



Cisco UCS Director Management Guide for IBM Storwize, Release 6.0

First Published: 2016-09-16

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2016 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface

Preface ix

Audience ix

Conventions ix

Related Documentation xi

Documentation Feedback xi

Obtaining Documentation and Submitting a Service Request xi

CHAPTER 1

New and Changed Information for this Release 1

New and Changed Information 1

CHAPTER 2

Overview 3

VersaStack 3

Discovering Devices for IBM Storwize Accounts 4

CHAPTER 3

Managing IBM Storwize and IBM SAN Volume Controller Storage Accounts 5

IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage Systems 5

IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Features 5

Managing IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage Accounts 6

Adding a Pod 6

Adding an IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller Storage Account 8

IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage System Reports 9

Managing System Tasks for IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage Accounts 10

CHAPTER 4**Managing Internal Drives 11**

- Internal Storage 11
- Internal Storage RAID Configurations 12

CHAPTER 5**Managing MDisks 13**

- MDisks 13
- Encryption Enabled Array MDisks 14
- Creating an Array MDisk 14
- Renaming an MDisk 15
- Deleting an Array MDisk 16
- Assigning an MDisk to a Storage Pool 16
- Unassigning an MDisk from the Storage Pool 17
- MDisk Tier 17
- Setting Tier 18
- Swapping Internal Drives 18
- Setting Spare Drives 19

CHAPTER 6**Managing Storage Pools 21**

- Parent Pools 21
- Child Pools 21
- Creating a Storage Pool 22
- Renaming a Storage Pool 22
- Deleting a Storage Pool 23
- Expanding a Storage Pool 23

CHAPTER 7**Managing Volumes 25**

- Volumes 25
 - File Volumes 26
 - Block Volumes 26
- Creating Volumes 27
- Adding Mirrored Copy 28
- Deleting Volumes 29
- Editing Volumes 29
- Expanding Volumes 30

Migrating Volumes	30
Renaming Volumes	31
Shrinking Volumes	31
Mapping to Host	32
Unmapping All Hosts	32
Unmapping A Single Host	33
View Details of Volume Mirrored Copies	33
Viewing Volume Mirrored Copies	33
Making Primary Volume	34
Splitting Volumes into New Volumes	34
Duplicating Volumes	35
Deleting Mirrored Copy	35

CHAPTER 8**Managing Hosts 37**

Hosts	37
Creating Hosts	38
Deleting Hosts	39
Renaming Hosts	40
Unmapping All Volumes	40
Importing Mappings	41
Duplicating Mappings	41
Deploying an ESXi Fibre Channel and Fibre Channel over Ethernet Host on IBM Storwize System	42
Deploying an ESXi iSCSI Host on IBM Storwize System	43
Creating and Adding Block Volume Datastore on the ESXi Host	45
Adding Block Datastore over ESXi on IBM Storwize System	46
Adding Compressed Datastore over ESXi on IBM Storwize System	46

CHAPTER 9**Managing File Storage 47**

File System Pools	47
File Systems	47
Creating Single Pool File Systems	48
Creating Compressed File Systems	48
Creating Migration-ILM File Systems	50
Mounting File Systems	52

- Unmounting File Systems 52
- Deleting File Systems 52
- File Sets 53
 - Creating File Sets 53
 - Editing File Sets 54
 - Deleting File Sets 55
- File Shares 55
 - Creating File Shares 55
 - Deleting File Shares 58
 - Activating and Deactivating File Shares 58
- Quotas 58
 - Creating Quotas 59
 - Editing Quotas 60
 - Deleting Quotas 61
- Snapshots 61
 - Snapshot Rules 61
 - Creating Snapshot Rules 61
 - Creating Snapshots 62
 - Deleting Snapshots 63
 - Configuring Snapshot Notifications 64

CHAPTER 10

- Managing Copy Services 65**
 - Copy Services 65
 - Creating FlashCopy Snapshot 66
 - Creating FlashCopy Backup 66
 - Creating FlashCopy Clone 67
 - Deleting FlashCopy 67
 - Starting FlashCopy Services 68
 - Stopping FlashCopy Services 68
 - Creating Remote Copy Relationships 68
 - Editing Remote Copy Relationships 69
 - Renaming Remote Copy Relationships 70
 - Deleting Remote Copy Relationships 70
 - Starting Remote Copy Relationships 71
 - Stopping Remote Copy Relationships 71

- Switching Remote Copy Relationships 71
- Adding Remote Copy Relationships to Remote Copy Consistent Groups 72
- Removing Remote Copy Relationships from Remote Copy Consistent Groups 72
- File Copy Services 72
 - Creating a Replication Destination 73
 - Removing a Replication Destination 74
- Consistency Groups 74
 - Creating FlashCopy and Remote Copy Consistency Groups 74
 - Renaming FlashCopy and Remote Copy Consistency Groups 75
 - Deleting FlashCopy and Remote Copy Consistency Groups 75

CHAPTER 11

- FlashCopy Mappings 77**
 - FlashCopy Mappings 77
 - FlashCopy Mapping States 78
 - Creating FlashCopy Mappings 78
 - FlashCopy Advanced Settings 79
 - Moving FlashCopy Mapping to FlashCopy Consistency Group 81
 - Starting FlashCopy Mapping 81
 - Stopping FlashCopy Mapping 82
 - Deleting FlashCopy Mapping 82

CHAPTER 12

- Workflows for IBM Storwize System 83**
 - Workflows for IBM Storwize and IBM SVC Storage Accounts 83



Preface

- [Audience, page ix](#)
- [Conventions, page ix](#)
- [Related Documentation, page xi](#)
- [Documentation Feedback, page xi](#)
- [Obtaining Documentation and Submitting a Service Request, page xi](#)

Audience

This guide is intended primarily for data center administrators who use Cisco UCS Director and who have responsibilities and expertise in one or more of the following:

- Server administration
- Storage administration
- Network administration
- Network security
- Virtualization and virtual machines

Conventions

Text Type	Indication
GUI elements	GUI elements such as tab titles, area names, and field labels appear in this font . Main titles such as window, dialog box, and wizard titles appear in this font .
Document titles	Document titles appear in <i>this font</i> .
TUI elements	In a Text-based User Interface, text the system displays appears in <i>this font</i> .

Text Type	Indication
System output	Terminal sessions and information that the system displays appear in <i>this font</i> .
CLI commands	CLI command keywords appear in this font . Variables in a CLI command appear in <i>this font</i> .
[]	Elements in square brackets are optional.
{x y z}	Required alternative keywords are grouped in braces and separated by vertical bars.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
<>	Nonprinting characters such as passwords are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

**Note**

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the document.

**Caution**

Means *reader be careful*. In this situation, you might perform an action that could result in equipment damage or loss of data.

**Tip**

Means *the following information will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information, similar to a Timesaver.

**Timesaver**

Means *the described action saves time*. You can save time by performing the action described in the paragraph.

**Warning****IMPORTANT SAFETY INSTRUCTIONS**

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Related Documentation

Cisco UCS Director Documentation Roadmap

For a complete list of Cisco UCS Director documentation, see the *Cisco UCS Director Documentation Roadmap* available at the following URL: http://www.cisco.com/en/US/docs/unified_computing/ucs/ucs-director/doc-roadmap/b_UCSDirectorDocRoadmap.html.

Cisco UCS Documentation Roadmaps

For a complete list of all B-Series documentation, see the *Cisco UCS B-Series Servers Documentation Roadmap* available at the following URL: <http://www.cisco.com/go/unifiedcomputing/b-series-doc>.

For a complete list of all C-Series documentation, see the *Cisco UCS C-Series Servers Documentation Roadmap* available at the following URL: <http://www.cisco.com/go/unifiedcomputing/c-series-doc>.

**Note**

The *Cisco UCS B-Series Servers Documentation Roadmap* includes links to documentation for Cisco UCS Manager and Cisco UCS Central. The *Cisco UCS C-Series Servers Documentation Roadmap* includes links to documentation for Cisco Integrated Management Controller.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to ucs-director-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see [What's New in Cisco Product Documentation](#).

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the [What's New in Cisco Product Documentation RSS feed](#). RSS feeds are a free service.



CHAPTER

1

New and Changed Information for this Release

- [New and Changed Information, page 1](#)

New and Changed Information

The following table provides an overview of the significant changes to this guide for this current release. The table does not provide an exhaustive list of all changes made to this guide or of all new features in this release.

Table 1: New Features and Changed Behavior in Cisco UCS Director, Release 6.0

Feature	What's New	Where Documented
IBM Storwize V5000 and IBM Flash System V9000 support	<ul style="list-style-type: none"> • Managing IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage Accounts • Managing MDisks • Managing Storage Pools • Managing Volumes • Managing Hosts • Managing File Storage • Managing Copy Services • FlashCopy Mappings 	<ul style="list-style-type: none"> • Managing IBM Storwize and IBM SAN Volume Controller Storage Accounts, on page 5 • Managing MDisks, on page 13 • Managing Storage Pools, on page 21 • Managing Volumes, on page 25 • Managing Hosts, on page 37 • Managing File Storage, on page 47 • Managing Copy Services, on page 65 • FlashCopy Mappings, on page 77



Overview

This chapter contains the following sections:

- [VersaStack, page 3](#)
- [Discovering Devices for IBM Storwize Accounts, page 4](#)

VersaStack

VersaStack is an integrated infrastructure solution based on the IBM Storwize family of virtualized storage technologies and the Cisco Unified Computing System (Cisco UCS) integrated infrastructure. This integrated infrastructure combines the performance and innovation of Cisco UCS integrated infrastructure, which includes Cisco UCS, Cisco Nexus, and Cisco UCS Director, with the performance and efficiency of the IBM Storwize system.

VersaStack is also Cisco Application Centric Infrastructure (ACI) ready, which is a holistic architecture with centralized automation and policy-driven application profiles. Cisco ACI delivers software flexibility with the scalability of hardware performance.

VersaStack is backed by Cisco Validated Designs that provides faster delivery of applications, greater IT efficiency, and less risk. The Cisco Validated Design for IBM VersaStack is targeted at virtualized infrastructure by combining the Cisco UCS integrated infrastructure and IBM Storwize with VMware vSphere hypervisor.

For more information, see [VersaStack Solution by Cisco and IBM](#).

Discovering Devices for IBM Storwize Accounts

Use the **Device Discovery Wizard** to discover devices for IBM Storwize and Cisco Nexus OS type accounts that are used for VersaStack integrated infrastructure.

-
- Step 1** On the menu bar, choose **Administration > Physical Accounts**.
 - Step 2** Click the **Physical Accounts** tab.
 - Step 3** Choose the storage (IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller) account.
 - Step 4** Click **Device Discovery**.
 - Step 5** On the **Overview** page of the **Device Discovery Wizard**, click **Next**.
 - Step 6** On the **Policy** page of the **Device Discovery Wizard**, specify a range or comma separated IP addresses for devices or account types in the **IP Addresses** field and add credential policies for the devices or account types (IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller).
 - Step 7** Click **Next**.
 - Step 8** On the **Discover and Assign** page of the **Device Discovery Wizard**, click **Discover** to perform a discovery and obtain a list of discovered devices.
 - Step 9** Choose the device that you want to add it to the pod.
 - Step 10** Choose the pod from the **Select Pod** drop-down list.
 - Step 11** Click **Add**.
 - Step 12** Click **Close**.
-



Managing IBM Storwize and IBM SAN Volume Controller Storage Accounts

This chapter contains the following sections:

- [IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage Systems](#), page 5
- [Managing IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage Accounts](#), page 6
- [IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage System Reports](#), page 9

IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage Systems

Cisco UCS Director supports IBM Storwize, IBM Storage Array (IBM Storwize V5000 and IBM Flash System V9000), and IBM SAN Volume Controllers (SVCs). These systems can manage a combination of an internal storage and external storage systems.

