

IBM System Storage N series

SnapDrive 7.0 for Windows Administration Guide for SAN Environments

Contents

Preface	
Supported features	12
Websites	12
Getting information, help, and service	12
Before you call	
Using the documentation	13
Hardware service and support	13
Firmware updates	13
How to send your comments	14
SnapDrive overview	
What SnapDrive does	15
List of current SnapDrive limitations	16
Recommendations for using SnapDrive	16
Understanding your SnapDrive components	
Understanding the Volume Shadow Copy Service	19
Understanding VSS	19
SnapDrive VSS requirements	20
Typical VSS backup process	20
Troubleshooting the VSS Hardware Provider	21
Establishing a connection to the storage system	
Managing iSCSI sessions	
iSCSI Software Initiator node naming requirements	24
Establishing an iSCSI session to a target	
Disconnecting an iSCSI target from a Windows host	
Disconnecting a session to an iSCSI target	27
Examining details of an iSCSI session	27
SnapDrive support for ESX iSCSI initiators	
ESX iSCSI initiator limitations	
Enabling storage system HTTP communication with SnapDrive	29
Using SnapDrive in Microsoft environments	
Understanding the new features of Windows Server 2012	30
CSV 2.0 in Windows Server 2012	

Hyper-V VSS backup changes with Windows Server 2012	30
SnapDrive limitations on Windows Server 2012	31
Data ONTAP DSM support for Windows MPIO	31
Cluster support	32
Support for Microsoft Cluster Shared Volumes	32
Verifying the cluster group owner	32
Changing the cluster group owner	33
Troubleshooting CSVs	33
Creating LUNs	34
Rules for creating LUNs	34
What volume mount points are	35
Creating a dedicated LUN	35
Creating a shared LUN	38
Windows Server 2008 and 2012 failover cluster support	42
Configuring a Windows Server 2008 and 2012 failover cluster witness	
disk	42
Creating a highly available Hyper-V virtual machine using SnapDrive	44
Support for creating disks on a virtual storage server operating in	
clustered Data ONTAP	45
GPT partition support	45
Configuring SnapDrive to create LUNs with Virtual Fibre Channel	46
Managing LUNs	46
How LUNs work	46
List of guidelines for connecting LUNs	47
Connecting to a LUN	48
Making drive letter or path modifications to a LUN	50
Guidelines for disconnecting or deleting LUNs	51
Disconnecting a LUN	52
Deleting a LUN	53
Deleting folders within volume mount points	54
Guidelines for resizing disks	54
Resizing a disk	55
Resizing a quorum disk	56
Managing LUNs not created in SnapDrive	56
Requirements for dynamically adding and removing pass-through disks	
on Hyper-V virtual machines	58

Managing Snapshot copies	60
Reasons for creating Snapshot copies	60
Restrictions on Snapshot copy creation	60
Creating a Snapshot copy	61
Scheduling Snapshot copies	62
Support for FlexClone volumes in SnapDrive	63
Snapshot copy cautions	63
Connecting to a LUN in a Snapshot copy	64
Data protection through archiving and restoring Snapshot backup copies .	65
Deleting a Snapshot copy	69
Managing space on storage system volumes	70
What SnapDrive fractional space reservation monitoring does	70
Configuring space reservation monitoring	70
Using the storage access control tool to enable thinly provisioned LUNs .	71
What Space Reclaimer does	71
Using SnapDrive in VMware environments	76
VMware support	76
VMware ESX server-related limitations	76
Enabling and disabling vCenter or ESX logon from SnapDrive MMC	76
Minimum vCenter privileges required for SnapDrive operations	77
Requirements for VMware vMotion support	77
Creating LUNs in VMware environments	78
Creating an RDM LUN on a guest OS	78
Using FC RDM LUNs in a Microsoft cluster	82
Managing LUNs in VMware environments	83
Connecting to an RDM LUN on a guest OS	84
Guideline for managing RDM LUNs not created in SnapDrive	86
Managing Snapshot backups in VMware environments	86
Support requirements for performing Snapshot copy operations in	
VMDKs on NFS and VMFS datastores	86
Snapshot copy support limitations on VMDKs	86
Troubleshooting VMDKs	87
Support requirements for space reclamation in VMDK files in NFS datastores	88
Additional Microsoft clusters on ESX documentation resources	89
Performing SnapVault and SnapMirror operations	90
Using SnapVault with SnapDrive	90

