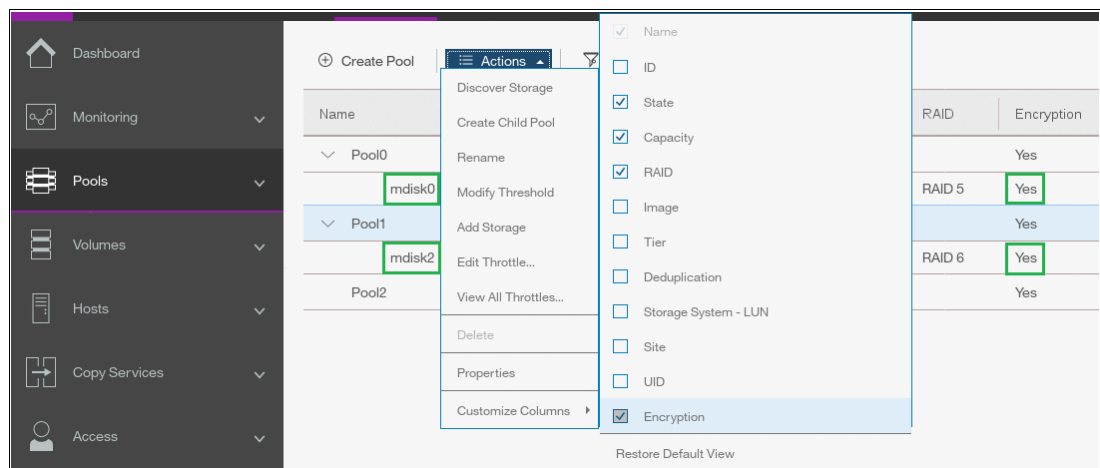


Figure 13-66 Array encryption state

You can also check the encryption state of an array by looking at its drives in **Pools** → **Internal Storage** view. The internal drives associated with an encrypted array are assigned

an encrypted property that can be seen by clicking an icon at the right edge of the table header row and selecting the **Encrypted** option from the menu, as shown in Figure 13-67.

Drive ID	Capacity	Use	Status	MDisk Name	Enclosure ID	Slot ID	↓	Encrypted	⋮
2	837.86 ...	Member	✓ Online	mdisk2	1			<input type="checkbox"/> Firmware Level	
5	837.86 ...	Member	✓ Online	mdisk2	1			<input checked="" type="checkbox"/> Enclosure ID	
22	837.86 ...	Candidate	✓ Online		1			<input checked="" type="checkbox"/> Slot ID	
6	837.86 ...	Candidate	✓ Online		1			<input type="checkbox"/> Part Number	
14	837.86 ...	Member	✓ Online	mdisk2	1			<input type="checkbox"/> Part Identity	
3	837.86 ...	Candidate	✓ Online		1			<input type="checkbox"/> Technology Type	
7	837.86 ...	Member	✓ Online	mdisk2	1			<input type="checkbox"/> RPM	
4	837.86 ...	Candidate	✓ Online		1			<input type="checkbox"/> Drive Class	
8	837.86 ...	Member	✓ Online	mdisk2	1			<input type="checkbox"/> Write Endurance Limit	
11	837.86 ...	Candidate	✓ Online		1			<input checked="" type="checkbox"/> Encrypted	
9	837.86 ...	Candidate	✓ Online		1			<input type="checkbox"/> Interface Speed	

Figure 13-67 Drive encryption state

13.8.4 Encrypted MDisks

See Chapter 4, “Storage pools” on page 139 for instructions on how to add external storage to a pool. Each MDisk belonging to external storage added to an encrypted pool or child pool is automatically encrypted using the pool or child pool key, unless the MDisk is detected or declared as self-encrypting.

The user interface gives no method to see which extents contain encrypted data and which do not. However, if a volume is created in a correctly configured encrypted pool, then all data written to this volume will be encrypted.

The extents could contain stale unencrypted data if the MDisk was earlier used for storage of unencrypted data. This is because file deletion only marks disk space as free, the data is not actually removed from the storage. So, if the MDisk is not self-encrypting and was a part of an unencrypted pool, and then was moved to an encrypted pool, then it will contain stale data from its previous life. Another failure mode is to misconfigure an external MDisk as self-encrypting, while in reality it’s not self-encrypting. At the same time, the MDisk will not encrypt the data, because it’s not self-encrypting, so we end up with unencrypted data on an extent in an encrypted pool.

However, all data written to any MDisk that’s a part of correctly configured encrypted pool, is going to be encrypted.

You can customise the MDisk by Pools view to show the object encryption state by clicking **Pools** → **MDisk by Pools**, selecting the menu bar, right-clicking it, and selecting the **Encryption Key** icon. Figure 13-68 on page 750 shows a case where self-encrypting MDisk is in an unencrypted pools.

Name	State	Capacity	Encryption
> MigrationPool_1024	✓ Online	<div><div></div></div> 74.00 GiB / 96.00 GiB (77%)	No
> MigrationPool_8192	✓ Online	<div><div></div></div> 10.00 GiB / 42.00 GiB (24%)	No
> Pool0	✓ Online	<div><div></div></div> 283.00 GiB / 553.00 GiB (51%)	No
▼ Pool1	✓ Online	<div><div></div></div> 0 bytes / 64.00 GiB (0%)	No
mdisk7	✓ Online	32.00 GiB	No
mdisk8	✓ Online	32.00 GiB	Yes

Figure 13-68 MDisk encryption state

Self-encrypting MDisks

When adding external storage to a pool, you should be exceptionally diligent when declaring the MDisk as self-encrypting. Correctly declaring an MDisk as self-encrypting avoids waste of resources, such as CPU time. However, when used improperly it may lead to unencrypted data at-rest.

To declare an MDisk as self-encrypting, select **Externally encrypted** when adding external storage in the **Assign Storage** view, as shown in Figure 13-69.

Figure 13-69 Externally encrypted MDisk

Spectrum Virtualize products can detect that an MDisk is self-encrypting by using the SCSI Inquiry page C2. MDisks provided by other Spectrum Virtualize products will report this page correctly. For these MDisks, the **Externally encrypted** box shown in Figure 13-69 will not be selected. However, when added, they are still considered as self-encrypting.

Note: You can override external encryption setting of an MDisk detected as self-encrypting and configure it as unencrypted using the CLI command `chmdisk -encrypt no`. However, you should only do so if you plan to decrypt the data on the backend or if the backend uses inadequate data encryption.

To check whether an MDisk has been detected or declared as self-encrypting, click **Pools** → **MDisk by Pools** and customize the view to show the encryption state by selecting the menu bar, right-clicking it, and selecting the **Encryption Key** icon, as shown in Figure 13-70.

Name	State	Capacity	Encryption
> MigrationPool_1024	✓ Online	<div><div></div></div> 74.00 GiB / 96.00 GiB (77%)	No
> MigrationPool_8192	✓ Online	<div><div></div></div> 10.00 GiB / 42.00 GiB (24%)	No
> Pool0	✓ Online	<div><div></div></div> 293.00 GiB / 553.00 GiB (53%)	No
▼ Pool1	✓ Online	<div><div></div></div> 0 bytes / 128.00 GiB (0%)	Yes
mdisk7	✓ Online	32.00 GiB	No
mdisk8	✓ Online	32.00 GiB	Yes
mdisk2	✓ Online	64.00 GiB	Yes

Figure 13-70 MDisk self-encryption state

Note that the value shown in the Encryption column shows the property of objects in respective rows. That means that in the configuration shown in Figure 13-70, Pool1 is encrypted, so every volume created from this pool will be encrypted. However, that pool is backed by three MDisks, out of which two are self-encrypting and one is not. Therefore, a value of “no” next to mdisk7 does not imply that encryption of Pool1 is in any way compromised. It only indicates that encryption of the data placed on mdisk7 will be done via software encryption, while data placed on mdisk2 and mdisk8 will be encrypted by the back-end storage providing these MDisks.

Note: You can change the self-encrypting attribute of an MDisk that is unmanaged or is a part of an unencrypted pool. However, you cannot change the self-encrypting attribute of an MDisk after it has been added to an encrypted pool.

13.8.5 Encrypted volumes

See Chapter 6, “Volume configuration” on page 269 for instructions on how to create and manage volumes. The encryption status of a volume depends on the pool encryption status. Volumes that are created in an encrypted pool are automatically encrypted.

You can modify Volumes view to show if the given volume is encrypted. Click **Volumes** → **Volumes** then click **Actions** → **Customize Columns** → **Encryption** to customize the view to show volumes encryption status, as shown in Figure 13-71 on page 752.

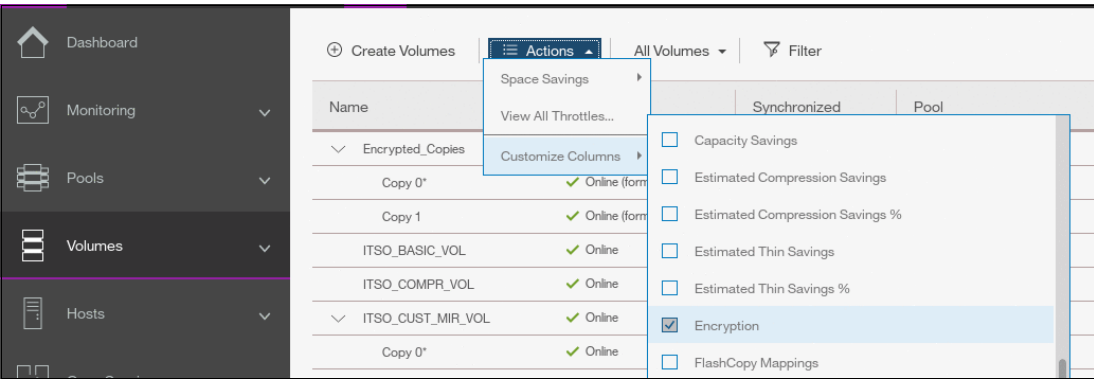


Figure 13-71 Volume view customization

Note that a volume is reported as encrypted only if all the volume copies are encrypted, as shown in Figure 13-72.

Create Volumes Actions All Volumes Copies							Showing 2 volumes Selecting 0 volumes	
Name	↑	State	Synchronized	Pool	Encryption	UID		
Encrypted_Copies	✓	Online (formatting)		Pool0	Yes	6005076300A60000008000000000000013		
Copy 0*	✓	Online (formatting)	Yes	Pool0	Yes	6005076300A60000008000000000000013		
Copy 1	✓	Online (formatting)	No	Pool1	Yes	6005076300A60000008000000000000013		
Mixed_Copies	✓	Online (formatting)		Pool1	No	6005076300A60000008000000000000012		
Copy 0*	✓	Online (formatting)	Yes	Pool1	Yes	6005076300A60000008000000000000012		
Copy 1	✓	Online (formatting)	No	Pool3	No	6005076300A60000008000000000000012		

Figure 13-72 Volume encryption status depending on volume copies encryption

When creating volumes make sure to select encrypted pools to create encrypted volumes, as shown in Figure 13-73 on page 753.