- **Internal Storage**—A collection of drives directly attached to the system that can be of different sizes and use different technologies.
- **External Storage**—A collection of RAID-protected storage systems that are connected to the system by the storage area network (SAN).

IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Features

Cisco UCS Director can manage the following IBM Storwize, IBM Storage Array, and IBM SVCs features:

- Internal Drives (Displays the internal drives report only)
- MDisks
- Storage Pools

- Volumes
- Hosts
- File Systems (Supported only in IBM Storwize V7000 Unified)
- File Sets (Supported only in IBM Storwize V7000 Unified)
- Shares (Supported only in IBM Storwize V7000 Unified)
- Quotas (Supported only in IBM Storwize V7000 Unified)
- Snapshots (Supported only in IBM Storwize V7000 Unified)
- Copy Services that includes FlashCopy, Remote Copy, and Consistency Groups for FlashCopy and Remote Copy

Managing IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage Accounts

Cisco UCS Director can manage IBM Storwize, IBM Storage Array, and IBM SVCs system features.

Navigate to **Administration > Physical Accounts**, and click the **Physical Accounts** tab to do the following tasks:

- **Refresh**—Refreshes the storage accounts in Cisco UCS Director.
- **Add**—The **Add Account** dialog box allows you to add a storage account of type IBM Storwize, IBM Storage Array, or IBM SVCs.
- **Device Discovery**—The **Discovery Device Wizard** helps you to discover devices or accounts and assign them to a pod.
- **View**—The **View Infrastructure Account** dialog box displays the storage account details.
- **Edit**—The **Edit Account** dialog box displays editable fields to edit a storage account.
- **Delete**—The **Delete Account** dialog box allows you to delete a storage account.
- **Test Connection**—The **Test Connectivity** dialog box displays the Connection Successful message after establishing a connection with IBM Storwize, IBM Storage Array, or IBM SVCs systems.

Adding a Pod

-
- Step 1** On the menu bar, choose **Administration > Physical Accounts**.
 - Step 2** Click the **Pods** tab.
 - Step 3** Click **Add**.
 - Step 4** In the **Add Pod** dialog box, complete the following fields:

Name	Description
Name field	A descriptive name for the pod.
Site drop-down list	Choose the site where you want to add the pod. If your environment does not include sites, you can omit this step.
Type drop-down list	<p>Choose the type of pod that you want to add. This can be one of the following:</p> <ul style="list-style-type: none"> • Flexpod • VersaStack • Generic • ExpressPod Medium • VSPEX • ExpressPod Small • Vblock <p>The nongeneric pod types accommodate only specific physical and virtual components. A generic pod does not require a specific pod license. You can add any type of physical or virtual component to a generic pod. For more information about bundled pod licenses (FlexPod, Vblock, and VSPEX), which include the necessary individual device licenses to run a pod, see the Cisco UCS Director Installation and Upgrade Guides.</p> <p>Note Only VersaStack and Generic pods are supported in the IBM accounts in Cisco UCS Director.</p>
Description field	(Optional) A description of the pod.
Address field	The physical location of the pod. For example, this field could include the city or other internal identification used for the pod.
Hide Pod check box	<p>Check the check box to hide the pod if you do not want it to show in the Converged Check View. You can continue to add or delete accounts from the pod.</p> <p>For example, you can use this check box to ensure that a pod that does not have any physical or virtual elements is not displayed in the Converged View.</p>

Step 5 Click **Add**.

What to Do Next

Add one or more accounts to the pod.

Adding an IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller Storage Account

You must create a storage account of type IBM Storwize, IBM Storage Array (IBM Storwize V5000 and IBM Flash System V9000), or IBM SAN Volume Controller to manage storage devices in Cisco UCS Director.

Step 1 On the menu bar, choose **Administration > Physical Accounts**.

Step 2 Click the **Physical Accounts** tab.

Step 3 Click **Add**.

Step 4 In the **Add Account** dialog box, complete the following fields:

Name	Description
Pod drop-down list	Choose the pod to which the storage account has to be added. For example, VersaStack.
Category drop-down list	Choose Storage . This is the type of infrastructure for the storage accounts
Account Type drop-down list	Choose an account of type IBM Storwize, IBM Storage Array (IBM Storwize V5000 and IBM Flash System V9000), or IBM SAN Volume Controller.

Step 5 Click **Submit**.

Step 6 In the **Add Account** dialog box, complete the following fields:

Name	Description
Pod	The pod to which the storage account is added.
Account Type	The storage account of type IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller.
Account Name field	A unique name that you want to assign to this account.
Description field	A description for this account.
IBM Storage Management Address field	The IP address of the IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller storage system.
User Name field	A user name assigned to access and establish connectivity with the IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller storage system.

Name	Description
Password field	A password assigned to access and establish connectivity with the IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller storage system.
Protocol drop-down list	The protocol that is used to access IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller storage system. By default, it is ssh .
Port field	By default, the port number displayed is 22.
Contact field	An e-mail address of the administrator or the person to contact for this account.
Location field	The location for this account.

Step 7 Click **Submit**.

Note To view the MDisk report, you should be a Security Administrator or Administrator of the IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller account.

What to Do Next

Choose an account of type IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller and click **Test Connection** that connects to the appropriate storage system. The **Test Connectivity** dialog box should display the Connection Successful message.

IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage System Reports

Cisco UCS Director user interface displays the following reports that helps you to manage IBM Storwize, IBM Storage Array, and IBM SVCs systems:

- Summary of IBM Storwize and IBM Storage array accounts
- IBM Storwize Features: Internal Drives, MDisk, Storage Pools, Volumes, Hosts, File Systems, File Sets, File Shares, Quotas, Snapshots, Snapshot Rules, Remote Copy, Remote Copy Consistency Groups, FlashCopy, FlashCopy Consistency Groups, and Flashcopy Mappings
- System Hardware
- FC ports
- Node Canisters
- System Tasks

Managing System Tasks for IBM Storwize, IBM Storage Array, and IBM SAN Volume Controller Storage Accounts

When you add a storage account of type IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller, a new system task (inventory collector) is created for each storage account in the IBM Storage Tasks folder. For removing the data from the database table of removed accounts in Cisco UCS Director, an IBM Storwize Account Cleanup task is also made available. The system task connects to IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller, collects an inventory of the system, and persists data in the Cisco UCS Director database. This collector adds to the system scheduler so that it can be called periodically as per the frequency configured in the collector (for example, one hour).

System Task is used not only for collecting inventory data from the storage devices, but also used.

For more information on how to manage system tasks in Cisco UCS Director, see the latest *Cisco UCS Director Administration Guide*.

To manage the system tasks of storage accounts, you can use the **System Tasks** tab in **Administration > System**.

- **Manage Task**—The **Manage Task** dialog box allows you to enable inventory collection for the selected storage account at specified intervals with respect to the system task policy or disable inventory collection.
- **Run Now**—The **Run Now** dialog box allows you to run the selected system task (inventory collection) and update the storage account information in Cisco UCS Director irrespective of the frequency configured in the collector.



Managing Internal Drives

This chapter contains the following sections:

- [Internal Storage, page 11](#)
- [Internal Storage RAID Configurations, page 12](#)

Internal Storage

Internal storage is the RAID-protected storage that is directly attached to the system. The internal drives can be of different sizes and use different technologies. The IBM Storwize, IBM Storage Array, and IBM SAN Volume Controllers support a range of enterprise-class, nearline-class, or flash drives and they automatically detect internal drives that are attached to them.

Internal drives can be designated as either a member, a candidate, or a spare.

- Member—The drive that is in used state.
- Candidate—The drive can be used to create an array or changed to a spare drive.
- Spare—The system uses this drive to replace a failed drive. Spares can also be used to replace faulty drives.

You can view the list of internal drives and their properties in the **Internal Drives** report for that selected account in Cisco UCS Director. The report displays the following physical and logical properties of all internal drives, and the details of drive slots of an IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller:

- Drive ID—The system assigned identifier for the drive.
- Status—Reports the availability between the system and the drive.
- Use—The current use or role of the drive. This attribute determines if the drive can be formed into an array.
- Technology Type—The technology type of the drive as either flash, enterprise, and nearline drives.
- Capacity—The total amount of usable space that is available on the drive.
- MDisk ID—The identifier for the MDisk.
- MDisk Name—The name of the MDisk that uses the drive.

- Member ID—The position of the drive in the array.
- Enclosure ID—The identifier for the enclosure that contains the drive slot.
- Slot ID—The location of the drive within its enclosure.

Internal Storage RAID Configurations

Internal drives are configured into RAID (Redundant Array of Independent Disks) arrays and presented as MDisks to the system. When a RAID member (drive) of an array fails, the system automatically replaces the failed member with a hot-spare drive and re-synchronizes the array to restore its redundancy.

- RAID-0—Provides no protection against drive failures which can be used for only temporary volumes. You require a minimum of two drives for this RAID level.
- RAID-1—Mirrors data for good performance and provides protection against drive failures. You require a minimum of two drives for this RAID level.
- RAID-5—Protects against a single drive failure. Data and one strip of parity are striped across all array members. You require a minimum of three drives for this RAID level and a maximum of 16 drives.
- RAID-6—Protects against two drive failures. Data and two strips of parity are striped across all array members. You require a minimum of five drives for this RAID level and a maximum of 16 drives.
- RAID-10—Provides good performance and protects against at least one drive failure. All data is mirrored on two array members. You require a minimum of two drives for this RAID level and a maximum of 16 drives.



Managing MDisks

This chapter contains the following sections:

- [MDisks, page 13](#)
- [Encryption Enabled Array MDisks, page 14](#)
- [Creating an Array MDisk, page 14](#)
- [Renaming an MDisk, page 15](#)
- [Deleting an Array MDisk, page 16](#)
- [Assigning an MDisk to a Storage Pool, page 16](#)
- [Unassigning an MDisk from the Storage Pool, page 17](#)
- [MDisk Tier, page 17](#)
- [Setting Tier, page 18](#)
- [Swapping Internal Drives, page 18](#)
- [Setting Spare Drives, page 19](#)

MDisks

An MDisk (managed disk) is a logical unit of the physical storage that consists of multiple physical disks. It provides usable blocks of physical storage to the system even if it does not have a one-to-one correspondence with a physical disk. 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 storage pools. The size of the extents that the MDisk is divided into depends on the attribute of the storage pool to which it has been added.

MDisks are either RAID arrays from an internal storage or volumes from an external storage system that are presented as a single logical disk to the storage area network (SAN). MDisks, either RAID arrays from an internal storage or volumes that are discovered from an external storage can be added to storage pools. MDisks are not visible to host systems.

Cisco UCS Director allows you to perform the following tasks on the selected MDisk:

- **Refresh**—Refreshes the list of MDisks for the selected IBM Storwize, IBM Storage Array, or IBM SVCs account.

- **Favorite**—Adds the MDisks report to the favorite list for the selected IBM Storwize, IBM Storage Array, or IBM SVCs account.
- **Create**—Creates a RAID array MDisk from an internal storage.
- **Rename**—Renames MDisks.
- **Unassign MDisk from Pool**—Removes an MDisk from a storage pool it is currently assigned to, even if it contains data on it. The system migrates the data to other MDisks in the storage pool.
- **Set Spare Goal**—Sets an optimal number of spares that are required to protect the array from drive failures.
- **Swap Drive**—Allows to swap internal drives into an MDisk or out of an MDisk.
- **Delete Array MDisk**—Deletes the RAID array MDisk, even if it contains data on it. The system migrates the data to other MDisks in the storage pool.
- **Select Tier**—Allows the MDisk to modify the existing enterprise-class, nearline-class, or flash drive technology.
- **Assign MDisk to Pool**—Assigns an MDisk to a storage pool.
- **Detect MDisks**—Identifies MDisks both from an internal and an external IBM Storwize, IBM Storage Array, or IBM SAN Volume Controller storage systems.