Considerations when using SnapVault	90
Initiating SnapVault backup jobs from SnapDrive in 7-Mode SAN	
environments	91
SnapVault operations supported in a clustered Data ONTAP	91
Using SnapMirror with SnapDrive for Windows	92
SnapMirror overview	92
Types of SnapMirror replication	92
Requirements for using SnapMirror with SnapDrive	94
Initiating replication manually	95
Connecting to a LUN in a mirrored destination volume	96
Restoring a volume on a SnapMirror destination	96
Recovering a cluster from shared LUNs on a SnapMirror destination	97
SnapDrive integration with OnCommand Unified Manager Core Package data	
protection capabilities	101
How SnapDrive integrates with OnCommand Unified Manager Core	
Package data protection capabilities	101
Dataset concepts	101
Configuring and managing access control	103
Support for storage system access control	103
Using storage system access control	103
Storage system access control reference	104
Enabling RBAC for use with SnapDrive	109
Using RBAC with the OnCommand Unified Manager Core Package	
server	110
Enabling RBAC on the storage system	111
Configuring SnapDrive for Windows to use RBAC	112
Creating SnapDrive user roles on DataFabric Manager server	112
Assigning roles to SnapDrive users on DataFabric Manager server	113
SnapDrive for Windows to DataFabric Manager role mappings	113
SnapDrive command-line reference	117
About sdcli commands	117
Executing sdcli commands	117
Common command switches	118
Configuration commands	120
The sysconfig list command	120
Dataset management commands	120

	dataset add_members	120
	dataset backup_add	121
	dataset backup_change_retention_type	121
	dataset backup_delete	122
	dataset backup_end	123
	dataset backup_get_metadata	123
	dataset backup_list	124
	dataset backup_set_metadata	124
	dataset backup_start	125
	dataset backup_status	125
	dataset backup_version_convert	126
	dataset create	126
	dataset create_local_backup	127
	dataset delete	127
	dataset dfm_request	128
	dataset get_available_policies	128
	dataset get_backup_location	129
	dataset get_backup_version_info	129
	dataset get_metadata	130
	dataset get_policy	130
	dataset get_retention_info	130
	dataset info	131
	dataset initiate_conformance	131
	dataset list_members	132
	dataset mount_backup	132
	dataset protect	133
	dataset remove_members	133
	dataset restore	134
	dataset restore_status	135
	dataset set_metadata	135
	dataset set_policy	136
	dataset transfer_now	136
	dataset vss_backup_end	137
	dataset vss_backup_prepare	137
Licens	e commands	138
	The license set command	138

The license list command	138
The license remove command	138
Initiator group management commands	139
The igroup list command	139
The igroup create command	139
The igroup rename command	140
The igroup delete command	140
Fractional space reservation monitoring commands	141
The spacemon list command	141
The spacemon set command	141
The spacemon snap_delta command	142
The spacemon snap_reclaimable command	142
The spacemon vol_info command	142
The spacemon delete command	143
Virtual Storage Console commands	143
The vsc_config list command	143
The vsc_config set command	143
The vsc_config delete command	143
Space reclamation commands	144
The spacereclaimer start command	144
The spacereclaimer stop command	144
The spacereclaimer analyze command	145
The spacereclaimer status command	145
Preferred IP address commands	145
The preferredIP set command	145
The preferredIP list command	146
The preferredIP delete command	146
iSCSI connection commands	146
The iscsi_target disconnect command	146
The iscsi_target list command	146
iSCSI initiator commands	147
The iscsi_initiator list command	147
The iscsi_initiator establish_session command	147
The iscsi_initiator terminate_session command	148
LUN commands	148
The disk create command	148

The disk connect command	150
The disk delete command	151
The disk disconnect command	151
The disk resize command	152
The disk expand command	153
disk add_initiator	154
disk remove_initiator	154
The disk list command	155
The disk add_mount command	156
The disk remove_mount command	156
The disk rename_flexclone command	156
Snapshot copy commands	157
The snap create command	157
The snap delete command	158
The snap list command	158
The snap mirror_list command	158
The snap mount command	159
The snap rename command	159
The snap restore command	160
The snap unmount command	161
The snap update_mirror command	162
The snap restore_volume_check command	162
The snap restore_volume command	162
SnapVault commands	163
The snapvault verify_configuration command	163
The snapvault snapshot_rename command	163
The snapvault snapshot_delete command	164
The snapvault archive command	164
The snapvault relationship_status command	164
The snapvault snap_list command	165
OnCommand commands	165
The oncommand_config list command	165
The oncommand_config set command	165
The oncommand_config delete command	166
The oncommand_config rbaccache command	166
Transport protocol commands	166