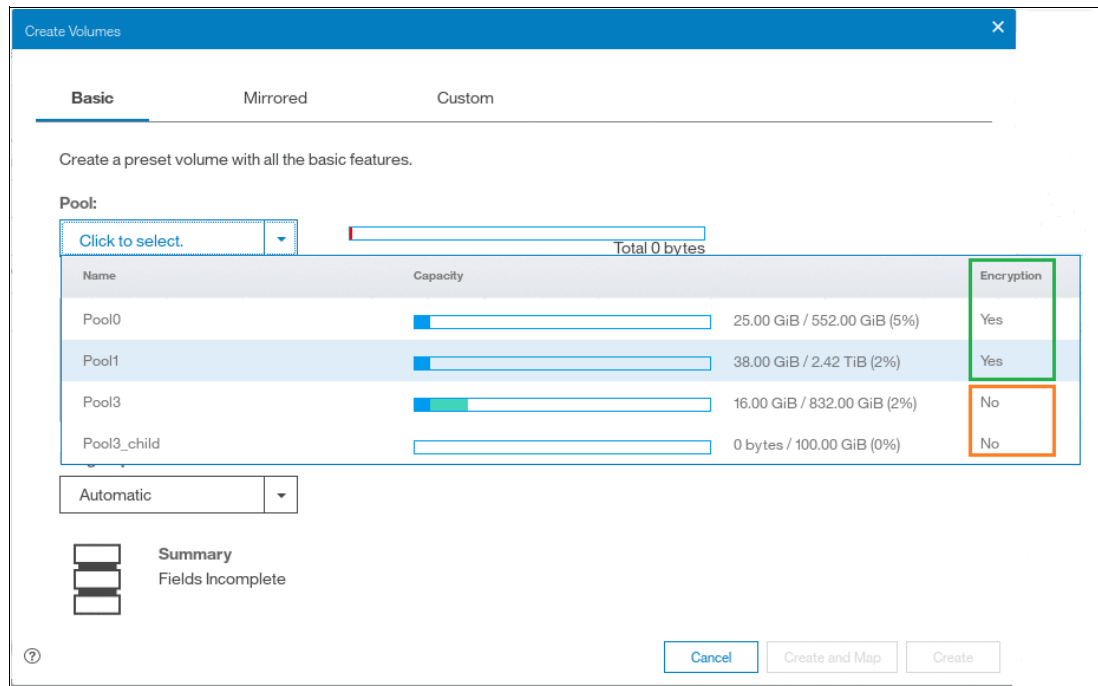


Figure 13-73 Create an encrypted volume by selecting an encrypted pool

You cannot change an existing unencrypted volume to an encrypted version of itself dynamically. However, this conversion is possible by using two migration options:

- ▶ Migrate a volume to an encrypted pool or child pool.
- ▶ Mirror a volume to an encrypted pool or child pool and delete the unencrypted copy.

For more information about either method, see Chapter 6, “Volume configuration” on page 269.

13.8.6 Restrictions

The following restrictions apply to encryption:

- ▶ Image mode volumes cannot be in encrypted pools.
- ▶ You cannot add external non self-encrypting MDisk to encrypted pools unless all nodes in the cluster support encryption.
- ▶ Nodes that cannot perform software encryption cannot be added to systems with encrypted pools that contain external MDisk that are not self-encrypting.

13.9 Rekeying an encryption-enabled system

Changing the master access key is a security requirement. *Rekeying* is the process of replacing current master access key with a newly generated one. The rekey operation works whether or not encrypted objects already exist. The rekeying operation requires access to a valid copy of the original master access key on an encryption key provider which you plan to rekey. Use the rekey operation according to the schedule defined in your organization's security policy and whenever you suspect that the key might have been compromised.

If you have both USB and key server enabled, then rekeying is done separately for each of the providers.

Important: Before you create a master access key, ensure that all nodes are online and that the current master access key is accessible.

Note: There is no method to directly change data encryption keys. If you need to change the data encryption key used to encrypt given data, then the only available method is to migrate that data to a new encrypted object (e.g. encrypted child pool). Because the data encryption keys are defined per encrypted object, such migration will force a change of the key used to encrypt that data.

13.9.1 Rekeying using a key server

Ensure that all the configured key servers can be reached by the system and that service IPs are configured on all your nodes.

To rekey the master access key kept on the key server provider, complete these steps:

1. Click **Settings** → **Security** → **Encryption**, ensure that **Encryption Keys** shows that all configured SKLM servers are reported as **Accessible**, as shown in Figure 13-74. Click on the Key Servers section label to expand the section.

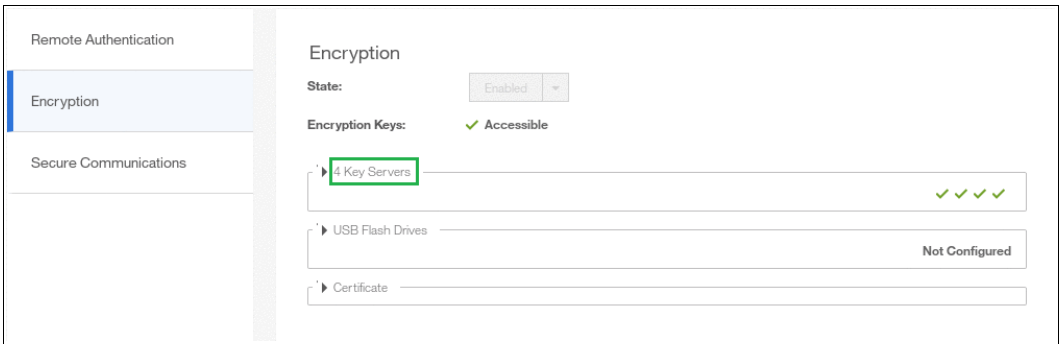


Figure 13-74 Locate Key Servers section on Encryption screen

2. Click **Rekey**, as shown in Figure 13-75 on page 755.

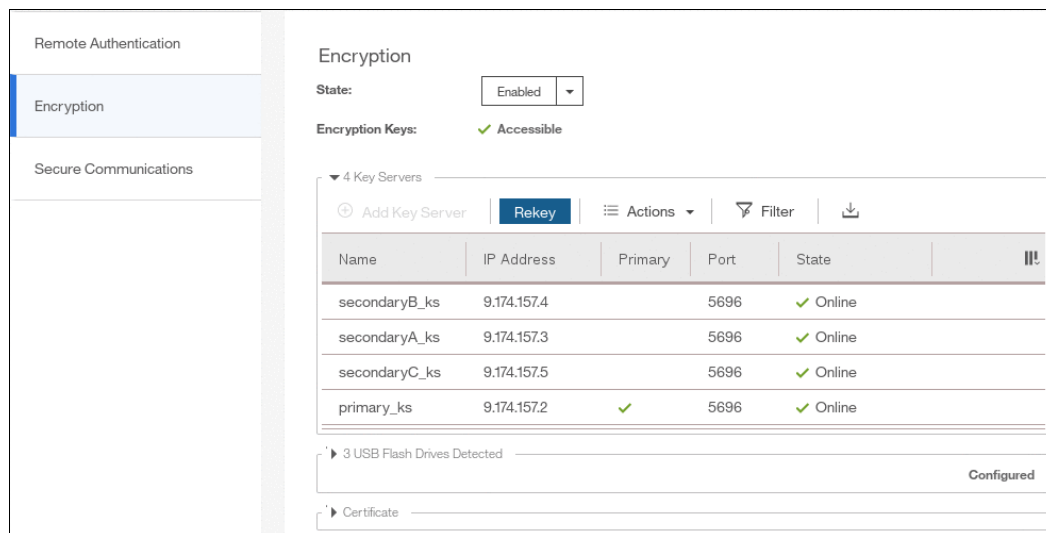


Figure 13-75 Start rekey on SKLM key server

- Click **Yes** in the next window to confirm the rekey operation, as shown in Figure 13-76.

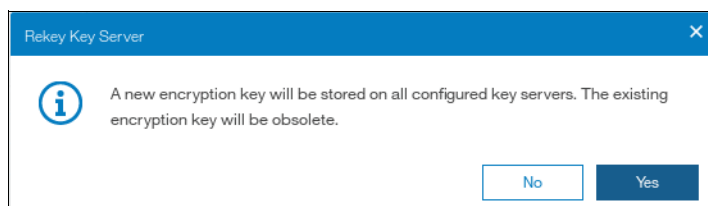


Figure 13-76 Confirm key server rekey operation

Note: The rekey operation is performed only on the primary key server configured in the system. If you have additional key servers configured apart from the primary one, they will not hold the updated encryption key until they obtain it from the primary key server. To restore encryption key provider redundancy after a rekey operation, replicate the encryption key from the primary key server to the secondary key servers.

You receive a message confirming the rekey operation was successful, as shown in Figure 13-77 on page 756.

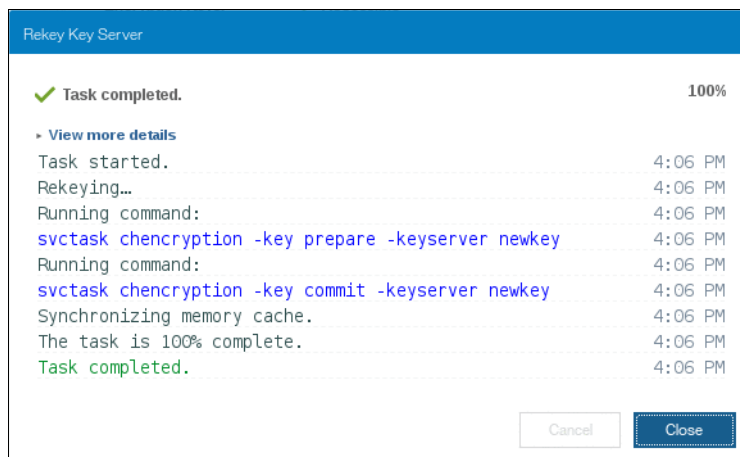


Figure 13-77 Successful key server rekey operation

13.9.2 Rekeying using USB flash drives

During the rekey process, new keys are generated and copied to the USB flash drives. These keys are then used instead of the current keys. The rekey operation fails if at least one of the USB flash drives does not contain the current key. To rekey the system, you need at least three USB flash drives to store the master access key copies.

After the rekey operation is complete, update all other copies of the encryption key, including copies stored on other media. Take the same precautions to securely store all copies of the new encryption key as when you were enabling encryption for the first time.

To rekey the master access key located on USB flash drives provider, complete these steps:

1. Click **Settings** → **Security** → **Encryption**. Click on the USB Flash Drives section label to expand the section as shown in Figure 13-78.

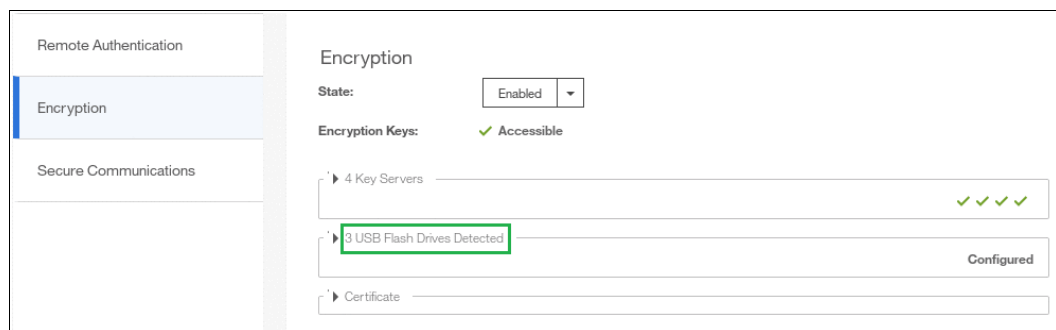


Figure 13-78 Locate USB Flash Drive section in the Encryption view

2. Verify that all USB drives plugged into the system are detected and show as Validated, as shown in Figure 13-79 on page 757. You need at least three USB flash drives, with at least one reported as Validated to process with rekey.




Figure 13-79 Start rekey on USB flash drives provider

3. If the system detects a validated USB flash drive and at least three available USB flash drives, new encryption keys are automatically copied on the USB flash drives, as shown in Figure 13-80. Click **Commit** to finalize the rekey operation.

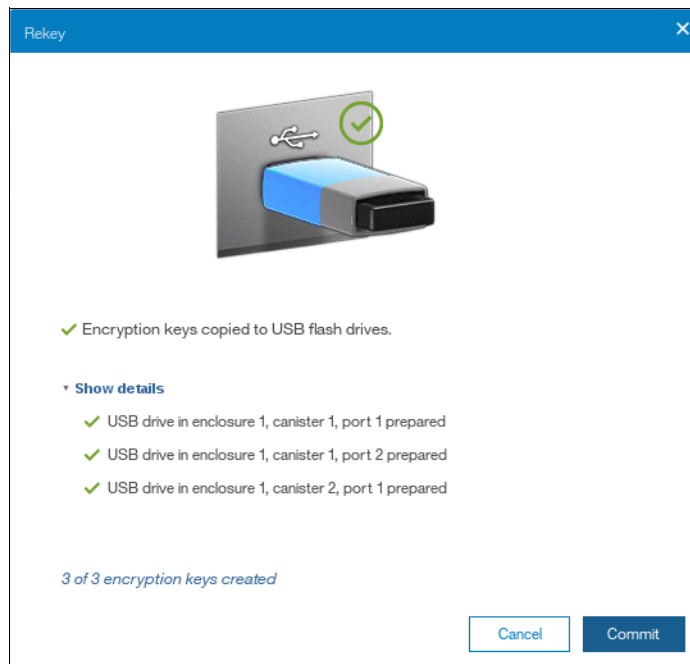


Figure 13-80 Writing new keys to USB flash drives

4. You should receive a message confirming the rekey operation was successful, as shown in Figure 13-81 on page 758.




Figure 13-81 Successful rekey operation using USB flash drives

13.10 Migrating between key providers

Once you have encryption enabled on the system it is possible to migrate from one key provider to the other.