Encryption Enabled Array MDisks

The IBM Storwize V7000 Gen2 system provides optional encryption of data at rest, which protects exposure of sensitive user data and user metadata that is stored on storage devices. For example, encryption protects exposure of sensitive data that is stored on discarded, lost, or stolen storage devices. Encryption of system data and system metadata is not required, and so system data and metadata are not encrypted.

You can activate the encryption license and enable encryption on the system to create encrypted array MDisks.

Creating an Array MDisk

An internal drive that is designated as a candidate can be used to create an array or a spare. You can create an encrypted array MDisk if encryption is enabled on your system.

-
- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type in which you want to create MDisks.
- Step 4** Click the **MDisks** tab.
- Step 5** Click **Create**.
- Step 6** In the **Create Array MDisk** dialog box, complete the following fields:

Name	Description
MDisk Name field	A unique name that you want to assign to the MDisk.

Name	Description
Drive(s) table	Drives that are marked as candidates are available for you to create an array. 1 Click Select and choose one or more internal drives. 2 Click Select .
RAID Level drop-down list	Choose the RAID level configuration. This field is not displayed when you select the account type as IBM Flash System V9000.
Storage Pool table	Choose the storage pool. 1 Click Select and choose a storage pool. 2 Click Select . This field is not displayed when you select the account type as IBM Flash System V9000.
Enable encryption check box	Check the check box to enable encryption.

Step 7 Click **Submit**.

Renaming an MDisk

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **MDisks** tab.
 - Step 5** Choose an MDisk that you want to rename.
 - Step 6** Click **Rename**.
 - Step 7** In the **Rename MDisk** dialog box, enter a new name for the MDisk.
 - Step 8** Click **Submit**.
-

Deleting an Array MDisk

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **MDisks** tab.
 - Step 5** Choose an MDisk that you want to delete.
 - Step 6** In the **Delete Array MDisk** dialog box, check the **Delete RAID array MDisk even if it has data on it. System migrates data to other MDisks in the pool** check box.
 - Step 7** Click **Submit**.
-

Assigning an MDisk to a Storage Pool

The mode of an MDisk can be either an array of drives that are located on the internal storage or unmanaged while the MDisk is not used by the storage pool. It can also be an image that is assigned to a volume with a direct mapping of data between the MDisk and the volume.

You can assign only unmanaged MDisks to a storage pool.

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **MDisks** tab.
 - Step 5** Choose an unmanaged MDisk that you want to assign to a storage pool.
 - Step 6** Click **Assign MDisk to Pool**.
 - Step 7** In the **Assign MDisk to Pool** dialog box, click **Select**.
 - Step 8** Check the check box to choose a storage pool to which you want to assign the MDisk.
 - Step 9** Click **Select**.
 - Step 10** Click **Submit**.
-

Unassigning an MDisk from the Storage Pool

-
- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **MDisks** tab.
- Step 5** Choose an MDisk that you want to unassign from the storage pool.
- Step 6** Click **Unassign MDisk from Pool**.
- Step 7** In the **Unassign MDisk from Pool** dialog box, check the **Unassign MDisk from storage pool even if it has data on it. System migrates data to other MDisks in the pool** check box.
- Step 8** Click **Submit**.
-

MDisk Tier

The IBM Storwize and IBM SAN Volume Controllers automatically and nondisruptively move frequently accessed data on MDisks that use enterprise-class or nearline-class drives to MDisks that use flash drives, thus placing such data in a faster tier of storage. The flash MDisks provide greater performance than enterprise-class or nearline-class MDisks.

IBM Storwize and IBM SAN Volume Controllers support the following tiers:

- Enterprise—The enterprise tier exists when enterprise-class MDisks are in the storage pool, such as those drives built from Serial Attached SCSI (SAS) drives.
- Nearline—The nearline tier exists when nearline-class MDisks are used in the storage pool, such as those drives built from nearline SAS drives.
- Flash—The flash tier exists when flash MDisks are in the storage pool.

Setting Tier

The tier indicates whether an MDisk uses an enterprise-class, nearline-class, or flash drives.

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **MDisks** tab.
 - Step 5** Choose an MDisk.
 - Step 6** Click **Select Tier**.
 - Step 7** In the **Select MDisk Tier** dialog box, choose the desired tier from the drop-down list.
 - Step 8** Click **Submit**.
-

Swapping Internal Drives

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **MDisks** tab.
 - Step 5** Choose an array MDisk in which you want to swap in and swap out an internal drive.
 - Step 6** Click **Swap Drive**.
 - Step 7** In the **Swap Drive** dialog box, click **Select** to choose an internal drive to swap out of an MDisk or swap into an MDisk.
 - Step 8** Check the check box that corresponds to an internal drive ID to swap out or swap into an internal drive.
 - Step 9** Click **Submit**.
-

Setting Spare Drives

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **MDisks** tab.
 - Step 5** Choose an MDisk.
 - Step 6** Click **Set Spare Goal**.
 - Step 7** In the **Set Spare Goal** dialog box, enter the number of spares you want to set.
 - Step 8** Click **Submit**.
-



CHAPTER 6

Managing Storage Pools

This chapter contains the following sections:

- [Parent Pools, page 21](#)
- [Child Pools, page 21](#)
- [Creating a Storage Pool, page 22](#)
- [Renaming a Storage Pool, page 22](#)
- [Deleting a Storage Pool, page 23](#)
- [Expanding a Storage Pool, page 23](#)

Parent Pools

The parent pools receive their capacity from MDisks. A storage pool is a collection of MDisks that jointly contain all of the data for a specified set of volumes. All MDisks in a storage pool are split into extents of the same size. 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. Volumes are created from the extents that are available in the pool.

Child Pools

The child pools are similar to the parent pools with similar properties and they can be used for volume copy operation. Instead of creating directly from MDisks, you can create child pools from the existing capacity that is allocated to a parent pool. You can create volumes that specifically use the capacity that is allocated to the child pool as with parent pools. The capacity of the child pool must be smaller than the free capacity that is available to the parent pool. The child pools are created with fully allocated physical capacity. The allocated capacity of the child pool is no longer reported as the free space of its parent pool. You can create and change the child pools only with the command-line interface of the IBM Storwize storage device but you can view the child pools and their properties in the management GUI of the IBM Storwize storage device.

You can drill down the Storage Pools report in Cisco UCS Director to view the child pool report.

Creating a Storage Pool

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod to which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Storage Pools** tab.
- Step 5** Click **Create**.
- Step 6** In the **Create Storage Pool** dialog box, complete the following fields:

Name	Description
Pool Name field	The storage pool name.
Extent Size drop-down list	Choose the size of an extent. It ranges from 16 MB to 8 GB.
Maximum Addressible Size field	4 PB.
Warning Threshold (0 to 100%) field	Set the warning threshold (in terms of percentage) to send out an alert to the administrator when the usage reaches the threshold value.
MDisks drop-down list	Choose the unmanaged (not used by the system) MDisk from the list to include in this pool.

- Step 7** Click **Submit**.

Renaming a Storage Pool

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod to which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Storage Pools** tab.
- Step 5** Choose the storage pool that you want to rename.
- Step 6** Click **Rename**.
- Step 7** In the **Rename Storage Pool** dialog box, enter the new name for the storage pool.
- Step 8** Click **Submit**.

Deleting a Storage Pool

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod to which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Storage Pools** tab.
 - Step 5** Choose the storage pool that you want to delete.
 - Step 6** Click **Delete**.
 - Step 7** In the **Delete Storage Pool** dialog box, check the **Delete all volumes, host mappings, and MDisks associated with pool** check box.
 - Step 8** Click **Submit**.
-

Expanding a Storage Pool

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Storage Pools** tab.
 - Step 5** Choose the storage pool that you want to expand.
 - Step 6** Click **Expand**.
 - Step 7** In the **Expand Pool** dialog box, click **Select**.
 - Step 8** Choose an MDisk that you want to use to expand the storage pool.
 - Step 9** Click **Submit**.
-



Managing Volumes

This chapter contains the following sections:

- [Volumes, page 25](#)
- [Creating Volumes, page 27](#)
- [Adding Mirrored Copy, page 28](#)
- [Deleting Volumes, page 29](#)
- [Editing Volumes, page 29](#)
- [Expanding Volumes, page 30](#)
- [Migrating Volumes, page 30](#)
- [Renaming Volumes, page 31](#)
- [Shrinking Volumes, page 31](#)
- [Mapping to Host, page 32](#)
- [Unmapping All Hosts, page 32](#)
- [Unmapping A Single Host, page 33](#)
- [View Details of Volume Mirrored Copies, page 33](#)

Volumes

One or more volumes are allocated physical storage space from each pool.

- 1 **File Volume**—A logical disk that the system presents it through the file systems. In comparison to the block volumes, it is a much simpler form of storage for users to access and share. A file volume is available using common file-level protocols, such as Server Message Block (SMB), Common Internet File System (CIFS), and Network File System (NFS), using network-attached storage (NAS) devices.
- 2 **Block Volume**—A logical disk that the system presents the physical storage space allocated from a storage pool to the attached hosts.

File Volumes

A file is the physical storage assigned to a file system pool. The file volume, in turn, is a logical disk that the system presents to the file systems.

- **File System Pools**—A file system pool consists of volumes of the same disk type. The file system pools are a type of storage pools that consists of a collection of disks that jointly contain all the data across a specified set of file systems.
- **File Systems**—A file system consists of one or more file system pools. The storage pools must be defined before creating the file system. You can create, edit, mount, unmount, set as destination, or remove a file system. You can select either one of the two preset policies for either a single file system pool or a migration, or create a custom policy.
- **File Sets**—A file set is a subset of a file system that provides granularity of functions such as snapshots or quotas within the file system.
- **File Shares**—A shared disk space that can be made available through the protocols specified during its creation. Hypertext Transfer Protocol (HTTP), Secure Copy Protocol (SCP), File Transfer Protocol (FTP), Common Internet File System (CIFS), and Network File System (NFS) shares can all be created, provided that the corresponding protocol is enabled for the system.
- **Quotas**—A quota is used to control the allocation of files and data blocks in a file system.
- **Snapshots**— A snapshot of an entire file system or of an independent file set can be created to preserve the contents of the file system or the independent file set at a single point in time.

Block Volumes

A block volume is a logical disk that the system presents to a host connected over a fibre channel or Ethernet network. Each volume copy is created from a set of extents in a storage pool. Each volume can have a maximum of two copies. Host applications access volumes, not MDisks or drives. To keep a volume accessible even when an MDisk on which it depends has become unavailable, a mirrored copy can be added to a selected volume. Volumes are presented to hosts by host mapping. Host mapping controls which hosts have access to specific volumes in the storage system.

You can create the following volume types:

- **Generic**—Creates volumes that use a set amount of capacity from the selected storage pool.
- **Mirror**—Creates volumes with two physical copies. Each copy can belong to a different storage pool that provide data protection from storage failures.
- **ThinMirror**—Creates volumes with two physical copies that provide data protection from storage failures while using only the capacity that is written by the host application.
- **ThinProvision**—Creates large-capacity volumes that use only the capacity that is written by the host application.
- **Compressed**—Creates volumes where data can be compressed when it is stored on a disk.
- **CompressedMirror**—Creates volumes where data can be compressed and stored with two physical copies that provide data protection from storage failures.