The transport_protocol list command	166
The transport_protocol set command	167
The transport_protocol delete command	168
Virtual server commands	168
The vsconfig list command	168
The vsconfig set command	168
The vsconfig dslist command	169
The vsconfig delete command	169
Hyper-V configuration commands	169
The hyperv_config list command	169
The hyperv_config set command	169
The hyperv_config delete command	170
Typical SnapDrive configurations	171
SnapDrive iSCSI configurations	171
Single host direct-attached to a single storage system using iSCSI	171
Single host attached to a single storage system through a GbE switch	171
Single host attached to a single storage system through a dedicated	
switch	172
Windows cluster connected to a storage system cluster through a	
dedicated GbE switch	173
SnapDrive FC configurations	173
Single host direct-attached to a single storage system using FC	173
Single host attached to a single storage system through an FC switch	174
Windows cluster attached to a storage system active/active	
configuration through an FC switch	174
SnapDrive MPIO configurations	175
Single host direct-attached to a single storage system using MPIO	175
Windows cluster attached to a storage system active/active	
configuration through a GbE switch using MPIO	176
Windows cluster attached to a storage system active/active	
configuration through an FC switch using MPIO	177
SAN booting with SnapDrive	177
What SAN booting is	177
How SnapDrive supports SAN booting	178
Copyright information	179
Trademark information	180

Index	33
-------	----

Preface

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in *Websites* on page 12).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

• A listing of currently available N series products and features can be found at the following web page:

www.ibm.com/storage/nas/

• The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

www.ibm.com/storage/support/nseries/

This web page also provides links to AutoSupport information as well as other important N series product resources.

• IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

www.ibm.com/systems/storage/network/interophome.html

• For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains

information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in *Websites* on page 12) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in *Websites* on page 12).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in *Websites* on page 12).

14 | SnapDrive 7.0 for Windows Administration Guide for SAN Environments

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to *starpubs@us.ibm.com*.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

SnapDrive overview

SnapDrive for Windows enables you to automate storage provisioning tasks and to manage data in Microsoft Windows environments. You can run SnapDrive on Windows hosts in either a physical or virtual environment.

What SnapDrive does

SnapDrive helps you automate storage provisioning tasks and manage data in SAN and SMB 3.0 Windows environments. You can run SnapDrive software on Windows hosts in either a physical or a virtual environment.

SnapDrive software integrates with Windows Volume Manager so that storage systems can serve as virtual storage devices for application data in Windows Server 2008 R2 and Windows Server 2012. It can also be used to provision storage for Windows virtual machines hosted on ESX hypervisors.

SnapDrive manages LUNs on a storage system, making these LUNs available as local disks on Windows hosts. This allows Windows hosts to interact with the LUNs just as if they belonged to a directly attached redundant array of independent disks (RAID).

SnapDrive PowerShell cmdlets support volume and share provisioning in SMB 3.0 environments. You can also use SnapDrive PowerShell cmdlets to create and manage Snapshot backups; manage mounting, restore, and discovery operations; and to troubleshoot.

SnapDrive provides the following additional features:

- · It enables online storage configuration, LUN expansion, and streamlined management.
- It enables connection of up to 168 LUNs.
- It integrates Data ONTAP Snapshot technology, which creates point-in-time images of data stored on LUNs.
- It works in conjunction with SnapMirror software to facilitate disaster recovery from either asynchronously or synchronously mirrored destination volumes.
- It enables SnapVault updates of qtrees to a SnapVault destination.
- It enables management of SnapDrive on multiple hosts.
- It enables support on Microsoft cluster configurations.
- It enables iSCSI session management.

List of current SnapDrive limitations

Some functionality is currently not supported in SnapDrive.