If you only have the USB key provider enabled, and you choose to enable the key server, then the GUI gives you an option to disable the USB key provider during key server configuration. Follow the procedure as described in 13.4.3, “Enabling encryption using key servers” on page 726. During the key server provider configuration the wizard will ask if the USB flash drives provider should be disabled, as shown in Figure 13-82.

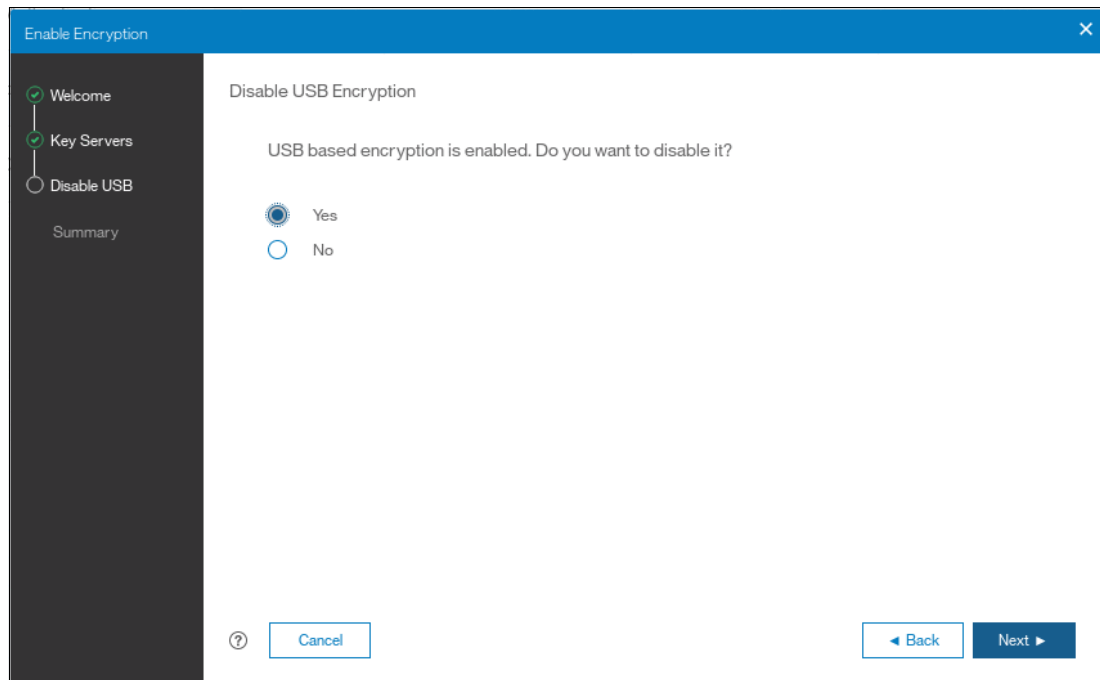


Figure 13-82 Disable the USB provider via encryption key server configuration wizard

Select **Yes** and continue with the procedure to migrate from USB to SKLM provider.

It is also possible to migrate from key server provider to USB provider or, if you have both providers enabled, to disable either of them. However, these operations are possible only via the CLI.

13.11 Disabling encryption

You are prevented from disabling encryption if there are any encrypted objects defined apart from self-encrypting MDisk. You can disable encryption in the same way whether you use USB flash drives, key server or both providers.

To disable encryption, follow these steps:

1. Click **Settings** → **Security** → **Encryption** and click **Enabled**. If no encrypted objects exist, then a drop-down menu is displayed. Click **Disabled** to disable encryption on the system. Figure 13-83 shows an example for a system with both encryption key providers configured.

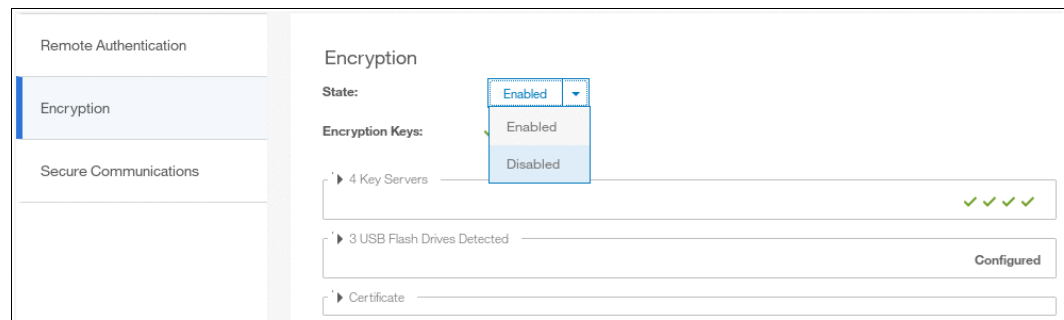


Figure 13-83 Disabling encryption on a system with both providers

2. You receive a message confirming encryption has been disabled. Figure 13-84 shows the message when using a key server.

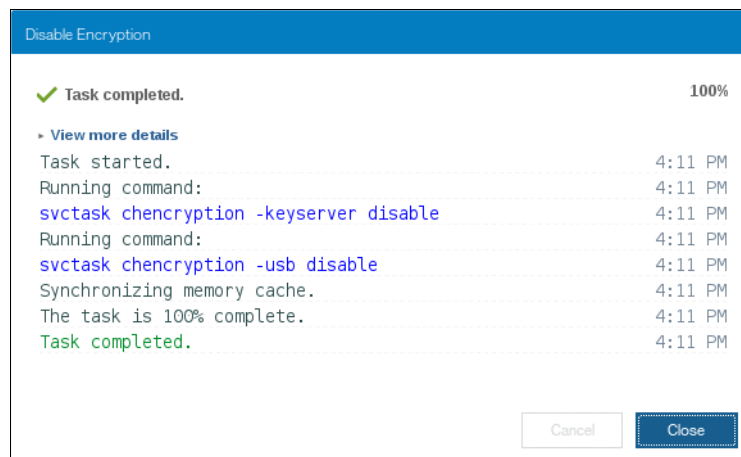


Figure 13-84 Encryption disabled

CLI setup and SAN Boot

This appendix describes the setup of the command-line interface (CLI) and provides extra information about the SAN Boot function.

Specifically, this appendix provides information about the following topics:

- ▶ “Command-line interface” on page 762
- ▶ “SAN Boot” on page 774

Command-line interface

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems have a powerful CLI, which offers even more functions than the graphical user interface (GUI). This section is not intended to be a detailed guide to the CLI because that topic is beyond the scope of this book. The basic configuration of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 CLI is covered. Example commands are described. However, the CLI commands are the same in the Lenovo Storage V series family, and more commands are available to manage internal storage. If a task completes in the GUI, the CLI command is always displayed in the details, as shown throughout this book.

Detailed CLI information is available at the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 web page under the command-line section, which is at the following address:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/tbrd_cli_strcli_4892pz.html

Basic setup

In the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 GUI, authentication is performed by using a user name and password. The CLI uses a Secure Shell (SSH) to connect from the host to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 system. Either a private and a public key pair or user name and password combination is necessary. The following steps are required to enable CLI access with SSH keys:

- ▶ A public key and a private key are generated together as a pair.
- ▶ A public key is uploaded to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems through the GUI.
- ▶ A client SSH tool must be configured to authenticate with the private key.
- ▶ A secure connection can be established between the client and the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

Secure Shell is the communication vehicle between the management workstation and the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems. The SSH client provides a secure environment from which to connect to a remote machine. It uses the principles of public and private keys for authentication. The system supports up to 32 interactive SSH sessions on the management IP address simultaneously. After 1 hour, a fixed SSH interactive session times out, which means that the SSH session is automatically closed. This session timeout limit is not configurable.

SSH keys are generated by the SSH client software. The SSH keys include a public key, which is uploaded and maintained by the clustered system, and a private key, which is kept private on the workstation that is running the SSH client. These keys authorize specific users to access the administration and service functions on the system. Each key pair is associated with a user-defined ID string that consists of up to 30 characters. Up to 100 keys can be stored on the system. New IDs and keys can be added, and unwanted IDs and keys can be deleted. To use the CLI, an SSH client must be installed on that system, the SSH key pair must be generated on the client system, and the client's SSH public key must be stored on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

The SSH client that is used in this book is PuTTY. Also, a PuTTY key generator can be used to generate the private and public key pair. The PuTTY client can be downloaded from the following address at no cost:

<http://www.chiark.greenend.org.uk>

Download the following tools:

- ▶ PuTTY SSH client: **putty.exe**
- ▶ PuTTY key generator: **puttygen.exe**

Generating a public and private key pair

To generate a public and private key pair, complete the following steps:

1. Start the PuTTY key generator to generate the public and private key pair (Figure A-1).




Figure A-1 PuTTY key generator

Ensure that the following options are used:

- SSH-2 RSA
- Number of bits in a generated key: 1024

2. Click **Generate** and move the cursor over the blank area to generate keys (Figure A-2).




Figure A-2 Generate keys

To generate keys: The blank area that is indicated by the message is the large blank rectangle on the GUI inside the section of the GUI that is labeled Key. Continue to move the mouse pointer over the blank area until the progress bar reaches the far right. This action generates random characters to create a unique key pair.

3. After the keys are generated, save them for later use. Click **Save public key** (Figure A-3).

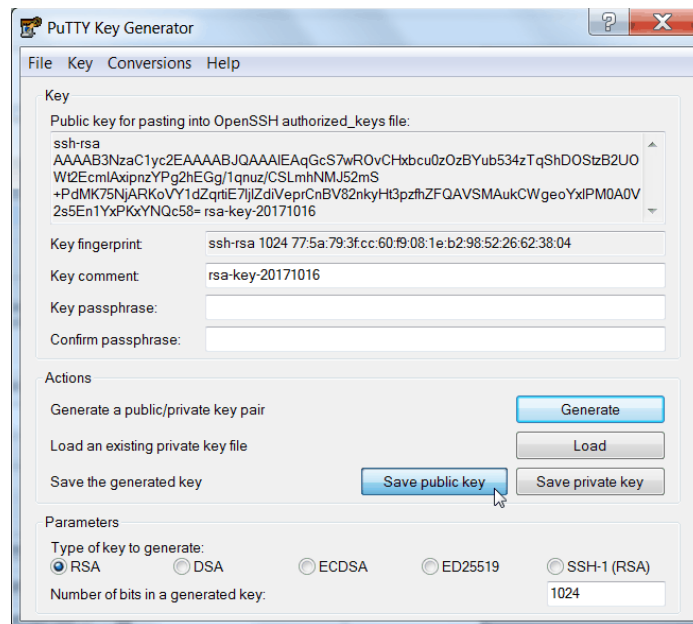


Figure A-3 Save public key

4. You are prompted for a name (for example, pubkey) and a location for the public key (for example, C:\Support Utils\PuTTY). Click **Save**.

Ensure that you record the name and location because the name and location of this SSH public key must be specified later.

Public key extension: By default, the PuTTY key generator saves the public key with no extension. Use the string “pub” for naming the public key, for example, superuser.pub, to easily differentiate the SSH public key from the SSH private key.