Creating Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Volumes** tab.
- Step 5** Click **Create**.
- Step 6** In the **Create Volume** dialog box, complete the following fields:

Name	Description
Volume Name field	The volume name. If you do not specify a name, the system assigns a name automatically.
Volume Type drop-down list	Choose the type of volume that you want to create.
Capacity field	Enter the volume capacity in terms of bytes, KB, MB, GB, or TB.
Capacity Unit drop-down list	Choose the unit in bytes, KB, MB, GB, or TB for the volume capacity.
Pool drop-down list	Choose the storage pool.
Secondary Pool drop-down list	Choose the secondary pool.
Mirror Sync Rate field	Enter the rate (%) at which copies (mirroring) can be synchronized. A value of zero prevents synchronization.
Preferred Node drop-down list	Choose the host. Note Hosts try to access the volumes through the preferred node. Each volume has a preferred node that balances the load between nodes. By default, the system automatically balances the load between nodes.
Real Capacity Unit drop-down list	Choose the real capacity that is allocated to the volume in terms of % of Virtual Capacity or GB.
Automatically Expand check box	Check this check box to automatically increase the real capacity as data is written to the volume.

Name	Description
Warning Threshold check box	Check this check box to manually increase the real capacity when the warning threshold exceeds in terms of % of Virtual Capacity. Note When the used capacity exceeds the warning threshold, an event is raised, indicating that the additional real capacity is required. You can manually increase the real capacity that is previously allocated to the volume.
Warning Threshold (% of Virtual Capacity) field	Displays the warning threshold set beyond which you have to manually increase the real capacity.
Thin Provisioned Grain Size drop-down list	Choose the thin-provisioned grain size. It is the size of a single chunk of data on the volume. Note Smaller grain sizes save data while larger grain sizes produce better performance.

Step 7 Click **Submit**.

Adding Mirrored Copy

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Volumes** tab.
- Step 5** Choose the volume that you want to create a copy of the volume in another storage pool.
- Step 6** Click **Add Mirrored Copy**.
- Step 7** In the **Add Mirrored Copy** dialog box, complete the following fields:

Name	Description
Volume Name	The volume name (Display only).
Current Pool	The current pool (Display only).
Volume Type drop-down list	Choose Generic or Thin-provisioned or Compressed type.
Select Pool drop-down list	Choose the storage pool in which you want to create the copy of the volume.

Step 8 Click **Submit**.

Deleting Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Volumes** tab.
- Step 5** Choose the volume that you want to delete.
- Step 6** Click **Delete**.
- Step 7** In the **Delete Volume** dialog box, check the **Force Delete** check box.
- Step 8** Click **Submit**.
-

Editing Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Volumes** tab.
- Step 5** Choose the volume that you want to edit.
- Step 6** Click **Edit**.
- Step 7** In the **Edit Volume** dialog box, complete the following fields:

Name	Description
Volume Name field	The volume name.
Capacity	The volume capacity (Display only).
I/O Group	Display only. The I/O groups through which the hosts can access the volume.
Cache Mode check box	Check this check box to enable the cache for faster I/O operations.

Name	Description
Sync Rate field	Enter the rate (%) at which copies (mirroring) can be synchronized. A value of zero prevents synchronization.

Step 8 Click **Submit**.

Expanding Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Choose the volume for which you want to increase the available capacity.
 - Step 6** Click **Expand**.
 - Step 7** In the **Expand Volume** dialog box, enter the capacity in the **Expand By** field.
 - Step 8** Choose the capacity unit from the **Unit** drop-down list.
 - Step 9** Click **Submit**.
-

Migrating Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Choose the volume from which you want to migrate data from the current storage pool to the new storage pool.
 - Step 6** Click **Migrate**.
 - Step 7** In the **Migrate Volume** dialog box, click **Select**.
 - Step 8** Check the check box to choose the new storage pool to which you want to migrate data.
 - Step 9** Click **Select**.
 - Step 10** Click **Submit**.
-

Renaming Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Choose the volume that you want to rename.
 - Step 6** Click **Rename**.
 - Step 7** In the **Rename Volume** dialog box, enter a new name for the volume.
 - Step 8** Click **Submit**.
-

Shrinking Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Choose the volume for which you want to decrease the available capacity.
 - Step 6** Click **Shrink**.
 - Step 7** In the **Shrink Volume** dialog box, enter the capacity in the **Shrink By** field.
 - Step 8** Choose the capacity unit from the **Unit** drop-down list.
 - Step 9** Click **Submit**.
-

Mapping to Host

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Choose the volume that you want to map to a host.
 - Step 6** Click **Map to Host**.
 - Step 7** In the **Map Volume to Host** dialog box, click **Select**.
 - Step 8** Check the check box to choose the host to which you want to map the volume.
 - Step 9** Click **Select**.
 - Step 10** Click **Submit**.
-

Unmapping All Hosts

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Choose the volume that you want to unmap from the host to which the volume is mapped.
 - Step 6** Click **Unmap All Hosts**.
 - Step 7** In the **Unmap Volume from All Hosts** dialog box, click **Submit**.
-

Unmapping A Single Host

-
- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Volumes** tab.
- Step 5** Choose the volume that you want to unmap from a single host to which the volume is mapped.
- Step 6** Click **Unmap Single Host**.
- Step 7** In the **Unmap Single Host** dialog box, click **Select**.
- Step 8** Choose the host from which you want to unmap the volume.
- Step 9** Click **Select**.
- Step 10** Click **Submit**.
-

View Details of Volume Mirrored Copies

Use the **Volume** tab and view the details of mirrored copies of volumes.

You can perform the following actions:

- **Make Primary Volume**
- **Split into New Volume**
- **Duplicate**
- **Delete Mirrored Copy**

Viewing Volume Mirrored Copies

-
- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Volumes** tab.
- Step 5** Choose the volume mirrored copy.
- Step 6** Click the down-arrow, and then click **View Details**.
- Step 7** Click **Volume Mirrored Copies** tab to perform actions on the volume mirrored copies.
-

Making Primary Volume

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Choose the volume copy that you want to make it as the primary volume.
 - Step 6** Click **Make Primary Volume**.
 - Step 7** Click **Select**.
 - Step 8** Check the check box to choose the volume copy that can be made as the primary volume.
 - Step 9** Click **Select**.
 - Step 10** Click **Submit**.
-

Splitting Volumes into New Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Choose the volume copy that you want to split into a new volume.
 - Step 6** Click **Split into New Volume**.
 - Step 7** In the **Splitting Volume into New Volume** dialog box, enter the new volume name.
 - Step 8** Click **Submit**.
-

Duplicating Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Choose the volume that you want to create a copy of the volume.
 - Step 6** Click **Duplicate**.
 - Step 7** In the **Duplicate Volume** dialog box, enter the new volume name.
 - Step 8** Click **Submit**.
-

Deleting Mirrored Copy

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Volumes** tab.
 - Step 5** Click the down-arrow, and then click **View Details**.
 - Step 6** In the **Volume Mirrored Copies** tab, choose the volume mirrored copy that you want to delete.
 - Step 7** Click **Delete Mirrored Copy**.
 - Step 8** In the **Delete Mirrored Copy** dialog box, check the **Force Delete** check box to delete the volume mirrored copy.
 - Step 9** Click **Submit**.
-



Managing Hosts

This chapter contains the following sections:

- [Hosts, page 37](#)
- [Creating Hosts, page 38](#)
- [Deleting Hosts, page 39](#)
- [Renaming Hosts, page 40](#)
- [Unmapping All Volumes, page 40](#)
- [Importing Mappings, page 41](#)
- [Duplicating Mappings, page 41](#)
- [Deploying an ESXi Fibre Channel and Fibre Channel over Ethernet Host on IBM Storwize System, page 42](#)
- [Deploying an ESXi iSCSI Host on IBM Storwize System, page 43](#)
- [Creating and Adding Block Volume Datastore on the ESXi Host, page 45](#)
- [Adding Block Datastore over ESXi on IBM Storwize System, page 46](#)
- [Adding Compressed Datastore over ESXi on IBM Storwize System, page 46](#)

Hosts

Hosts connect to the storage system either through a fibre channel interface or an IP network. The system does not automatically present volumes to the attached hosts. You must map each volume to a particular host object to enable the volume to be accessed through the worldwide port names (WWPNs) or iSCSI names that are associated with the host object. A host object is a logical object that represents a list of WWPNs and a list of iSCSI names, which means that the host object can have both WWPNs and iSCSI names. These WWPNs and iSCSI names identify the interfaces that the hosts use to communicate with the storage system. iSCSI names can be either iSCSI qualified names (IQNs) or extended unique identifiers (EUIs).

- **iSCSI hosts**—An iSCSI-attached host is a system that is connected to the storage through an IP network. When you configure a host on the storage system, you can create a host object that is identified by iSCSI names that are used to communicate with the system.

- Fibre Channel hosts—A fibre channel host is an application server that is connected to the system through the fibre channel network. Before the host can access the volumes, you must create a fibre channel host object on the system. You can configure the host object with the WWPN of the host machine and specify which volumes the host machine can access.

Creating Hosts

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Hosts** tab.
- Step 5** Click **Create**.
- Step 6** In the **Create Host** dialog box, complete the following fields:

Name	Description
Host Type drop-down list	Choose the type of host that you want to connect to the storage system.
Host Name field	The host name.
Fibre Channel Port Definitions field	Enter the world wide port name (WWPN) that identifies the host to the system. A WWPN consists of 16 (hexadecimal) digits. In case of multiple ports, separate port numbers by a comma. See Fibre Channel Port Definitions in the FC Ports tab.
iSCSI Port Definitions field	Enter the iSCSI name that can be either in the iSCSI qualified names (IQNs) or in the extended unique identifiers (EUIs) format. In case of multiple ports, separate port numbers by a comma.
Use CHAP Authentication (all ports) check box	Check this check box for CHAP authentication for iSCSI-attached hosts.
Use CHAP Authentication Name field	Enter a name for the shared secret for CHAP authentication.
I/O Group	Choose one or more I/O groups for the host to access volumes from the selected I/O group(s).

Name	Description
Host Type drop-down list	Choose one of the following host types: <ul style="list-style-type: none"> • Generic—Choose this type for most supported host systems. • HP/UX—Choose this type for HP host systems. • OpenVMS—Choose this type for OpenVMS host systems.. • TPGS (Target Port Group Support) —Choose this type for Solaris or Apple host systems that require a load balancing.

Step 7 Click **Submit**.

Deleting Hosts

Step 1 On the menu bar, choose **Physical > Storage**.

Step 2 In the left pane, navigate to the pod in which the storage account is added.

Step 3 Choose the storage account type.

Step 4 Click the **Hosts** tab.

Step 5 Choose the host that you want to delete.

Step 6 Click **Delete**.

Step 7 Check the **Delete hosts even it has volumes are mapped to them. These volumes will no longer be accessible to hosts** check box.

Step 8 Click **Submit**.