- A LUN managed by SnapDrive cannot be configured as a "dynamic" disk (a storage device that is divided into volumes rather than partitions); it can serve only as a "basic" disk (a storage device for host-side application data).
- A LUN cannot be configured as an extended partition. SnapDrive supports only a single, primary partition on a LUN.
- LUNs created in OnCommand System Manager or at the storage system command line cannot be managed unless certain steps are taken to prepare these disks for SnapDrive.
- SnapDrive supports LUNs on qtrees, but you cannot manage quotas from SnapDrive. LUNs can be created within a qtree and quota limits for that qtree are enforced; therefore, you cannot create a LUN or expand an existing LUN beyond the quota limit set for that qtree.
- SnapDrive supports the use of SnapMirror to replicate volumes but not individual qtrees.
- SnapDrive does not support the creation of CSVs with a drive letter or a volume mount point on Windows Server 2008 R2 and Windows Server 2012.
- When you have made Snapshot copies from SIS clones, you cannot make new Snapshot copies until you have completed background sharing work.
 SIS clones made from Snapshot copies prevents creation of new Snapshot copies until background sharing work is completed.
- When you are operating in a clustered Data ONTAP environment, SnapDrive must be installed on all nodes in the cluster.

Recommendations for using SnapDrive

Follow these recommendations whenever you use SnapDrive for Windows.

- Use SnapDrive to create and manage all the LUNs on your storage system.
- Use the OnCommand System Manager to set up and configure your storage system and to provision volumes for use with SnapDrive.
- If you want to dedicate all free space on a volume to LUNs, set the snap reserve setting on the storage system to 0 percent.
- Place all LUNs connected to the same host on a dedicated volume accessible by just that host.
- Unless you can be sure that name resolution publishes only the storage system interface you intend, configure each network interface by IP address, rather than by name.
- If you use Snapshot copies, you cannot use the entire space on a storage system volume to store your LUN.

The storage system volume hosting the LUN should be the size of all the LUNs on the volume, with enough additional space for the Snapshot copies of the volume. The additional space should be based on the change rate of the LUNs in the volume and the retention policy for the Snapshot copies.

• Do not create any LUNs in /vol/vol0. This is a storage system limitation. This volume is used by Data ONTAP to administer the storage system and should not be used to contain any LUNs.

Understanding your SnapDrive components

Several components are integrated into the SnapDrive software and are automatically installed. These components enable you to manage LUNs, Windows volumes, or SMB shares. You can use these components together to enable SnapDrive workflows, including provisioning, Snapshot copy management, backup, restore, and mounting operations.

The following SnapDrive components are integrated in the software and are automatically installed during installation.

SnapDrive "snap-in"

This software module integrates with Microsoft Management Console (MMC) to provide a graphical interface for managing LUNs on the storage system. The module does the following:

- · Resides in the Windows Server computer management storage tree
- Provides a native MMC snap-in user interface for configuring and managing LUNs
- · Supports remote administration so that you can manage SnapDrive on multiple hosts
- · Provides SnapMirror integration
- · Provides AutoSupport integration, including event notification

SnapDrive command-line interface

The sdcli.exe utility enables you to manage LUNs from the command prompt of the Windows host. You can perform the following tasks with the sdcli.exe utility:

- Enter individual commands
- Run management scripts

PowerShell cmdlets

The SnapDrive PowerShell cmdlets enable you to perform provisioning, Snapshot copy management and backup, restore, and mounting operations in an SMB 3.0 environment.

SnapDrive supports PowerShell versions 2.0 and later.

Underlying SnapDrive service

This software interacts with software on the storage system to facilitate LUN management for the following:

- A host
- · Applications running on a host

Data ONTAP Volume Shadow Copy Service (VSS) Hardware Provider on Windows Server hosts

The Data ONTAP VSS Hardware Provider is a module of the Microsoft VSS framework. The Data ONTAP Hardware Provider enables VSS Snapshot technology on the storage system when SnapDrive is installed on Windows Server hosts.

Understanding the Volume Shadow Copy Service

The Data ONTAP VSS Hardware Provider is installed with SnapDrive for Windows and can be used with Microsoft Volume Shadow Copy Service.

Understanding VSS

Volume Shadow Copy Service (VSS) is a feature of Microsoft Windows Server that coordinates data servers, backup applications, and storage management software to support the creation and management of consistent backups.