5. Click **Save private key** (Figure A-4).

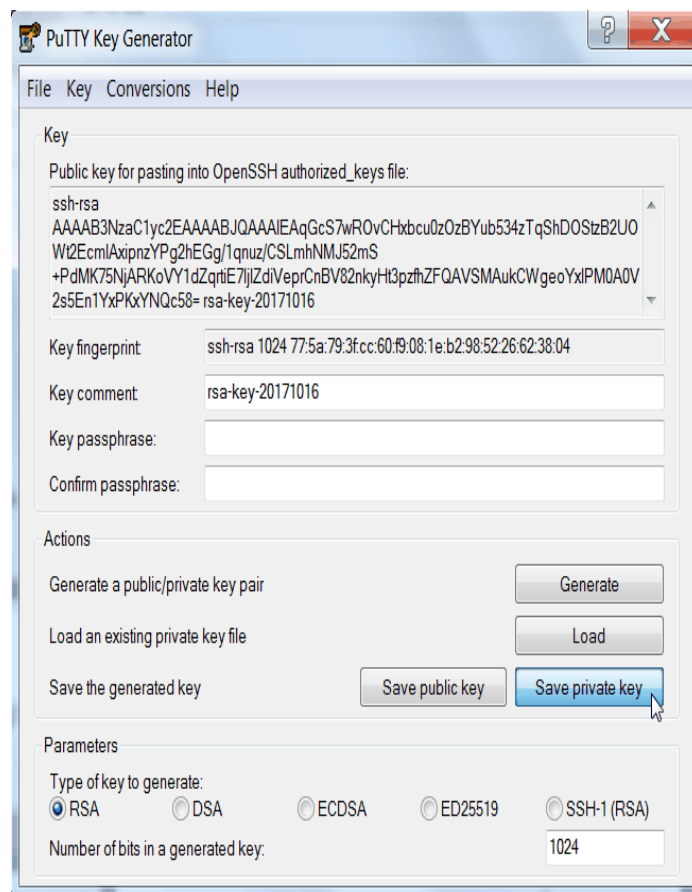


Figure A-4 Save private key

6. You are prompted with a warning message (Figure A-5). Click **Yes** to save the private key without a passphrase.

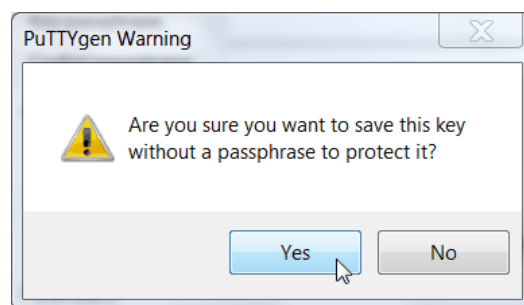


Figure A-5 Confirm the security warning

7. When you are prompted, enter a name (for example, i cat), select a secure place as the location, and click **Save**.

Key generator: The PuTTY key generator saves the private key with the PPK extension.

8. Close the PuTTY key generator.

Uploading the SSH public key to the storage

After you create your SSH key pair, upload your SSH public key onto the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems. Complete the following steps:

1. On the System Overview, click the **Access** functional icon and select **Users** in the GUI menu (Figure A-6).

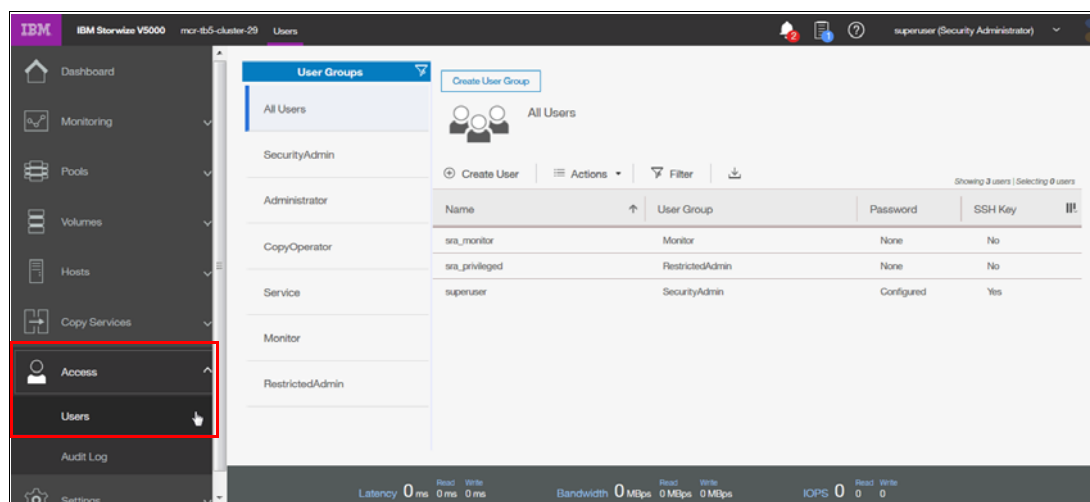


Figure A-6 Click Users on the Access menu

2. Under User Groups, select **All Users**. Right-click the user name for which you want to upload the key and click **Properties** (Figure A-7).

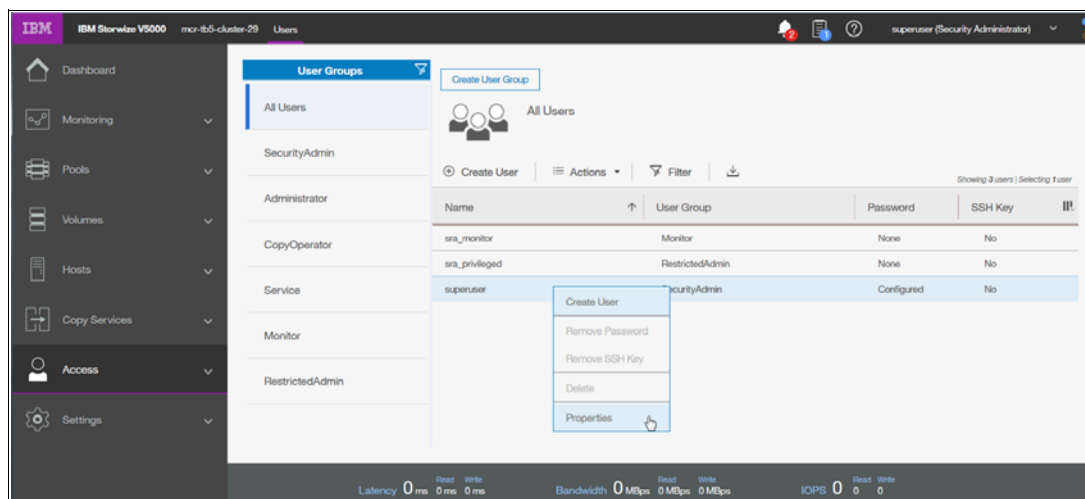


Figure A-7 Superuser properties

3. To upload the public key, click **Browse**, and select the folder where you stored the public SSH key (Figure A-8).

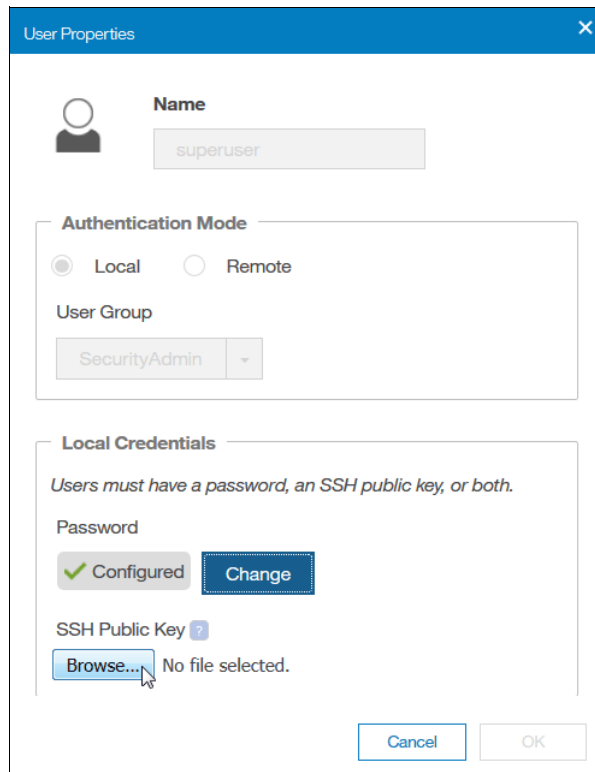


Figure A-8 Select public key

4. Select your public key, and click **Open** (Figure A-9).




Figure A-9 Selection of the public SSH key

5. Click **OK**, as shown in Figure A-10 on page 769. The key is uploaded.

User Properties

Name
superuser

Authentication Mode
☒ Local ☐ Remote
 User Group: SecurityAdmin

Local Credentials
 Users must have a password, an SSH public key, or both.
 Password: ✓ Configured Change
 SSH Public Key: Browse... Superuser pubkey

Cancel OK

Figure A-10 Select the public key

6. Check in the GUI whether the SSH key was successfully imported. See Figure A-11.

IBM Storage V5000 mcr-ft5-cluster-29 Users

User Groups
 All Users
 SecurityAdmin
 Administrator
 CopyOperator
 Service
 Monitor
 RestrictedAdmin

Create User Group

All Users

Showing 3 users | Selecting 1 user

Name	User Group	Password	SSH Key
sra_monitor	Monitor	None	No
sra_privileged	RestrictedAdmin	None	No
superuser	SecurityAdmin	Configured	Yes

Latency 0 ms Read 0 ms Write 0 ms Bandwidth 0 MBps Read 0 MBps Write 0 MBps IOPS 0 Read 0 Write 0

Figure A-11 SSH key was successfully imported

Configuring the SSH client

Before you can use the CLI, you must configure the SSH client:

1. Start PuTTY. The PuTTY Configuration window opens (Figure A-12).

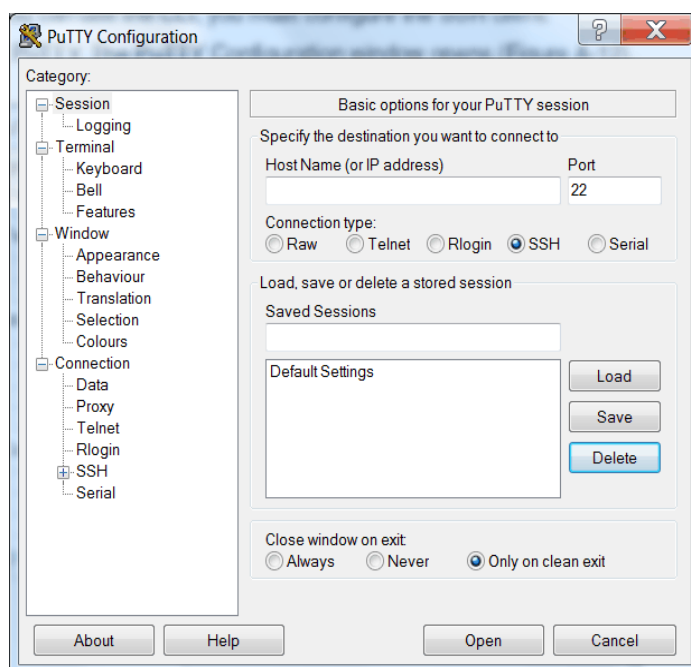


Figure A-12 PuTTY

In the right pane, select **SSH** as the connection type. Under the “Close window on exit” section, select **Only on clean exit**, which ensures that if any connection errors occur, they are displayed on the user’s window, see Figure A-13.

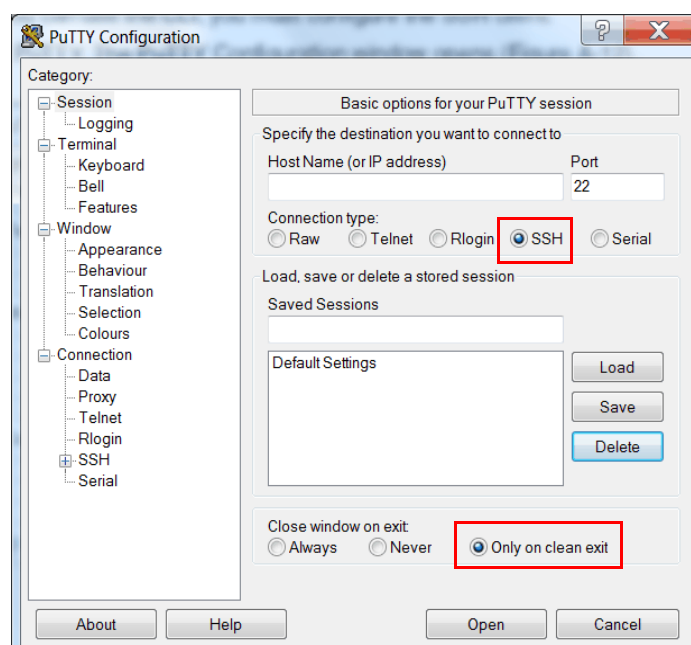


Figure A-13 Select SSH + Only on clean exit

2. In the Category pane, on the left side of the PuTTY Configuration window (Figure A-14), click **Connection** → **SSH** to open the PuTTY Configuration window Options controlling SSH connections view.