Renaming Hosts

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Hosts** tab.
 - Step 5** Choose the host for which you want to change the name of the host.
 - Step 6** Click **Rename**.
 - Step 7** In the **Rename Host** dialog box, enter a new host name.
 - Step 8** Click **Submit**.
-

Unmapping All Volumes

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Hosts** tab.
 - Step 5** Choose the host from which you want to unmap all volumes that are mapped to the host.
 - Step 6** Click **Unmap All Volumes**.
 - Step 7** Click **Submit**.
-

Importing Mappings

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Hosts** tab.
 - Step 5** Choose the target host from which you want to import the host mappings.
 - Step 6** Click **Import Mappings**.
 - Step 7** In the **Import Mappings** dialog box, click **Select**.
 - Step 8** Check the check box to chose the source host to which you want to import the host mappings from the target host.
 - Step 9** Click **Select**.
 - Step 10** Click **Submit**.
-

Duplicating Mappings

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Hosts** tab.
 - Step 5** Choose the source host from which you want to duplicate the host mappings to one or more targets.
 - Step 6** Click **Duplicate Mappings**.
 - Step 7** In the **Duplicate Mappings** dialog box, click **Select**.
 - Step 8** Check the check box to choose one or more target hosts in which you want to duplicate host mappings from the source.
 - Step 9** Click **Select**.
 - Step 10** Click **Submit**.
-

Deploying an ESXi Fibre Channel and Fibre Channel over Ethernet Host on IBM Storwize System

- Step 1** On the menu bar, choose **Policies > Orchestration**.
- Step 2** In the **Orchestration** pane, the **Workflows** tab is displayed by default.
- Step 3** In the left pane, choose the **IBM Storwize** folder and click the arrow next to the folder to show all the workflows.
- Step 4** Double-click the **Deploy ESXi FC host on IBM Storwize V7000** or **Deploy ESXi FCoE host on IBM Storwize V7000** workflow that opens in the **Workflow Designer**.

Tasks	Description
Create UCS Service Profile	This task allows to create a new UCS Service Profile.
Select UCS Server	This task allows selects UCS Server from the Server Pool or the list of servers according to the conditions specified.
Associate UCS Server Profile	This task allows you to select the UCS Service Profile and associate it with the selected UCS Server.
Set up PXE Boot	This task allows you to specify PXE Boot parameters required.
Create IBM Storwize Volume	This task allows you to create a new volume on the host. The volumes are built from extents in the storage pools, and hence the size can be increased or decreased.
Create IBM Storwize Host (FC or FCoE)	This task allows you to create the host.
Map IBM Storwize Volume To Host	This task allows you to create a new mapping between the volume and the host, which makes the volume to be accessible for input/output operations to the specified host.
Generic Configure SAN Zoning	This task allows you to create SAN zones.
Power On UCS Server	This task lets you to power on the UCS Server and this task uses the UCS Service Profile to turn the power on for the UCS Server. (This task can be used only when there is a UCS Service Profile associated with the UCS Server).
Monitor PXE Boot	This task monitors the PXE Boot that specifies the maximum amount of time to wait for PXE Boot.
Modify UCS Service Profile Boot Policy	This task allows you to specify the new boot policy for the selected UCS Service Profile (modify the UCS Service Profile with the new boot policy that allows to modify the boot order for the UCS Service Profile.)

Tasks	Description
Wait for Specified Duration	This is a generic task that you can use it in the workflow to wait for a specified duration.
Reset UCS Server	This task allows you to perform a hard reset on the UCS Server irrespective of the UCS Service Profile associated with the server.
Wait for Specified Duration	This is a generic task that you can use it in the workflow to wait for a specified duration.
Register Host with VCenter	This task allows you to register a new host within the selected cloud (vCenter) and also place the host exactly where you want to it be on the cluster/datacenter.
IBM Storwize Thin Provision Volume Datastore	This task allows you to mount a thin-provision volume datastore on the ESXi host.

Step 5 Click **Execute Now**.

Deploying an ESXi iSCSI Host on IBM Storwize System

Step 1 On the menu bar, choose **Policies > Orchestration**.

Step 2 In the **Orchestration** pane, the **Workflows** tab is displayed by default.

Step 3 In the left pane, choose the **IBM Storwize** folder and click the arrow next to the folder to show all the workflows.

Step 4 Double-click the **Deploy ESXi iSCSI Host on IBM Storwize V7000** workflow that opens in the **Workflow Designer**.

Tasks	Description
Create UCS Service Profile	This task allows to create a new UCS Service Profile.
Add vNic to UCS Service Profile	This task allows you to add vNIC to the UCS Service Profile.
Add iSCSI vNIC to UCS Service Profile	This task allows you to add iSCSI vNIC to the UCS Service Profile.
IBM Storwize iSCSI Boot Target	This task allows you to configure iSCSI boot targets.

Tasks	Description
Create Service Profile iSCSI Boot Policy	This task allows you to configure the iSCSI Boot Policy for the existing UCSM Service Profile. You can set the iSCSI Boot Parameters and create iSCSI Static Targets.
Associate UCS Server Profile	This task allows you to select the UCS Service Profile and associate it with the selected UCS Server.
Create IBM Storwize Volume	This task allows you to create a new volume on the host. The volumes are built from extents in the storage pools, and hence the size can be increased or decreased. You can Create volumes of type Generic, Thin-Provision, Mirrored and Compressed volumes.
Create IBM Storwize Host (iSCSI)	This task allows you to create a host entry by using a WWPAN/WWNN number of the remote server and associate one or more HBA WWPANs or iSCSI names with a logical host object.
Map IBM Storwize Volume To Host	This task allows you to create a new mapping between the volume and the host, which makes the volume to be accessible for input/output operations to the specified host.
Set up PXE Boot	This task allows you to specify PXE Boot parameters required like OS type for the Host, Host Name, Server IP Address Range, MAC Address, NetMask, Gateway, Name Server, Password and Time Zone.
Power On UCS Server	This task lets you to power on the UCS Server and this task uses the UCS Service Profile to turn the power on for the UCS Server. (This task can be used only when there is a UCS Service Profile associated with the UCS Server)
Monitor PXE Boot	This task monitors the PXE Boot that specifies the maximum amount of time to wait for PXE Boot.
Modify UCS Service Profile Boot Policy	This task allows you to specify the new boot policy for the selected UCS Service Profile (modify the UCS Service Profile with the new boot policy that allows to modify the boot order for the UCS Service Profile.)
Wait for Specified Duration	This is a generic task that you can use it in the workflow to wait for a specified duration.
Restet UCS Server	This task allows you to perform a hard reset on the UCS Server irrespective of the UCS Service Profile associated with the server.

Tasks	Description
Wait for Specified Duration	This is a generic task that you can use it in the workflow to wait for a specified duration.
Register Host with vCenter	This task allows you to register a new host within the selected cloud (vCenter) and also place the host exactly where you want to it be on the cluster/datacenter.

Step 5 Click **Execute Now**.

Creating and Adding Block Volume Datastore on the ESXi Host

Step 1 On the menu bar, choose **Policies > Orchestration**.

Step 2 In the **Orchestration** pane, the **Workflows** tab is displayed by default.

Step 3 In the left pane, choose the **IBM Storwize** folder and click the arrow next to the folder to show all the workflows.

Step 4 Double-click the **Create_and_Add_IBM_LUN_as_VMFS_Datastore** workflow that opens in the **Workflow Designer**.

Tasks	Description
Create IBM Storwize Volume	This task allows you to create a new volume on the host. The volumes are built from extents in the storage pools, and hence the size can be increased or decreased.
IBM Storwize Block Volume Datastore	This task allows you to mount a block volume datastore on the ESXi host

Step 5 Click **Execute Now**.

Adding Block Datastore over ESXi on IBM Storwize System

Step 1 On the menu bar, choose **Policies > Orchestration**.

Step 2 In the **Orchestration** pane, the **Workflows** tab is displayed by default.

Step 3 In the left pane, choose the **IBM Storwize** folder and click the arrow next to the folder to show all the workflows.

Step 4 Double-click the **Add new block datastore to ESXI on V7000** workflow that opens in the **Workflow Designer**.

Tasks	Description
Generic Configure SAN Zoning	This task allows you to create SAN zones.
IBM Storwize Block Volume Datastore	This task allows you to mount a block volume datastore on the ESXi host.

Step 5 Click **Execute Now**.

Adding Compressed Datastore over ESXi on IBM Storwize System

Step 1 On the menu bar, choose **Policies > Orchestration**.

Step 2 In the **Orchestration** pane, the **Workflows** tab is displayed by default.

Step 3 In the left pane, choose the **IBM Storwize** folder and click the arrow next to the folder to show all the workflows.

Step 4 Double-click the **Add new Compressed datastore to ESXI on V7000** workflow that opens in the **Workflow Designer**.

Tasks	Description
Generic Configure SAN Zoning	This task allows you to create SAN zones.
IBM Storwize Compressed Block Volume Datastore	This task allows you to mount a compressed block volume datastore on the ESXi host.

Step 5 Click **Execute Now**.



CHAPTER 9

Managing File Storage

This chapter contains the following sections:

- [File System Pools, page 47](#)
- [File Systems, page 47](#)
- [File Sets, page 53](#)
- [File Shares, page 55](#)
- [Quotas, page 58](#)
- [Snapshots, page 61](#)

File System Pools

A file system pool consists of volumes of the same disk type. File system pools are a type of storage pools that consists of a collection of disks that jointly contain all the data across a specified set of file systems. Policies are used to define file placement and migration to different file system pools.

File Systems

You can create, mount, unmount, set as destination, or remove a file system. A file system consists of one or more file system pools.

You can select one of the following presets to create a file system:

- **Single pool**—Creates a file system with a single pool.
- **Compressed**—Creates a file system with a compressed data pool and a placement policy that places incompressible files to the system pool.
- **Migration-ILM**—Creates a file system that migrates data between a main pool called System and a secondary pool.
- **Custom**—Creates a file system with customized settings.

Creating Single Pool File Systems

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **File Systems** tab.
- Step 5** Click **Single Pool**.
- Step 6** In the **Single Pool** dialog box, complete the following fields:

Name	Description
File System Name field	Enter a unique name for the file system.
Owner	Enter the name of the owner or the UID of the owner of the path. It can be a user name or a combination of domain and user name (for example, admin or domain\admin).
File System Pool Name	Choose the file system pool.
Storage Pool table	Choose the storage pool. <ol style="list-style-type: none"> 1 Click Select and choose the storage pool. 2 Click Select .
File System Size	Enter the size of the file system.
File System Size Units	Display only. The size of the file system in terms of GB.

- Step 7** Click **Submit**.

Creating Compressed File Systems

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **File Systems** tab.
- Step 5** Click **Compressed**.
- Step 6** In the **Create Compressed File System** dialog box, complete the following fields:

Name	Description
File System Name field	Enter a unique name for the compressed file system.
Owner	Displays the name of the owner or the UID of the owner of the path. It can be a user name or a combination of domain and user name (for example, admin or domain\admin). The owner can also be the root.
Owning Group	Displays the name or GID of the group that is associated with the path. The owner group can also be the root.
File System Pool Settings	
System (non compressed) table	Choose the storage pool. 1 Click Select and choose the storage pool. 2 Click Select .
Size (GB) field	Enter the size of the file system.
System (metadata) table Note To enable metadata replication, select a pool from System (metadata)	Choose the storage pool. 1 Click Select and choose the storage pool. 2 Click Select .
Size (GB) field	Enter the size of the file system.
System (compressed) table	Choose the storage pool. 1 Click Select and choose the storage pool. 2 Click Select .
Size (GB) field	Enter the size of the file system.

Step 7 Click **Submit**.

Creating Migration-ILM File Systems

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **File Systems** tab.
- Step 5** Click **Migration-ILM**.
- Step 6** In the **Create Migration-ILM File System** dialog box, complete the following fields:

Name	Description
File System Name field	Enter a unique name for the Migration-ILM file system.
Owner	root
Owning Group	root
File System Pool Settings	
System table	Choose the storage pool. <ol style="list-style-type: none"> 1 Click Select and choose the storage pool. 2 Click Select.
Size (GB) field	Enter the size of the file system.
System (metadata) table Note To enable metadata replication, select a pool from System (metadata)	Choose the storage pool. <ol style="list-style-type: none"> 1 Click Select and choose the storage pool. 2 Click Select.
Size (GB) field	Enter the size of the file system.
Name field	Enter a unique name.
Compressed check box	Check this check box to create a compressed file system.
Storage Pool	Choose the storage pool. <ol style="list-style-type: none"> 1 Click Select and choose the storage pool. 2 Click Select.
Size (GB) field	Enter the size of the file system.