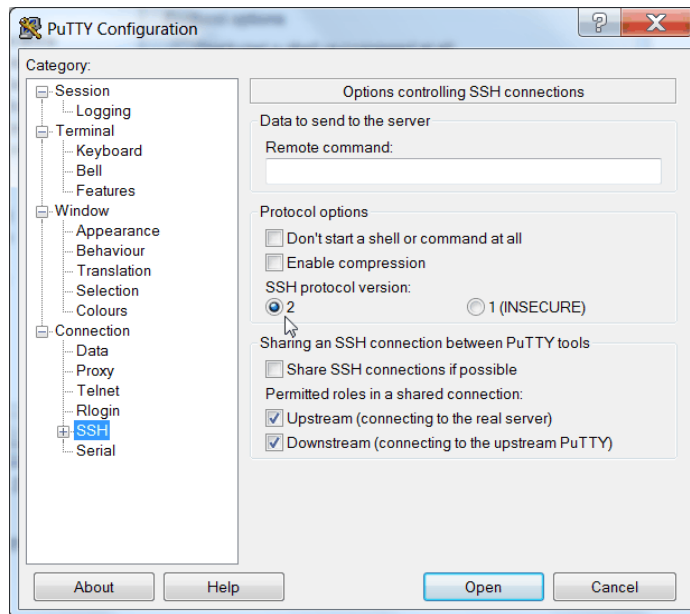


Figure A-14 SSH protocol version 2

Under Preferred SSH protocol version, select **2**.

3. In the Category pane on the left, click **Connection** → **SSH** → **Auth**, as shown in Figure A-15. More options are displayed for controlling SSH authentication.

In the Private key file for authentication field, either browse to or type the fully qualified directory path and file name of the SSH client private key file, which was created previously (for example, C:\Support Utils\putty\privat.PPK).

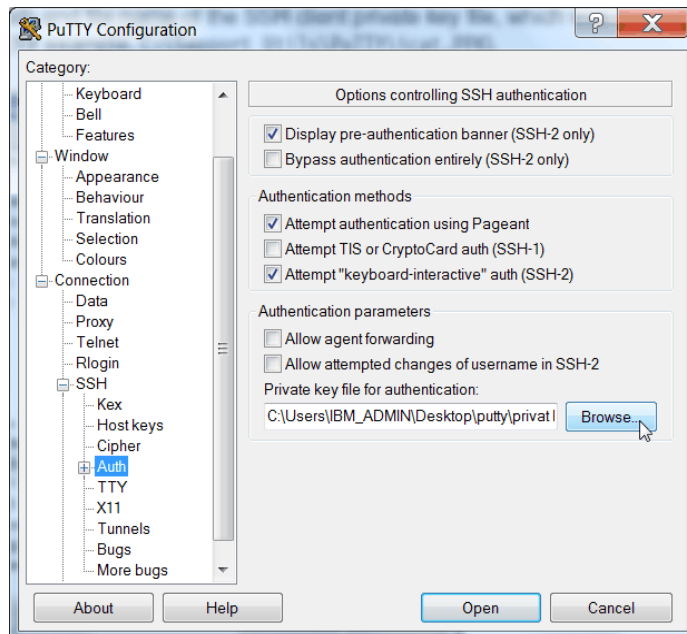


Figure A-15 SSH authentication

4. In the Category pane, click **Session** to return to the PuTTY Configuration window Basic options for your PuTTY session view (Figure A-16).
5. Enter the following information in these fields in the right pane:
 - Host Name (or IP address): Specify the host name or system IP address of the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 clustered system.
 - Saved Sessions: Enter a session name.

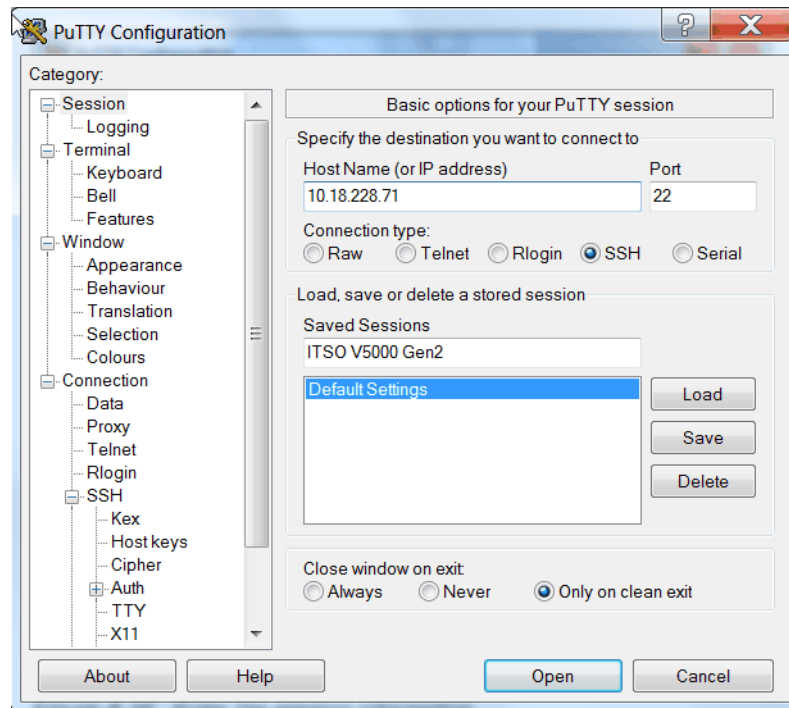


Figure A-16 Enter the session information

6. Click **Save** to save the new session (Figure A-17).

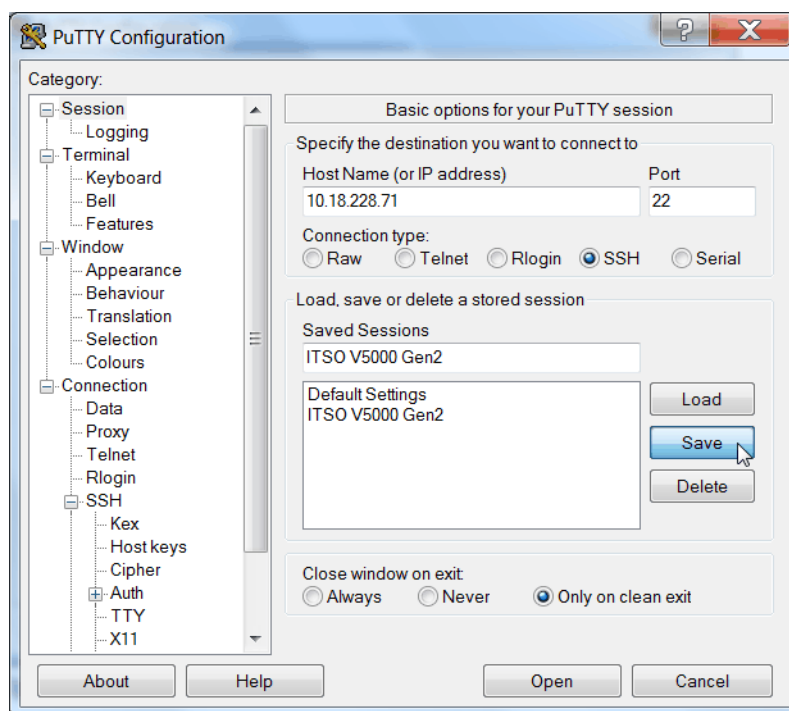


Figure A-17 Save the new session

7. Figure A-18 shows the saved PUTTY session. Select the new session and click **Open**.

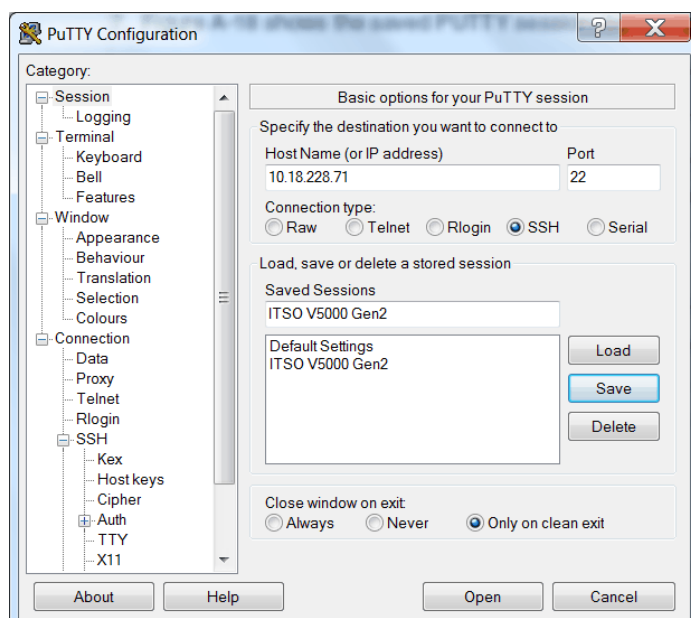


Figure A-18 Saved PUTTY session

8. If a PuTTY Security Alert window opens. Confirm it by clicking **Yes** (Figure A-19 on page 774).

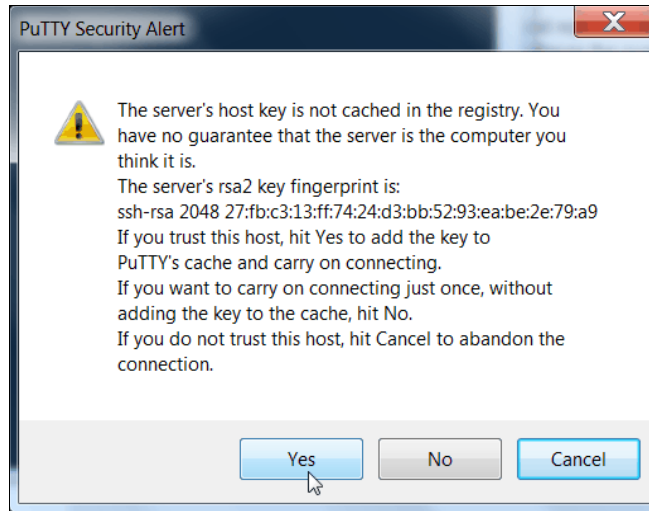


Figure A-19 Confirm the security alert

9. PuTTY now connects to the system and prompts you for a user name to log in as. Enter `Superuser` as the user name (Example A-1) and click Enter.

Example A-1 Enter user name

```
login as: Superuser
Authenticating with public key "putty public key"
IBM_2077:ITS0 V5000Gen2:Superuser>
```

The tasks to configure the CLI for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 administration are complete.

SAN Boot

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support SAN Boot for Microsoft Windows, VMware, and many other operating systems. SAN Boot support can change, so regularly check the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 interoperability matrix at this address:

<https://datacentersupport.lenovo.com/us/en/products/storage/lenovo-storage/v3700v2/6535/documentation>

<https://datacentersupport.lenovo.com/us/en/products/storage/lenovo-storage/v5030/6536/documentation>

More information about SAN Boot is also available in the *Multipath Subsystem Device Driver User's Guide*, which is available at the following address:

<http://www.ibm.com/support/docview.wss?rs=503&context=HW26L&uid=ssg1S7000303>

Enabling SAN Boot for Windows

Complete the following procedure if you want to install a Windows host by using SAN Boot:

1. Configure the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems so that only the boot volume is mapped to the host.
2. Configure the Fibre Channel storage area network (SAN) so that the host sees only one Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems node port. Multiple paths during installation are not supported.
3. Configure and enable the host bus adapter (HBA) BIOS.
4. Install the operating system by using the normal procedure, selecting the volume as the partition on which to install.

HBAs: You might need to load an additional HBA device driver during installation, depending on your Windows version and the HBA type.

5. Install Subsystem Device Driver Device Specific Module (SDDDSM) after the installation completes.
6. Modify your SAN zoning to allow multiple paths.
7. Check your host to see whether all paths are available.
8. Set redundant boot devices in the HBA BIOS to enable the host to boot when its original path fails.

Enabling SAN Boot for VMware

Complete the following steps if you want to install a VMware ESX host by using SAN Boot:

1. Configure the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems so that only the boot volume is mapped to the host.
2. Configure the Fibre Channel SAN so that the host sees only one Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 system node port. Multiple paths during installation are not supported.
3. Configure and enable the HBA BIOS.
4. Install the operating system by using the normal procedure, selecting the volume as the partition on which to install.

HBAs: You might need to load an additional HBA device driver during installation, depending on your ESX level and the HBA type.

5. Modify your SAN zoning to allow multiple paths.
6. Check your host to see whether all paths are available and modify the multipath policy, if required.

Windows SAN Boot migration

If your host runs Windows Server 2008, Windows 2012 or Windows 2016 operating system and uses existing SAN Boot images that are controlled by storage controllers, you can migrate these images to image-mode volumes that are controlled by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems.

SAN Boot procedures: For SAN Boot procedures for other operating systems, check the Lenovo Information Center for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v3700.doc/lenovo_vseries.html

Complete the following steps to migrate your existing SAN Boot images:

1. If the existing SAN Boot images are controlled by an Lenovo storage controller that uses the Lenovo Subsystem Device Driver (SDD) as the multipathing driver, you must use SDD V1.6 or later. Run the SDD **datapath set bootdiskmigrate 2077** command to prepare the host for image migration. For more information, see the Multipath Subsystem Device Driver documentation.
2. Shut down the host.
3. Complete the following configuration changes on the storage controller:
 - a. Write down the Small Computer System Interface (SCSI) logical unit number (LUN) ID that each volume is using, for example, boot LUN SCSI ID 0, Swap LUN SCSI ID 1, and Database LUN SCSI ID 2.
 - b. Remove all of the image-to-host mappings from the storage controller.
 - c. Map the existing SAN Boot image and any other disks to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems.
4. Change the zoning so that the host can see the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 I/O group for the target image mode volume.
5. Complete the following configuration changes on the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems:
 - a. Create an image mode volume for the managed disk (MDisk) that contains the SAN Boot image. Use the MDisk unique identifier to specify the correct MDisk.
 - b. Create a host object and assign the host HBA ports.
 - c. Map the image mode volume to the host by using the same SCSI ID as before. For example, you might map the boot disk to the host with SCSI LUN ID 0.
 - d. Map the swap disk to the host, if required. For example, you might map the swap disk to the host with SCSI LUN ID 1.
6. Change the boot address of the host by completing the following steps:
 - a. Restart the host and open the HBA BIOS utility of the host during the booting process.
 - b. Set the BIOS settings on the host to find the boot image at the worldwide port name (WWPN) of the node that is zoned to the HBA port.
7. If SDD V1.6 or later is installed and you run **bootdiskmigrate** in step 1, reboot your host, update SDDDSM to the current level, and go to step 14 on page 777. If SDD V1.6 is not installed, go to step 8.
8. Modify the SAN zoning so that the host sees only one path to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.
9. Boot the host in single-path mode.
10. Uninstall any multipathing driver that is not supported for the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 systems hosts that run the applicable Windows Server operating system.
11. Install SDDDSM.

12. Restart the host in single-path mode. Ensure that SDDDSM was correctly installed.
13. Modify the SAN zoning to enable multipathing.
14. Rescan the drives on your host and check that all paths are available.
15. Reboot your host and enter the HBA BIOS.
16. Configure the HBA settings on the host. Ensure that all HBA ports are boot-enabled and that they can see both nodes in the I/O group that contains the SAN Boot image.
Configure the HBA ports for redundant paths.
17. Exit the BIOS utility and finish starting the host.

Map any additional volumes to the host, as required.

Terminology

This appendix summarizes the controller firmware and Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 terms that are commonly used in this book.

To see the complete set of terms that relate to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, see Lenovo Information Center at:

http://systemx.lenovofiles.com/help/topic/com.lenovo.storage.v5030.8.1.0.doc/lenovo_vseries.html

Array

An ordered collection, or group, of physical devices (disk drive modules) that are used to define logical volumes or devices. An array is a group of drives designated to be managed with a Redundant Array of Independent Disks (RAID).

Asymmetric virtualization

Asymmetric virtualization is a virtualization technique in which the virtualization engine is outside the data path and performs a metadata-style service. The metadata server contains all the mapping and locking tables, and the storage devices contain only data. See also “Symmetric virtualization” on page 793

Asynchronous replication

Asynchronous replication is a type of replication in which control is given back to the application as soon as the write operation is made to the source volume. Later, the write operation is made to the target volume. See also “Synchronous replication” on page 793.

Automatic data placement mode

Automatic data placement mode is an Easy Tier operating mode in which the host activity on all the volume extents in a pool are “measured,” a migration plan is created, and then automatic extent migration is performed.

Back end

See “Front end and back end” on page 785.

Caching I/O Group

The caching I/O Group is the I/O Group in the system that performs the cache function for a volume.

Call home

Call home is a communication link that is established between a product and a service provider. The product can use this link to call IBM or another service provider when the product requires service. With access to the machine, service personnel can perform service tasks, such as viewing error and problem logs or initiating trace and dump retrievals.

Canister

A canister is a single processing unit within a storage system.

Capacity licensing

Capacity licensing is a licensing model that licenses features with a price-per-terabyte model. Licensed features are FlashCopy, Metro Mirror, Global Mirror, and virtualization. See also “FlashCopy” on page 784, “Metro Mirror” on page 788, and “Virtualization” on page 794.

Chain

A set of enclosures that are attached to provide redundant access to the drives inside the enclosures. Each control enclosure can have one or more chains.

Challenge Handshake Authentication Protocol

Challenge Handshake Authentication Protocol (CHAP) is an authentication protocol that protects against eavesdropping by encrypting the user name and password.

Channel extender

A channel extender is a device that is used for long-distance communication that connects other storage area network (SAN) fabric components. Generally, channel extenders can involve protocol conversion to asynchronous transfer mode (ATM), Internet Protocol (IP), or another long-distance communication protocol.

Child pool

Administrators can use child pools to control capacity allocation for volumes that are used for specific purposes. Rather than being created directly from managed disks (MDisks), child pools are created from existing capacity that is allocated to a parent pool. As with parent pools, volumes can be created that specifically use the capacity that is allocated to the child pool. Child pools are similar to parent pools with similar properties. Child pools can be used for volume copy operation. Also, see “Parent pool” on page 789.

Cloud Container

Cloud Container is a virtual object that includes all of the elements, components or data that are common to a specific application or data.

Cloud Service Provider

Cloud Service Provider (CSP) is the company or organization that provides off- and on-premises cloud services such as storage, server, network, and so on. Controller firmware has built in software capabilities to interact with Cloud Providers such as IBM SoftLayer, Amazon S3 and deployments of OpenStack Swift.

Cloud Tenant

Cloud Tenant is a group or an instance that provides common access with the specific privileges to a object, software or data source.

Clustered system (Lenovo Storage V3700 V2, V3700 V2 XP, and V5030)

A clustered system, formerly known as a cluster, is a group of up to four Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 canisters (two in each system) that presents a single configuration, management, and service interface to the user.

Cold extent

A cold extent is an extent of a volume that does not get any performance benefit if it is moved from a hard disk drive (HDD) to a Flash disk. A cold extent also refers to an extent that needs to be migrated onto an HDD if it is on a Flash disk drive.

Compression

Compression is a function that removes repetitive characters, spaces, strings of characters, or binary data from the data that is being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for data. See also “RACE engine” on page 790.

Compression accelerator

A compression accelerator is hardware onto which the work of compression is off-loaded from the microprocessor.

Configuration node

While the cluster is operational, a single node in the cluster is appointed to provide configuration and service functions over the network interface. This node is termed the configuration node. This configuration node manages the data that describes the clustered-system configuration and provides a focal point for configuration commands. If the configuration node fails, another node in the cluster transparently assumes that role.

Consistency Group

A Consistency Group is a group of copy relationships between virtual volumes or data sets that are maintained with the same time reference so that all copies are consistent in time. A Consistency Group can be managed as a single entity.

Container

A container is a software object that holds or organizes other software objects or entities.

Contingency capacity

For thin-provisioned volumes that are configured to automatically expand, the unused real capacity that is maintained. For thin-provisioned volumes that are not configured to automatically expand, the difference between the used capacity and the new real capacity.

Copied state

Copied is a FlashCopy state that indicates that a copy was triggered after the copy relationship was created. The Copied state indicates that the copy process is complete and the target disk has no further dependency on the source disk. The time of the last trigger event is normally displayed with this status.

Counterpart SAN

A counterpart SAN is a non-redundant portion of a redundant SAN. A counterpart SAN provides all of the connectivity of the redundant SAN, but without 100% redundancy. Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 canisters are typically connected to a “redundant SAN” that is made up of two counterpart SANs. A counterpart SAN is often called a SAN fabric.

Cross-volume consistency

A consistency group property that ensures consistency between volumes when an application issues dependent write operations that span multiple volumes.

Data consistency

Data consistency is a characteristic of the data at the target site where the dependent write order is maintained to ensure the recoverability of applications.

Data encryption key

The data encryption key is used to encrypt data and it is created automatically when an encrypted object, such as an array, a pool, or a child pool, is created. It is stored in secure memory and it cannot be viewed or changed. The data encryption key is encrypted using the master access key.

Data migration

Data migration is the movement of data from one physical location to another physical location without the disruption of application I/O operations.

Data reduction pool

Data Reduction pools are specific types of pools where more control over volumes capacity is given to specific hosts (for example VMware VAAI/VASA/VVOL, Microsoft ODX). These hosts are able to return unused space for reuse. With standard pools, the system is not aware of any unused space on host-allocated volumes.

Dependent write operation

A write operation that must be applied in the correct order to maintain cross-volume consistency.

Directed Maintenance Procedure

The fix procedures, which are also known as Directed Maintenance Procedures (DMPs), ensure that you fix any outstanding errors in the error log. To fix errors, from the Monitoring panel, click **Events**. The Next Recommended Action is displayed at the top of the Events window. Select **Run This Fix Procedure** and follow the instructions.

Discovery

The automatic detection of a network topology change, for example, new and deleted nodes or links.

Disk tier

MDisks (logical unit numbers (LUNs)) that are presented to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 likely have different performance attributes because of the type of disk or RAID array on which they are installed. MDisks can be on 15,000 revolutions per minute (RPM) Fibre Channel (FC) or serial-attached SCSI (SAS) disk, Nearline SAS, or

Serial Advanced Technology Attachment (SATA), or even Flash Disks. Therefore, a storage tier attribute is assigned to each MDisk, and the default is *generic_hdd*.

Distributed RAID or DRAID

An alternative RAID scheme where the number of drives that are used to store the array can be greater than the equivalent, typical RAID scheme. The same data stripes are distributed across a greater number of drives, which increases the opportunity for parallel I/O and hence improves overall array performance. See also “Rebuild area” on page 791.

Easy Tier

Easy Tier is a volume performance function within the Lenovo Storage V series family that provides automatic data placement of a volume’s extents in a multitiered storage pool. The pool normally contains a mix of Flash Disks and HDDs. Easy Tier measures host I/O activity on the volume’s extents and migrates hot extents onto the Flash Disks to ensure the maximum performance.

Encryption key

The encryption key, also known as master access key, is created and stored on USB flash drives or on a key server when encryption is enabled. The master access key is used to decrypt the data encryption key.

Encryption key server

An internal or external system that receives and then serves existing encryption keys or certificates to a storage system.

Encryption of data at rest

Encryption of data at rest is the inactive encryption data that is stored physically on the storage system.

Evaluation mode

The evaluation mode is an Easy Tier operating mode in which the host activity on all the volume extents in a pool are “measured” only. No automatic extent migration is performed.

Event (error)

An event is an occurrence of significance to a task or system. Events can include the completion or failure of an operation, user action, or a change in the state of a process.