Name	Description
Migration start threshold (%) field	Starts migrating data automatically when the used capacity exceeds the specified value.
Migration stop threshold (%) field	Stops migrating data automatically when the used capacity decreases below the specified value.
Enable Migration Schedule	<p>Check this check box that enables to schedule migration.</p> <ul style="list-style-type: none"> • Type—Allows to choose the frequency to schedule migration. • Time of Day—Allows to choose the time (00:00) of the day for all types of migration. • Day of the Week—Allows to choose the time (00:00) of the day and one of the days by their names in the week, if Once a Week is the Type selected. • Days of the Week—Allows to choose the time (00:00) of the day and multiple days by their names in the week, if Multiple Days a Week is the Type selected. • Day of Month—Allows to choose the time (00:00) of the day and one of the days in the month, if Once a Month is the Type selected. • Days of Month—Allows to choose the time (00:00) of the day and multiple days by their names, if Multiple Days in a Month is the Type selected.

Step 7 Click **Submit**.

Mounting File Systems

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **File Systems** tab.
 - Step 5** Choose the file system that you want to mount.
 - Step 6** Click **Mount**.
 - Step 7** In the **Mount File System** dialog box, click **Submit**.
-

Unmounting File Systems

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **File Systems** tab.
 - Step 5** Choose the file system that you want to unmount.
 - Step 6** Click **Unmount**.
 - Step 7** In the **Unmount File System** dialog box, click **Submit**.
-

Deleting File Systems

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **File Systems** tab.
 - Step 5** Choose the file system that you want to delete.
 - Step 6** Click **Delete**.
 - Step 7** In the **Confirm File System Deletion** dialog box, enter Yes in the text box.
 - Step 8** Click **Submit**.
-

File Sets

A file set is a subset of a file system that provides granularity of functions such as snapshots or quotas within the file system. File sets provide the ability to partition a file system to allow administrative operations at a finer granularity than the entire file system. There are two types of file sets: dependent and independent.

You can create a preset file set, either a basic or custom type.

- **Basic**—If you are creating a basic file set, specify a path and name, and select if it is a dependent or independent file set type.
- **Custom**—If you are creating a custom file set, specify a path and name, and select the file set type (independent or dependent). Depending on the file set type, you can specify quota or snapshot settings.

Creating File Sets

Step 1 On the menu bar, choose **Physical > Storage**.

Step 2 In the left pane, navigate to the pod in which the storage account is added.

Step 3 Choose the storage account type.

Step 4 Click the **File Sets** tab.

Step 5 Click **Create**.

Step 6 In the **Create File Set** dialog box, complete the following fields:

Name	Description
File Set Type drop-down list	Choose the preset file set type.
Junction Path drop-down list	Choose the path of junction. The junction path must be on one of the file systems.
Subdirectory field	Enter the name of the subdirectory to which the file set belongs.
Name field	Enter the name of the new file set.
Owner field	Enter the owner of path or the UID of the owner of the path. It can be a user name or a combination of domain and user name (for example, admin or domain\admin).
Group field	Enter the group or the GID of the group that is associated with the path.
Comments field	Enter comments, if any.
Type drop-down list	Choose the file set type.

Name	Description
Maximum number of inodes field	Specify the maximum number of inodes that the file set can use. This value depends on the available inodes in the file system. Note The maximum number of inodes that need to be available to create a file set is 1024.
Allocated number of inodes field	Specify the number of inodes that are allocated when the file set is created.
Set Quota check box	Check this check box to set a quota on the new file set.
Soft Limit field and Sized drop-down list	Enter the soft limit for the quota.
Hard Limit field and Sized drop-down list	Enter the hard limit for the quota.
Snapshot drop-down list	Choose the snapshot rule.

Step 7 Click **Submit**.

Editing File Sets

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **File Sets** tab.
 - Step 5** Choose the file set that you want to edit.
 - Step 6** Click **Edit**.
 - Step 7** In the **Edit File Set** dialog box, enter details in the editable fields, and then click **Submit**.
-

Deleting File Sets

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **File Sets** tab.
 - Step 5** Choose the file set that you want to delete.
 - Step 6** Click **Delete**.
 - Step 7** In the **Confirm File Set Deletion** dialog box, type Yes in the text box.
 - Step 8** Click **Submit**.
-

File Shares

A share is a shared disk space available through the protocols specified during its creation. You can create Hypertext Transfer Protocol (HTTP), Secure Copy Protocol (SCP), File Transfer Protocol (FTP), Common Internet File System (CIFS), and Network File System (NFS) shares, provided that the corresponding protocol is enabled for the system.

Creating File Shares

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **File Shares** tab.
 - Step 5** Click **Create**.
 - Step 6** In the **Create File Share** dialog box, complete the following fields:

Name	Description
Share Type drop-down list	Choose the file share type.
Path drop-down list	Choose the file share path.
Share Name field	the name of the new file share.
Owner field	Enter the owner of path or the UID of the owner of the path. It can be a user name or a combination of domain and user name (for example, admin or domain\admin).

Name	Description
Group field	Enter the group or the GID of the group that is associated with the path.
Create a New Independent File Set check box	Check this check box to create a new independent file set that allows for snapshots, user and group quotas
Supported Protocols check boxes	<p>Check one or more check boxes to enable the protocols for the file share:</p> <ol style="list-style-type: none"> 1 HTTP 2 FTP 3 SCP 4 Enable CIFS. See CIFS Settings below. 5 Enable NFS. See NFS Settings below.
CIFS Settings	
Read Only Access check box	Check the check box to enable read-only access. If selected, none of the users and user groups is allowed to modify the existing files or create new files on this share.
Browsable by the File System Browser check box	Check the check box to display the file share in a browser that supports file systems, for example, Windows Explorer.
Hide Objects check box	Check the check box to hide files and directories for users who have no read permissions.
Comments field	Enter comments, if any.
Display Access Control Lists with CIFS check box	Check the check box to enable that displays CIFS access control lists through the CIFS protocol.
Cache Policy drop-down list	Choose the caching policy for the client-side offline caching.

Name	Description
Share Permissions	<p>To specify the access restriction to the file shares, do the following:</p> <ol style="list-style-type: none"> 1 Click the + icon. 2 In the Add Entry to Share Permissions dialog box, enter the name of the user or group, or the security identifier (SID) for access to the file share. 3 Choose either to allow or deny access to the file share. 4 Choose the value for the selected type of access. 5 Click Submit.
NFS Settings	
Add NFS Clients	<p>To add NFS clients, do the following:</p> <ol style="list-style-type: none"> 1 Click the + icon. 2 In the Add Entry to NFS Clients dialog box, enter the client host name or IP address. 3 Check the Read Only Access check box that enables only read-only access. 4 Check the Root Squash check box that enables mapping either the root UID or all client UIDs to the anonymous UID. 5 Check the Secure check box to specify the type of security to use when authenticating NFS clients. The default value is System. If the default value is selected, the UNIX UIDs and GIDs are used to authenticate users. 6 Check the All Squash check box. 7 Enter the anonymous UID for the root user in the Anonymous UID field. The default value is -2. 8 Enter the anonymous GID for the root user in the Anonymous GID field. The default value is -2. 9 Click Submit.

Step 7 Click **Submit**.

Deleting File Shares

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **File Shares** tab.
 - Step 5** Choose the file share that you want to delete.
 - Step 6** Click **Delete**.
 - Step 7** In the **Delete File Share** dialog box, click **Submit**.
-

Activating and Deactivating File Shares

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **File Shares** tab.
 - Step 5** Choose the file share that you want to activate or deactivate.
 - Step 6** Click **Activate/Deactivate**.
 - Step 7** Click **Submit**.
-

Quotas

You can define the quota for users, user groups, and file sets in the system. When you set a quota, you need to specify the soft and hard limits.

Creating Quotas

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Quotas** tab.
- Step 5** Choose the file system in which you want to create a quota.
- Step 6** Click **Create**.
- Step 7** In the **Create Quota** dialog box, complete the following fields:

Name	Description
Type drop-down list	Choose the level (file sets, user groups, and users) at which you want to set the quota.
Path drop-down list	Choose the file system path (including the sub directory) to set the quota.
Soft Limit field	Enter the soft limit for the quota.
Soft Limit Unit drop-down list	Choose the soft limit unit.
Hard Limit field	Enter the hard limit for the quota.
Hard Limit Unit drop-down limit	Choose the hard limit unit.
Quota Scope	Choose the file system path to set the quota.
Specify Group IDs and quota sizes	<p>In the Specify Group IDs and Quota Sizes table:</p> <ol style="list-style-type: none"> 1 Click the + icon. 2 In the Add Entry to Specify Group IDs and Quota Sizes dialog box, enter the group ID. 3 Enter the soft limit. 4 Choose the soft limit unit. 5 Enter the hard limit. 6 Choose the hard limit unit. 7 Click Submit.

Name	Description
Specify User IDs and quota sizes	<p>In the Specify User IDs and Quota Sizes table:</p> <ol style="list-style-type: none"> 1 Click the + icon. 2 In the Add Entry to Specify User IDs and Quota Sizes dialog box, enter the user ID. 3 Enter the soft limit. 4 Choose the soft limit unit. 5 Enter the hard limit. 6 Choose the hard limit unit. 7 Click Submit.

Step 8 Click **Submit**.

Editing Quotas

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Quotas** tab.
- Step 5** Choose the type of quota that you want to edit.
- Step 6** Click **Edit**.
- Step 7** In the **Edit Quota** dialog box, edit the quota sizes that you have already specified for soft limit and hard limit.
- Step 8** Click **Submit**.
-

Deleting Quotas

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Quotas** tab.
 - Step 5** Choose the type of quota that you want to delete.
 - Step 6** Click **Delete**.
 - Step 7** In the **Delete Quota** dialog box, click **Submit**.
-

Snapshots

You can create a snapshot of an entire file system or of an independent file set to preserve the contents of the file system or the independent file set quota at a single point in time. Snapshots can still be created or deleted manually at any time. The storage space is required only to retain a copy of all of the data blocks that have been changed or deleted after the time of the snapshot, and is charged against the file system or the independent file set quota.

You cannot create a snapshot of a dependent file set.

Snapshot Rules

You can configure (create, change, or delete) the snapshot scheduling rules to schedule the creation and deletion of snapshots. You can specify the frequency to automatically take snapshots and retain these snapshots in the system. You can take snapshots hourly, daily, weekly, or monthly. The frequency in which the snapshots are taken can also be defined for specific days of the week, specific weeks of the month, or specific months. A snapshot rule can be associated with one or more file sets, and a file set can have one or more rules defined for it.

Creating Snapshot Rules

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Snapshot Rules** tab.
- Step 5** Click **Create**.
- Step 6** In the **Create Snapshot Rule** dialog box, complete the following fields:

Name	Description
Name field	Enter the snapshot rule name.
Frequency drop-down list	Choose the frequency to automatically create and retain snapshots.
Frequency Options	Specify the options according to the selection of the frequency.
Retention—For each past time period, enter the minimum number of snapshots to be kept for this snapshot rule.	
Hours field	Specify the hour.
Days field	Specify the day.
Weeks field	Specify the week.
Months field	Specify the month.
Prefix field	Specifies the prefix for the snapshots to be taken with this rule. The prefix is added to the date and time to identify the rule that created the snapshot. If the prefix is not specified, only the date and time is used. A prefix should be between 1 and 12 characters long and cannot contain any white space.

Step 7 Click **Submit**.