Event code

An event code is a value that is used to identify an event condition to a user. This value might map to one or more event IDs or to values that are presented on the service panel. This value is used to report error conditions to Lenovo and to provide an entry point into the service guide.

Event ID

An event ID is a value that is used to identify a unique error condition that was detected by the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. An event ID is used internally in the cluster to identify the error.

Excluded condition

The excluded condition is a status condition. It describes an MDisk that the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 have decided is no longer sufficiently reliable to be

managed by the cluster. The user must issue a command to include the MDisk in the cluster-managed storage.

Extent

An extent is a fixed-size unit of data that is used to manage the mapping of data between MDisk and volumes. The size of the extent can range 16 MB - 8 GB in size.

External storage

External storage refers to managed disks (MDisks) that are SCSI logical units that are presented by storage systems that are attached to and managed by the clustered system.

Failback

Failback is the restoration of an appliance to its initial configuration after the detection and repair of a failed network or component.

Failover

Failover is an automatic operation that switches to a redundant or standby system or node in a software, hardware, or network interruption. See also Failback.

Feature activation code

An alphanumeric code that activates a licensed function on a product.

Fibre Channel port logins

Fibre Channel (FC) port logins refer to the number of hosts that can see any one Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 port. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 have a maximum limit per node port of FC logins that are allowed.

Field-replaceable unit

Field-replaceable units (FRUs) are individual parts that are replaced entirely when any one of the unit's components fails. They are held as spares by the Lenovo service organization.

FlashCopy

FlashCopy refers to a point-in-time copy where a virtual copy of a volume is created. The target volume maintains the contents of the volume at the point in time when the copy was established. Any subsequent write operations to the source volume are not reflected on the target volume.

FlashCopy mapping

A FlashCopy mapping is a continuous space on a direct-access storage volume, which is occupied by or reserved for a particular data set, data space, or file.

FlashCopy relationship

See FlashCopy mapping.

FlashCopy service

FlashCopy service is a copy service that duplicates the contents of a source volume on a target volume. In the process, the original contents of the target volume are lost. See also "Point-in-time copy" on page 789.

Flash drive

A data storage device that uses solid-state memory to store persistent data.

Flash module

A modular hardware unit containing flash memory, one or more flash controllers, and associated electronics.

Front end and back end

The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 take MDisks to create pools of capacity from which volumes are created and presented to application servers (hosts). The MDisks are in the controllers at the back end of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 to the back-end controller zones. The volumes that are presented to the hosts are in the front end of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

Global Mirror

Global Mirror (GM) is a method of asynchronous replication that maintains data consistency across multiple volumes within or across multiple systems. Global Mirror is generally used where distances between the source site and target site cause increased latency beyond what the application can accept.

Global Mirror with change volumes

Change volumes are used to record changes to the primary and secondary volumes of a remote copy relationship. A FlashCopy mapping exists between a primary and its change volume and a secondary and its change volume.

Grain

A grain is the unit of data that is represented by a single bit in a FlashCopy bitmap (64 KiB or 256 KiB) in the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. A grain is also the unit to extend the real size of a thin-provisioned volume (32 KiB, 64 KiB, 128 KiB, or 256 KiB).

Hop

One segment of a transmission path between adjacent nodes in a routed network.

Host bus adapter

A host bus adapter (HBA) is an interface card that connects a server to the SAN environment through its internal bus system, for example, PCI Express. Typically it is referred to the Fibre Channel adapters.

Host ID

A host ID is a numeric identifier that is assigned to a group of host FC ports or Internet Small Computer System Interface (iSCSI) host names for LUN mapping. For each host ID, SCSI IDs are mapped to volumes separately. The intent is to have a one-to-one relationship between hosts and host IDs, although this relationship cannot be policed.

Host mapping

Host mapping refers to the process of controlling which hosts have access to specific volumes within a cluster (host mapping is equivalent to LUN masking).

Hot extent

A hot extent is a frequently accessed volume extent that gets a performance benefit if it is moved from an HDD onto a Flash Disk.

Hot Spare node (SVC only)

Hot Spare Node is an online SVC node defined in a cluster but not in any IO group. In case of a failure of any of online nodes in any IO group of cluster, it will be automatically swapped by this Spare node. Once the recovery of an original node has finished, the Spare node gets back to the standby spare status. This feature is not available for Lenovo Storage V3700 V2, V3700 V2 XP, and V5030.

HyperSwap

Pertaining to a function that provides continuous, transparent availability against storage errors and site failures, and is based on synchronous replication.

Image mode

Image mode is an access mode that establishes a one-to-one mapping of extents in the storage pool (existing LUN or (image mode) MDisk) with the extents in the volume.

Image volume

An image volume is a volume in which a direct block-for-block translation exists from the managed disk (MDisk) to the volume.

I/O Group

Each pair of Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 canisters is known as an input/output (I/O) Group. An I/O Group has a set of volumes that are associated with it that are presented to host systems. Each Lenovo Storage V3700 V2, V3700 V2 XP, or V5030 canister is associated with exactly one I/O Group. The canister in an I/O Group provide a failover and failback function for each other. A Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 cluster consists of two I/O groups.

Internal storage

Internal storage refers to an array of managed disks (MDisks) and drives that are held in Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 enclosures.

Internet Small Computer System Interface qualified name

Internet Small Computer System Interface (iSCSI) qualified name (IQN) refers to special names that identify both iSCSI initiators and targets. IQN is one of the three name formats that is provided by iSCSI. The IQN format is `iqn.<yyyy-mm>.<reversed domain name>`. For example, the default for a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 canisters can be in the following format:

```
iqn.1986-03.com.ibm:2076.<clustername>.<nodename>
```

Internet storage name service

The Internet storage name service (iSNS) protocol that is used by a host system to manage iSCSI targets and the automated iSCSI discovery, management, and configuration of iSCSI and FC devices. It was defined in Request for Comments (RFC) 4171.

Inter-switch link hop

An inter-switch link (ISL) is a connection between two switches and counted as one ISL hop. The number of hops is always counted on the shortest route between two N-ports (device

connections). In a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 environment, the number of ISL hops is counted on the shortest route between the pair of canister that are farthest apart. The Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 support a maximum of three ISL hops.

Input/output group

A collection of volumes and canister relationships that present a common interface to host systems. Each pair of canister is known as an input/output (I/O) group.

iSCSI initiator

An initiator functions as an iSCSI client. An initiator typically serves the same purpose to a computer as a SCSI bus adapter would, except that, instead of physically cabling SCSI devices (like hard drives and tape changers), an iSCSI initiator sends SCSI commands over an IP network.

iSCSI session

An iSCSI Initiator and an iSCSI Target talk with each other and this conversation called an iSCSI Session.

iSCSI target

An iSCSI target is a storage resource located on an Internet Small Computer System Interface (iSCSI) server.

Latency

The time interval between the initiation of a send operation by a source task and the completion of the matching receive operation by the target task. More generally, latency is the time between a task initiating data transfer and the time that transfer is recognized as complete at the data destination.

Least recently used

Least recently used (LRU) pertains to an algorithm used to identify and make available the cache space that contains the data that was least recently used.

Licensed capacity

The amount of capacity on a storage system that a user is entitled to configure.

License key

An alphanumeric code that activates a licensed function on a product.

License key file

A file that contains one or more licensed keys.

Lightweight Directory Access Protocol

Lightweight Directory Access Protocol (LDAP) is an open protocol that uses TCP/IP to provide access to directories that support an X.500 model and that does not incur the resource requirements of the more complex X.500 Directory Access Protocol (DAP). For example, LDAP can be used to locate people, organizations, and other resources in an Internet or intranet directory.

Local and remote fabric interconnect

The local fabric interconnect and the remote fabric interconnect are the SAN components that are used to connect the local and remote fabrics. Depending on the distance between the two fabrics, they can be single-mode optical fibers that are driven by long wave (LW) gigabit interface converters (GBICs) or small form-factor pluggables (SFPs), or more sophisticated components, such as channel extenders or special SFP modules that are used to extend the distance between SAN components.

Local fabric

The local fabric is composed of SAN components (switches, cables, and so on) that connect the components (nodes, hosts, and switches) of the local cluster together.

Logical unit and logical unit number

The logical unit (LU) is defined by the SCSI standards as a logical unit number (LUN). LUN is an abbreviation for an entity that exhibits disk-like behavior, for example, a volume or an MDisk.

Machine signature

A string of characters that identifies a system. A machine signature might be required to obtain a license key.

Managed disk

A managed disk (MDisk) is a SCSI disk that is presented by a RAID controller and managed by Lenovo Storage V3700 V2, V3700 V2 XP, and V5030. The MDisk is not visible to host systems on the SAN.

Managed disk group (storage pool)

See “Storage pool (managed disk group)” on page 793.

Metro Global Mirror

Metro Mirror Global is a cascaded solution where Metro Mirror synchronously copies data to the target site. This Metro Mirror target is the source volume for Global Mirror that asynchronously copies data to a third site. This solution has the potential to provide disaster recovery with no data loss at Global Mirror distances when the intermediate site does not participate in the disaster that occurs at the production site.

Metro Mirror

Metro Mirror (MM) is a method of synchronous replication that maintains data consistency across multiple volumes within the system. Metro Mirror is generally used when the write latency that is caused by the distance between the source site and target site is acceptable to application performance.

Mirrored volume

A mirrored volume is a single virtual volume that has two physical volume copies. The primary physical copy is known within the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 as copy 0 and the secondary copy is known within the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 as copy 1.

Node canister

A node canister is a hardware unit that includes the node hardware, fabric and service interfaces, and serial-attached SCSI (SAS) expansion ports. Node canisters are specifically

recognized on Lenovo Storage V series products. In SVC all these components are spread within the whole system chassis, so we usually do not consider node canisters in SVC, but just the node as a whole.

Node rescue

The process by which a node that has no valid software installed on its hard disk drive can copy software from another node connected to the same Fibre Channel fabric.

NPIV

NPIV or N_Port ID Virtualization is a Fibre Channel feature whereby multiple Fibre Channel node port (N_Port) IDs can share a single physical N_Port.

Object Storage

Object storage is a general term that refers to the entity in which an Cloud Object Storage (COS) organize, manage and store with units of storage or just *objects*.

Oversubscription

Oversubscription refers to the ratio of the sum of the traffic on the initiator N-port connections to the traffic on the most heavily loaded ISLs, where more than one connection is used between these switches. Oversubscription assumes a symmetrical network, and a specific workload that is applied equally from all initiators and sent equally to all targets. A symmetrical network means that all the initiators are connected at the same level, and all the controllers are connected at the same level.

Parent pool

Parent pools receive their capacity from MDisks. All MDisks in a pool are split into extents of the same size. Volumes are created from the extents that are available in the pool. You can add MDisks to a pool at any time either to increase the number of extents that are available for new volume copies or to expand existing volume copies. The system automatically balances volume extents between the MDisks to provide the best performance to the volumes. See also “Child pool” on page 780.

Partnership

In Metro Mirror or Global Mirror operations, the relationship between two clustered systems. In a clustered-system partnership, one system is defined as the local system and the other system as the remote system.

Point-in-time copy

A point-in-time copy is the instantaneous copy that the FlashCopy service makes of the source volume. See also “FlashCopy service” on page 784.

Preparing phase

Before you start the FlashCopy process, you must prepare a FlashCopy mapping. The preparing phase flushes a volume's data from cache in preparation for the FlashCopy operation.

Primary volume

In a stand-alone Metro Mirror or Global Mirror relationship, the target of write operations issued by the host application.

Private fabric

Configure one SAN per fabric so that it is dedicated for node-to-node communication. This SAN is referred to as a private SAN.

Public fabric

Configure one SAN per fabric so that it is dedicated for host attachment, storage system attachment, and remote copy operations. This SAN is referred to as a public SAN. You can configure the public SAN to allow Lenovo Storage V series family node-to-node communication also. You can optionally use the `-localportfcmask` parameter of the `chsystem` command to constrain the node-to-node communication to use only the private SAN.

Quorum disk

A disk that contains a reserved area that is used exclusively for system management. The quorum disk is accessed when it is necessary to determine which half of the clustered system continues to read and write data. Quorum disks can either be MDisk or drives.