Creating Snapshots

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Snapshots** tab.
- Step 5** Click **Create**.
- Step 6** In the **Create Snapshot** dialog box, complete the following fields:

Name	Description
Snapshot Type drop-down field	Choose the snapshot type either to create manually or schedule the creation of snapshots.
Path field	Specify the file path.
Snapshot name	Enter the snapshot name. Use the format @GMT-dddd.MM.dd-HH.mm.ss if Windows Volume Shadow Copy Service is used to restore files. (For example, @GMT-2015.01.23-09.46.12)
Snapshot Schedules	Choose one or more snapshot schedules.

Step 7 Click **Submit**.

Deleting Snapshots

Step 1 On the menu bar, choose **Physical > Storage**.

Step 2 In the left pane, navigate to the pod in which the storage account is added.

Step 3 Choose the storage account type.

Step 4 Click the **Snapshots** tab.

Step 5 Click **Delete**.

Step 6 In the Snapshot Deletion dialog box, click **Delete**.

If the snapshots are used with NDMP (Network Data Management Protocol) backups, ensure that you deactivate and reactivate the NDMP configuration after you delete the snapshots.

Step 7 Click **Submit**.

Configuring Snapshot Notifications

Use this procedure to specify the snapshots events that you want them to be notified.

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Snapshots** tab.
- Step 5** Click **Configure** .
- Step 6** In the **Configure Snapshot** dialog box, complete the following fields:

Name	Description
Concurrent Rules check box	Check this check box to notify when another rule is in progress while a new rule is attempting to start.
Rule-created snapshot creation fails check box	Check this check box to notify when the creation of a snapshot that is invoked by a rule fails.
A snapshot deletion fails check box	Check this check box to notify when the deletion of a snapshot that is invoked by a rule fails.
The number of pending operations exceeds check box	Check this check box to notify when the number of pending operations for a rule exceed a specified value.
The total number of pending rule operations exceeds check box	Check this check box to notify when the total number of pending operations for a rule exceed a specified value.
A snapshot operation exceeds time limit (minutes)	Check this check box to notify when the duration of a snapshot exceeds the specified time limit.

- Step 7** Click **Submit**.



Managing Copy Services

This chapter contains the following sections:

- [Copy Services, page 65](#)
- [Consistency Groups, page 74](#)

Copy Services

Both IBM Storwize and IBM SAN Volume Controllers provide Copy Services functions that enable to copy volumes for all supported hosts that are connected to IBM Storwize and IBM SAN Volume Controllers.

FlashCopy

The FlashCopy function creates an instant, point-in-time copy from a source volume to a target volume, which associates the source volume and the target volume in a FlashCopy mapping.

- **Snapshot**—The snapshot preset creates a point-in-time view of the production data. The snapshot is not intended to be an independent copy but is used to maintain a view of the production data at the time the snapshot is created. Therefore, the snapshot holds only the data from regions of the production volume that changed since the snapshot was created. Since the snapshot preset uses thin provisioning, only the capacity that is required for the changes is used.
- **Clone**—The clone preset creates an exact replica of the volume, which can be changed without impacting the original volume. After the copy completes, the mapping that was created by the preset is automatically deleted.
- **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.

Remote Copy

Metro Mirror and Global Mirror are two types of remote-copy operations that enable you to set up a relationship between two volumes, where updates made to one volume are mirrored on the other volume. The volumes can be on either the same system or on two different systems. To use Metro Mirror and Global Mirror functions, you must have the remote-copy license installed on each enclosure that you plan to use these functions.

- The synchronous copy service (Metro Mirror)—Provides a consistent copy of a source volume on a target volume so that the copy is continuously updated. Data is written to the target volume synchronously after it is written to the source volume, both of which can belong to the same instance or different instances.
- The asynchronous copy service (Global Mirror)—Provides a consistent copy of a source volume on a target volume so that the copy is continuously updated, but the copy might not contain the most recent updates. Data is written to the target volume asynchronously after it is written to the source volume, both of which can belong to the same instance or different instances.

Creating FlashCopy Snapshot

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy** tab.
 - Step 5** Choose the volume that you want to create a snapshot.
 - Step 6** Click **Snapshot**.
 - Step 7** In the **Create Snapshot** dialog box, click **Submit**.
-

Creating FlashCopy Backup

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy** tab.
 - Step 5** Choose the volume that you want to create a backup.
 - Step 6** Click **Backup**.
 - Step 7** In the **Create Backup** dialog box, click **Submit**.
-

Creating FlashCopy Clone

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy** tab.
 - Step 5** Choose the volume that you want to clone.
 - Step 6** Click **Clone**.
 - Step 7** In the **Create Clone** dialog box, click **Submit**.
-

Deleting FlashCopy

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy** tab.
 - Step 5** Choose the volume that you want to delete.
 - Step 6** Click **Delete**.
 - Step 7** Check the **Delete the Volume even it has host mappings or is used in FlashCopy mappings or remote copy relationships** check box, if required .
 - Step 8** In the **Delete** dialog box, click **Submit**.
-

Starting FlashCopy Services

Before You Begin

Create FlashCopy consistency group and FlashCopy mappings using one of the FlashCopy presets (Snapshot, Clone, or Backup).

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy Consistency Groups** tab.
 - Step 5** Choose the FlashCopy Consistency Group, which is empty.
 - Step 6** Click **Start**.
 - Step 7** In the **Start Consistency Group** dialog box, click **Submit**.
-

Stopping FlashCopy Services

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy Consistency Groups** tab.
 - Step 5** Choose the FlashCopy Consistency Group.
 - Step 6** Click **Stop**.
 - Step 7** In the **Stop Consistency Group** dialog box, click **Submit**.
-

Creating Remote Copy Relationships

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type (IBM Storwize or IBM SAN Volume Controller).
 - Step 4** Click the **Remote Copy** tab.
 - Step 5** Click **Create**.
 - Step 6** In the **Create Remote Copy Relationships** dialog box, complete the following fields:

Name	Description
Copy Type drop-down list	Choose the remote copy type.
Name field	Enter the name for the selected remote copy type.
Master Volume field	Choose the master volume.
Auxiliary volume field	Choose the secondary volume.
Sync Status drop-down list	Choose if the volumes are already synchronization or not.
Copy status drop-down list	Choose if you want the volume to be copied or not.

Step 7 Click **Submit**.

Editing Remote Copy Relationships

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Remote Copy** tab.
 - Step 5** Choose the remote copy type that you want to edit.
 - Step 6** Click **Edit**.
 - Step 7** In the **Edit Remote Copy Relationship** dialog box, choose the remote copy type that you want to change to.
 - Step 8** For the Global Mirror type, choose the cyclic mode.
 - Step 9** For the Global Mirror type, enter the cyclic period in seconds.
 - Step 10** Click **Submit**.
-

Renaming Remote Copy Relationships

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Remote Copy** tab.
 - Step 5** Choose the remote copy relationship that you want to rename.
 - Step 6** Click **Rename**.
 - Step 7** In the **Rename Remote Copy Relationship** dialog box, enter a new relationship name for the remote copy.
 - Step 8** Click **Submit**.
-

Deleting Remote Copy Relationships

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Remote Copy** tab.
 - Step 5** Choose the remote copy relationship that you want to delete.
 - Step 6** Click **Delete**.
 - Step 7** In the **Delete Remote Copy Relationship** dialog box, check the **Force Delete** check box to delete the relationship even when the data in the target volume is inconsistent, or if the target volume has other dependencies.
 - Step 8** Click **Submit**.
-

Starting Remote Copy Relationships

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Remote Copy** tab.
 - Step 5** Choose the remote copy relationship that you want to start.
 - Step 6** Click **Start**.
 - Step 7** In the **Start Remote Copy Relationship** dialog box, click **Submit**.
-

Stopping Remote Copy Relationships

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Remote Copy** tab.
 - Step 5** Choose the remote copy relationship that you want to stop.
 - Step 6** Click **Stop**.
 - Step 7** In the **Stop Remote Copy Relationship** dialog box, click **Submit**.
-

Switching Remote Copy Relationships

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **Remote Copy** tab.
 - Step 5** Choose the remote copy relationship that you want to switch the write access and primary volume between master auxiliary volumes.
 - Step 6** Click **Switch**.
 - Step 7** In the **Switch Remote Copy Relationship** dialog box, Click **Submit**.
-

Adding Remote Copy Relationships to Remote Copy Consistent Groups

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Remote Copy** tab.
- Step 5** Click **Add**.
- Step 6** In the **Add to Consistency Group** dialog box, click **Select**.
- Step 7** Choose the consistent group to which you want to add the remote copy relationship with the consistent group.
- Step 8** Click **Select**.
- Step 9** Click **Submit**.
-

Removing Remote Copy Relationships from Remote Copy Consistent Groups

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **Remote Copy** tab.
- Step 5** Choose the remote copy relationship that you want to remove from the consistency group.
- Step 6** Click **Remove**.
- Step 7** In the **Remove Remote Copy Relationship** dialog box, click **Submit**.
-

File Copy Services

The IBM Storwize V7000 Unified and the IBM SAN Volume Controller (SVCs) systems support the following methods to replicate data to and between different file systems:

- 1 Replicate file system—File system replication provides asynchronous replication of the file systems on one system to another system located remotely over an IP network. Replicating a file system creates copies of a file system between two systems, separated by some geographical distance to provide disaster recovery and business continuity. Asynchronous replication is normally used between the source and target systems where distance might affect the response time because of bandwidth shortages.

You can use the current system as the source or target for replication.

- if you want to use the current system as the source system for replication, the source system requires the management/admin IP address of the target system's management node and node pair definitions.

Before entering these source settings, you must first go to the target system and specify the target settings.

- If you want to use the current system as the target system for replication, the current (target) system requires the management/admin IP address of the source system's management node to establish the connection with the source system. You must go to the source system and specify the source settings.
- 2 Remote caching—Remote caching provides transparent data distribution among data centers and multiple remote locations over a wide area network (WAN). Remote caching fetches files on demand from the home system file set to the remote cluster cache in real time during normal operations. The data created, maintained, updated, and changed on the home system can be viewed and used on a cache system located anywhere in the WAN. Users can see only the files they have permission to see and access the files as if they are present locally. The data is revalidated when the user tries to access it, to ensure that the most recent file versions are available that can be prefetched by setting policies or the files can be pulled into cache on demand.

Creating a Replication Destination

Before You Begin

Configure the asynchronous replication relationship between the two (source and target) systems before configuring file system replication.

-
- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **File Systems** tab.
 - Step 5** Choose the file system for which you want to create a replication destination.
 - Step 6** Click **Replication Destination**.
 - Step 7** In the **Replication Destination** dialog box, enter the ID value of the source cluster (Source cluster ID) that uses the local path for replication.
 - Step 8** Click **Submit**.
-

Removing a Replication Destination

-
- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **File Systems** tab.
- Step 5** Choose the file system for which you want to remove the replication destination at the specified path.
- Step 6** Click **Replication Destination**.
- Step 7** In the **Replication Destination** dialog box, click **Submit**.
-

Consistency Groups

A consistency group is a container for FlashCopy mappings, Global Mirror relationships, and Metro Mirror relationships. You can add many mappings or relationships to a consistency group, however FlashCopy mappings, Global Mirror relationships, and Metro Mirror relationships cannot appear in the same consistency group. When you use a consistency group, you can perform copy operations on the entire group instead of the individual mappings or relationships.

- FlashCopy consistency groups—You can group FlashCopy mappings into a consistency group so that they can be updated at the same time.
- Metro Mirror and Global Mirror consistency groups—You can group Metro Mirror or Global Mirror relationships into a consistency group so that they can be updated at the same time.

Creating FlashCopy and Remote Copy Consistency Groups

-
- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **FlashCopy Consistency Groups** tab or the **Remote Copy Consistency Groups** tab.
- Step 5** Click **Create**.
- Step 6** In the **Create this Consistency Group** dialog box, enter the name for the consistency group.
- Step 7** Click **Submit**.
-

Renaming FlashCopy and Remote Copy Consistency Groups

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy Consistency Groups** tab or the **Remote Copy Consistency Groups** tab.
 - Step 5** Choose the consistency group that you want to rename.
 - Step 6** Click **Rename**.
 - Step 7** In the **Rename Consistency Group** dialog box, enter the new name for the consistency group.
 - Step 8** Click **Submit**.
-

Deleting FlashCopy and Remote Copy Consistency Groups

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy Consistency Groups** tab or the **Remote Copy Consistency Groups** tab.
 - Step 5** Choose the consistency group that you want to delete.
 - Step 6** Click **Delete**.
 - Step 7** In the **Delete this Consistency Group** dialog box, click **Submit**.
-



FlashCopy Mappings

This chapter contains the following sections:

- [FlashCopy Mappings, page 77](#)
- [FlashCopy Mapping States, page 78](#)
- [Creating FlashCopy Mappings, page 78](#)
- [Moving FlashCopy Mapping to FlashCopy Consistency Group, page 81](#)
- [Starting FlashCopy Mapping, page 81](#)
- [Stopping FlashCopy Mapping, page 82](#)
- [Deleting FlashCopy Mapping, page 82](#)

FlashCopy Mappings

A FlashCopy mapping defines the relationship between a source volume and a target volume that can be created between any two volumes in a system. The source and target volumes must be of equal size and the volumes need not be in the same I/O group or pool.

The FlashCopy function creates an instant copy of a volume. To create an instant copy of a volume, you must first create a mapping between the source volume (the disk that is copied) and the target volume (the disk that receives the copy).

When a FlashCopy operation starts, no data is actually copied at the time, but a checkpoint is made of the source volume. The checkpoint creates a bitmap that indicates that no part of the source volume is copied. Each bit in the bitmap represents one region (a grain) of the source volume. If new data is written to the source or target volume, the existing data on the source is copied to the target volume before the new data is written to the source or target volume.

After a FlashCopy operation starts, the read operations to the source volume continue to occur. The bitmap is updated to mark that the grain of the source volume is copied so that the write operations to the same grain do not re-copy the data later. During a read operation to the target volume, the bitmap is used to determine if the grain is copied. If the grain is already copied, the data is read from the target volume. If the grain is not copied, the data is read from the source volume.

FlashCopy Mapping States

At any point in time, a FlashCopy mapping is in one of the following states:

- 1 **Idle or Copied**—In this state, the source and target volumes act as independent volumes even if a mapping exists between the two. If the mapping is incremental, and the background copy is complete, the mapping records only the differences between the source and target volumes. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target volumes will be offline. Read and write caching is enabled for both the source and the target volumes.
- 2 **Copying**—In this state, the copy is in progress. Read and write caching is enabled on the source and the target volumes.
- 3 **Stopped**— In this state, 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 records the 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.

Creating FlashCopy Mappings

- Step 1** On the menu bar, choose **Physical > Storage**.
- Step 2** In the left pane, navigate to the pod in which the storage account is added.
- Step 3** Choose the storage account type.
- Step 4** Click the **FlashCopy Mappings** tab.
- Step 5** Click **Create**.
- Step 6** In the **Create FlashCopy Mapping** dialog box, complete the following fields:

Name	Description
Name field	Enter the name for the FlashCopy Mapping.
Master Volume drop-down list	Choose the master volume.
Auxili Auxiliary Volume drop-down list	Choose the auxiliary volume that has the same size as the master volume that was previously selected.
Select a Preset check box	Check this check box to create a FlashCopy mapping for the FlashCopy types. Note If you do not check the Select a Preset check box, you can proceed with the Advanced Setting.

Name	Description
Type drop-down list	Choose one of the following FlashCopy types: <ul style="list-style-type: none"> • Snap shot • Clone • Backup
Advanced Setting	
Background Copy Rate field	Specify the copy rate to determine the priority that is given to the copy process. Note A faster rate increases the priority of the copy process, which might affect the performance of other operations.
Incremental check box	Check this check box to copy the data that is modified since the initial FlashCopy mapping.
Delete mapping after completion check box	Check this check box to automatically delete the mapping when the mapping reaches the idle_or_copied state and the progress is 100%.
Clearing Rate field	Specify the clean rate to minimize the amount of time if a mapping is in the stopping state. Note If the mapping has not completed, the target volume is offline while the mapping is stopping.

Step 7 Click **Submit**.

FlashCopy Advanced Settings

Incremental FlashCopy Mappings

In an incremental FlashCopy mapping, the initial mapping copies all of the data from the source volume to the target volume. Subsequent FlashCopy mappings copy only the data that is modified since the initial FlashCopy mapping.



Note

You can define a FlashCopy mapping as incremental only when you create the FlashCopy mapping.

Background Copy Rate

You can specify a copy rate when you create a FlashCopy mapping. When the mapping is in the copying state, the copy rate determines the priority that is given to the background copy process. If you want a copy of the whole source volume so that a mapping can be deleted and still be accessed from the target volume, you must copy all the data that is on the source volume to the target volume.

When a mapping is started and the copy rate is greater than zero, the unchanged data is copied to the target volume, and the bitmap is updated to show that the copy has occurred. After a certain period of time, the length of which depends on the priority that is determined by the copy rate and the size of the volume, the whole volume (source) is copied to the target. The mapping returns to the `idle_or_copied` state and you can now restart the mapping at any time to create a new copy at the target.

While the mapping is in the copying state, you can set the copy rate to zero and the clean rate to a value other than zero to minimize the amount of time a mapping is in the stopping state. If the copy rate is zero, only the data that changes on the source is copied to the target. The target never contains a copy of the whole source unless every extent is overwritten at the source. You can use this copy rate when you require a temporary copy of the source.

Clean Rate

You can specify a clean rate when you create a FlashCopy mapping. The clean rate is used to control the rate where the data is copied from the target volume of the mapping to the target volume of a mapping that is either the latest copy of the target volume, or is the next oldest copy of the source volume.

You can use the clean rate in the following scenarios:

- The mapping is in the stopping state.
- The mapping is in the copying state and has a copy rate of zero.
- The mapping is in the copying state and the background copy has completed.

You can use the clean rate to minimize the amount of time that a mapping is in the stopping state. If the mapping has not completed, the target volume is offline while the mapping is in the stopping state. The target volume remains offline until the mapping is restarted. You can stop the mapping at any time after it has been started. If the target volume already contains a complete copy of the source volume, this action makes the target inconsistent and the target volume goes offline. The target volume remains offline until the mapping is restarted.

Moving FlashCopy Mapping to FlashCopy Consistency Group

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy Mappings** tab.
 - Step 5** Choose the FlashCopy mapping, if the mapping is in one of the following states: stopped, idle, or copied.
 - Step 6** Click **Move To Group**.
 - Step 7** In the **Move FlashCopy Mapping to Consistency Group** dialog box, click **Select**.
 - Step 8** Check the check box to which FlashCopy Consistency Group you want to move the FlashCopy mapping.
 - Step 9** Click **Submit**.
-

Starting FlashCopy Mapping

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy Mappings** tab.
 - Step 5** Choose the FlashCopy mapping, if the mapping is in one of the following states: stopped, idle, or copied.
 - Step 6** Click **Start**.
 - Step 7** In the **Start FlashCopy Mapping** dialog box, click **Submit**.
-

Stopping FlashCopy Mapping

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy Mappings** tab.
 - Step 5** Choose the FlashCopy mapping, if the mapping is in the copying state.
 - Step 6** Click **Stop**.
 - Step 7** In the **Stop FlashCopy Mapping** dialog box, click **Submit**.
-

Deleting FlashCopy Mapping

- Step 1** On the menu bar, choose **Physical > Storage**.
 - Step 2** In the left pane, navigate to the pod in which the storage account is added.
 - Step 3** Choose the storage account type.
 - Step 4** Click the **FlashCopy Mappings** tab.
 - Step 5** Choose the FlashCopy mapping that you want to delete.
 - Step 6** Click **Delete**.
You can delete the FlashCopy mapping even when the data on the target volume is inconsistent or if the target volume has other dependencies.
 - Step 7** Check the **Force Delete** check box, if required.
 - Step 8** In the **Delete FlashCopy Mapping** dialog box, click **Submit**.
-



Workflows for IBM Storwize System

This chapter contains the following sections:

- [Workflows for IBM Storwize and IBM SVC Storage Accounts](#), page 83

Workflows for IBM Storwize and IBM SVC Storage Accounts

Cisco UCS Director defines a set of workflows in the **IBM Storwize Tasks** folder for IBM Storwize and IBM SAN Volume Controllers under the **Workflows** tab. On the menu bar, choose **Policies > Orchestration** and launch the Workflow Designer. In the left pane, choose the **Physical Storage Tasks** folder from the **Available Tasks** pane.

The IBM Storwize Data Store Tasks folder contains the following workflow tasks:

- **IBM Storwize Add NFS Datastore**—This task adds NFS datastores in the path specified.
- **IBM Storwize Block Volume Datastore**—This task creates datastores for block volume and mounts it into zoned host of the vCenter.
- **IBM Storwize Compressed Volume Datastore**—This task creates datastores for compressed volume and mounts it into zoned host of the vCenter.
- **IBM Storwize Thin Provisioned Volume Datastore**—This task creates datastores for thin provisioned volume and mounts it into zoned host of the vCenter.
- **IBM Storwize iSCSI Boot Target**—This task configures iSCSI boot targets.

The following folders contain tasks for actions that you can perform on managed IBM Storwize and IBM SAN Volume Controller features:

- **IBM Storwize Custom SSH Command**—This task allows you to run IBM Storwize CLI on IBM Storwize devices.
- **IBM Storwize FC Consistency Groups Tasks**—This folder includes tasks that can be used to create, rename, and delete fibre channel consistency groups.
- **IBM Storwize FlashCopy Mapping Tasks**—This folder includes tasks that can be used to create, delete, start, and stop FlashCopy mappings and move mappings to FlashCopy consistency groups.

- IBM Storwize File Sets Tasks—This folder includes tasks that can be used to create a new file set and delete file sets.
- IBM Storwize File Systems Tasks—This folder includes tasks that can be used to create a file system of single pool, compressed, and migration-ILM types, replicate a file system, delete a file system, mount, and unmount file systems.
- IBM Storwize File Shares Tasks—This folder includes tasks that can be used to create, delete file shares, and activate or deactivate file shares.
- IBM Storwize FlashCopy Tasks—This folder includes tasks that can be used to create backup, clone, and snapshot of a FlashCopy and delete a FlashCopy.
- IBM Storwize Host Tasks—This folder includes tasks that can be used to create, rename, delete an IBM Storwize host, unmap all volumes from the host, duplicate host mappings, and import host mappings.
- IBM Storage MDisk Tasks—This folder includes tasks that can be used to rename MDisks, delete an MDisk, assign an MDisk to a storage pool, unassign an MDisk from the storage pool, set spare goal, and set tier.
- IBM Storwize (Storage) Pool tasks—This folder includes tasks that can be used to create, rename, and delete storage pools.
- IBM Storwize Remote Copy Tasks—This folder includes tasks that can be used to create, delete, start, and stop remote copy relationships.
- IBM Storwize Snapshots Tasks—This folder includes tasks that can be used to create and delete snapshots.
- IBM Storwize Snapshot Rules Tasks—This folder includes tasks that can be used to create and delete snapshot rules.
- IBM Storwize Volume Tasks—This folder includes tasks that can be used to create, modify, rename, delete, duplicate, expand, shrink, migrate, split volumes, map volumes to a host, unmap volumes from the host, make volumes as primary, and add mirror copy of volumes.