Quorum index

The quorum index is the pointer that indicates the order that is used to resolve a tie. Nodes attempt to lock the first quorum disk (index 0), followed by the next disk (index 1), and finally the last disk (index 2). The tie is broken by the node that locks them first.

RACE engine

The RACE engine compresses data on volumes in real time with minimal effect on performance. See “Compression” on page 781 or “Real-time Compression”.

Real capacity

Real capacity is the amount of storage that is allocated to a volume copy from a storage pool.

Real-time Compression

Real-time Compression is an IBM integrated software function for storage space efficiency. The RACE engine compresses data on volumes in real time with minimal effect on performance. See also “RACE engine”.

Redundant Array of Independent Disks

Redundant Array of Independent Disks (RAID) refers to two or more physical disk drives that are combined in an array in a certain way, which incorporates a RAID level for failure protection or better performance. The most common RAID levels are 0, 1, 5, 6, and 10. Some storage administrators refer to the RAID group as TRAIID - Traditional RAID.

RAID 0

RAID 0 is a data striping technique that is used across an array and no data protection is provided.

RAID 1

RAID 1 is a mirroring technique that is used on a storage array in which two or more identical copies of data are maintained on separate mirrored disks.

RAID 10

RAID 10 is a combination of a RAID 0 stripe that is mirrored (RAID 1). Therefore, two identical copies of striped data exist; no parity exists.

RAID 5

RAID 5 is an array that has a data stripe, which includes a single logical parity drive. The parity check data is distributed across all the disks of the array.

RAID 6

RAID 6 is a RAID level that has two logical parity drives per stripe, which are calculated with different algorithms. Therefore, this level can continue to process read and write requests to all of the array's virtual disks in the presence of two concurrent disk failures.

Read intensive drives

The Read Intensive (RI) flash drives (SSD drives) that are available on Lenovo Storage V3700 V2, V3700 V2 XP, and V5030, and Lenovo Storage V7000 are one Drive Write Per Day (DWPD) Read Intensive drives.

Rebuild area

Reserved capacity that is distributed across all drives in a redundant array of drives. If a drive in the array fails, the lost array data is systematically restored into the reserved capacity, returning redundancy to the array. The duration of the restoration process is minimized because all drive members simultaneously participate in restoring the data. See also "Distributed RAID or DRAID" on page 783.

Redundant storage area network

A redundant SAN is a SAN configuration in which there is no single point of failure (SPoF); therefore, data traffic continues no matter what component fails. Connectivity between the devices within the SAN is maintained (although possibly with degraded performance) when an error occurs. A redundant SAN design is normally achieved by splitting the SAN into two independent counterpart SANs (two SAN fabrics), so that if one path of the counterpart SAN is destroyed, the other counterpart SAN path keeps functioning.

Relationship

In Metro Mirror or Global Mirror, a relationship is the association between a master volume and an auxiliary volume. These volumes also have the attributes of a primary or secondary volume.

Reliability, availability, and serviceability

Reliability, availability, and serviceability (RAS) are a combination of design methodologies, system policies, and intrinsic capabilities that, when taken together, balance improved hardware availability with the costs that are required to achieve it.

Reliability is the degree to which the hardware remains free of faults. Availability is the ability of the system to continue operating despite predicted or experienced faults. Serviceability is how efficiently and nondisruptively broken hardware can be fixed.

Remote fabric

The remote fabric is composed of SAN components (switches, cables, and so on) that connect the components (nodes, hosts, and switches) of the remote cluster together. Significant distances can exist between the components in the local cluster and those components in the remote cluster.

Remote Support Server and Client

Remote Support Client is a software toolkit that resides in Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 and opens a secured tunnel to the Remote Support Server. Remote

Support Server resides in the Lenovo network and collects key health check and troubleshooting informations required by Lenovo support personnel.

Secondary volume

Pertinent to remote copy, the volume in a relationship that contains a copy of data written by the host application to the primary volume.

Secure Sockets Layer certificate

Secure Sockets Layer (SSL) is the standard security technology for establishing an encrypted link between a web server and a browser. This link ensures that all data passed between the web server and browsers remain private and to be able to create an SSL connection a web server requires an SSL Certificate.

Security Key Lifecycle Manager

Security Key Lifecycle Manager (SKLM) centralizes, simplifies, and automates the encryption key management process to help minimize risk and reduce operational costs of encryption key management.

Serial-attached SCSI

Serial-attached Small Computer System Interface (SAS) is a method that is used in accessing computer peripheral devices that employs a serial (one bit at a time) means of digital data transfer over thin cables. The method is specified in the American National Standard Institute standard called SAS. In the business enterprise, SAS is useful for access to mass storage devices, particularly external hard disk drives.

Service Location Protocol

The Service Location Protocol (SLP) is an Internet service discovery protocol that enables computers and other devices to find services in a local area network (LAN) without prior configuration. It was defined in the request for change (RFC) 2608.

Small Computer System Interface (SCSI)

Small Computer System Interface (SCSI) is an ANSI-standard electronic interface with which personal computers can communicate with peripheral hardware, such as disk drives, tape drives, CD-ROM drives, printers, and scanners, faster and more flexibly than with previous interfaces.

Snapshot

A snapshot is an image backup type that consists of a point-in-time view of a volume.

Solid-state disk

A solid-state disk (SSD) or Flash Disk is a disk that is made from solid-state memory and therefore has no moving parts. Most SSDs use NAND-based flash memory technology. It is defined to the Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 as a disk tier `generic_ssd`.

Space efficient

See “Thin provisioning” on page 794.

Spare

An extra storage component, such as a drive or tape, that is predesignated for use as a replacement for a failed component.

Spare goal

The optimal number of spares that are needed to protect the drives in the array from failures. The system logs a warning event when the number of spares that protect the array drops below this number.

Space-efficient volume

For more information about a space-efficient volume, see “Thin-provisioned volume” on page 794.

Stand-alone relationship

In FlashCopy, Metro Mirror, and Global Mirror, relationships that do not belong to a consistency group and that have a null consistency-group attribute.

Statesave

Binary data collection that is used for a problem determination by Lenovo service support.

Storage area network or SAN

A storage area network (SAN) is a dedicated storage network that is tailored to a specific environment, which combines servers, systems, storage products, networking products, software, and services.

Storage pool (managed disk group)

A storage pool is a collection of storage capacity, which is made up of managed disks (MDisks), that provides the pool of storage capacity for a specific set of volumes. A storage pool can contain more than one tier of disk, which is known as a multitier storage pool and a prerequisite of Easy Tier automatic data placement.

Striped

Pertaining to a volume that is created from multiple managed disks (MDisks) that are in the storage pool. Extents are allocated on the MDisks in the order specified.

Support Assistant

A function that is used to provide support personnel access to the system to complete troubleshooting and maintenance tasks.

Symmetric virtualization

Symmetric virtualization is a virtualization technique in which the physical storage, in the form of a Redundant Array of Independent Disks (RAID), is split into smaller chunks of storage known as extents. These extents are then concatenated, by using various policies, to make volumes. See also “Asymmetric virtualization” on page 779.

Synchronous replication

Synchronous replication is a type of replication in which the application write operation is made to both the source volume and target volume before control is given back to the application. See also “Asynchronous replication” on page 779.

Thin-provisioned volume

A thin-provisioned volume is a volume that allocates storage when data is written to it.

Thin provisioning

Thin provisioning refers to the ability to define storage, usually a storage pool or volume, with a “logical” capacity size that is larger than the actual physical capacity that is assigned to that pool or volume. Therefore, a thin-provisioned volume is a volume with a virtual capacity that differs from its real capacity.

Throttles

Throttling is a mechanism to control the amount of resources that are used when the system is processing I/Os on supported objects. The system supports throttles on hosts, host clusters, volumes, copy offload operations, and storage pools. If a throttle limit is defined, the system either processes the I/O for that object, or delays the processing of the I/O to free resources for more critical I/O operations.

Transparent Cloud Tiering

Transparent Cloud Tiering is a separately installable feature of IBM Spectrum Scale that provides a native cloud storage tier.

T10 DIF

T10 DIF is a *Data Integrity Field* (DIF) extension to SCSI to enable end-to-end protection of data from host application to physical media.

Unique identifier

A unique identifier (UID) is an identifier that is assigned to storage-system logical units when they are created. It is used to identify the logical unit regardless of the logical unit number (LUN), the status of the logical unit, or whether alternate paths exist to the same device. Typically, a UID is used only once.

Virtualization

In the storage industry, virtualization is a concept in which a pool of storage is created that contains several storage systems. Storage systems from various vendors can be used. The pool can be split into volumes that are visible to the host systems that use them. See also “Capacity licensing” on page 780.

Virtualized storage

Virtualized storage is physical storage that has virtualization techniques applied to it by a virtualization engine.

Virtual local area network

Virtual local area network (VLAN) tagging separates network traffic at the layer 2 level for Ethernet transport. The system supports VLAN configuration on both IPv4 and IPv6 connections.

Virtual Storage Area Network

A virtual Storage Area Network (VSAN) is a logical fabric entity defined within the storage area network (SAN). It can be defined on a single physical SAN switch or across multiple physical switched or directors. In VMware terminology the vSAN is defined as a logical layer of storage capacity built from physical disk drives attached directly into the ESXi hosts. This solution is not considered for the scope of our publication.

Vital product data

Vital product data (VPD or VDP) is information that uniquely defines system, hardware, software, and microcode elements of a processing system.

Volume

A volume is a Lenovo Storage V3700 V2, V3700 V2 XP, and V5030 logical device that appears to host systems that are attached to the SAN as a SCSI disk. Each volume is associated with exactly one I/O Group. A volume has a preferred node within the I/O Group.

Volume copy

A volume copy is a physical copy of the data that is stored on a volume. Mirrored volumes have two copies. Non-mirrored volumes have one copy.

Volume protection

To prevent active volumes or host mappings from inadvertent deletion, the system supports a global setting that prevents these objects from being deleted if the system detects that they have recent I/O activity. When you delete a volume, the system checks to verify whether it is part of a host mapping, FlashCopy mapping, or remote-copy relationship. In these cases, the system fails to delete the volume, unless the **-force** parameter is specified. Using the **-force** parameter can lead to unintentional deletions of volumes that are still active. Active means that the system detected recent I/O activity to the volume from any host.

Write-through mode

Write-through mode is a process in which data is written to a storage device at the same time that the data is cached.

Notices

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

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

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO 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. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

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

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

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

Note: This document is based on an IBM Redbooks publication. The content was used with permission.

Trademarks

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

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo(logo)®

Lenovo®

The following terms are trademarks of other companies:

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

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

Hyper-V, Internet Explorer, Microsoft, Microsoft Edge, SQL Server, Windows, Windows Server, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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



Implementing the Lenovo Storage V3700 V2 and V5030 Systems with IBM Spectrum Virtualize V8.1

Provides an overview of Lenovo Storage V3700 V2, V2 XP and V5030 systems

Introduces configuration setup for the Lenovo Storage V3700 V2, V2 XP and V5030 systems

Explains storage terminologies and concepts for various applications

Describes trouble shooting and monitoring of the Lenovo Storage V3700 V2, V2 XP and V5030

Organizations of all sizes face the challenge of managing massive volumes of increasingly valuable data. But storing this data can be costly, and extracting value from the data is becoming more difficult. IT organizations have limited resources but must stay responsive to dynamic environments and act quickly to consolidate, simplify, and optimize their IT infrastructures. The Lenovo Storage V3700 V2, V2 XP and V5030 systems provide a smarter solution that is affordable, easy to use, and self-optimizing, which enables organizations to overcome these storage challenges.

These storage systems deliver efficient, entry-level configurations that are designed to meet the needs of small and midsize businesses. Designed to provide organizations with the ability to consolidate and share data at an affordable price, the Lenovo Storage V3700 V2, V2 XP and V5030 offer advanced software capabilities that are found in more expensive systems.

This book is intended for pre-sales and post-sales technical support professionals and storage administrators. It applies to the Lenovo Storage V3700 V2, V3700 V2 XP and V5030 with IBM Spectrum Virtualize V8.1.



**BUILDING
TECHNICAL
INFORMATION
BASED ON
PRACTICAL
EXPERIENCE**

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.