VSS coordinates Snapshot copy-based backup and restore operations and includes these components:

VSS Requestor	The VSS requestor is a backup application, such as SnapManager for Microsoft Exchange or NTBackup. It initiates VSS backup and restore operations. The requestor also specifies Snapshot copy attributes for backups it initiates.
VSS Writer	The VSS writer owns and manages the data to be captured in the Snapshot copy. Microsoft Exchange 2003 is an example of a VSS writer.
VSS provider	The VSS provider is responsible for creating and managing the Snapshot copy. A provider can be either a hardware provider or a software provider:
	 A hardware provider integrates storage array-specific Snapshot copy and cloning functionality into the VSS framework. The Data ONTAP VSS Hardware Provider integrates the SnapDrive service and storage systems running Data ONTAP into the VSS framework.
	Note: The Data ONTAP VSS Hardware Provider is installed automatically as part of the SnapDrive software installation.
	• A software provider implements Snapshot copy or cloning functionality in software that is running on the Windows system.
	Note: To ensure that the Data ONTAP VSS Hardware Provider works properly, do not use the VSS software provider on Data ONTAP LUNs. If you use the VSS software provider to create Snapshot copies on a Data ONTAP LUN, you cannot delete that LUN by using the VSS Hardware Provider.

SnapDrive VSS requirements

To use VSS with SnapDrive for Windows, your storage system and SnapDrive host must meet minimum requirements.

• Your storage system must be running at least Data ONTAP 7.2.

Note: In versions of Data ONTAP prior to 7.3, Snapshot copies taken after a shadow copy are locked due to the existence of LUN clones in the previous Snapshot copies, making them impossible to delete. In Data ONTAP 7.3, this restriction is removed, so SnapDrive is able to delete any Snapshot copies.

• The Virtual Disk Service must be running on your Windows host.

Typical VSS backup process

A typical backup using SnapManager for Microsoft Exchange, Microsoft Exchange 2013, and the Data ONTAP VSS Hardware Provider is outlined in the following process.

- 1. SnapManager determines which LUNs it wants to capture and verifies that Exchange 2013 is present as a valid writer.
- 2. SnapManager initiates the shadow copy process.
- **3.** VSS informs Exchange and the Data ONTAP VSS Hardware Provider that a shadow copy is starting. Exchange stops writing to disk.
- 4. VSS ensures that NTFS is in a consistent state.
- 5. VSS requests the Data ONTAP VSS Hardware Provider to create a shadow copy.
- 6. The Data ONTAP VSS Hardware Provider requests SnapDrive to create a Snapshot copy of the storage system volume that contains the specified LUN.
- 7. SnapDrive requests that the storage system create a Snapshot copy of the specified volume.
- **8.** When the shadow copy is complete, VSS returns NTFS to a normal state and informs Exchange that it can resume disk writes.
- **9.** VSS manages the shadow copy of the LUN based on the attributes specified by the requestor. For example, VSS could mount the LUN in a Snapshot copy. In a case, however, in which SnapManager is the requestor, SnapManager tells VSS to forget about the shadow copy it just created. This enables SnapManager to have complete control of the Snapshot copy.

Troubleshooting the VSS Hardware Provider

If you attempt to create a backup on a storage system running Data ONTAP, and a Snapshot copy is not created on the storage system, you can troubleshoot the VSS Hardware Provider in several ways.

About this task

There can be many providers installed on the same Windows host, including the VSS software provider, which is always installed. The provider used is determined by either the Requestor or VSS, not the provider. If the first choice provider is not available, an alternative can be silently substituted.

To make a Snapshot copy on the storage system, the Data ONTAP VSS Hardware Provider must be used. If a Snapshot copy on the storage system is not created successfully, verify that the Data ONTAP VSS Hardware Provider was used to create the Snapshot copy.

Only the Data ONTAP VSS Hardware Provider can take a Snapshot copy on a storage system. When you use a VSS requestor, such as SnapManager for Microsoft Exchange or NTBackup, to back up a LUN backed by a storage system running Data ONTAP, the Data ONTAP VSS Hardware Provider must be used for the Snapshot copy to succeed.

Steps

- 1. View the installed providers and verify that the Data ONTAP VSS Hardware Provider is installed.
- 2. Verify that the Data ONTAP VSS Hardware Provider was used to create the Snapshot copy and that it was completed successfully.
- **3.** Verify your VSS configuration.

Viewing installed VSS providers

You can view the VSS providers installed on your host.

Steps

- Select Start > Run and enter the following command to open a Windows command prompt: cmd
- 2. At the prompt, enter the following command:

vssadmin list providers

The output should be similar to the following:

```
Provider name: `Data ONTAP VSS Hardware Provider'
Provider type: Hardware
Provider ID: {ddd3d232-a96f-4ac5-8f7b-250fd91fd102}
Version: 7.0.0.xxxx
```

Verifying that the VSS Hardware Provider was used successfully

You can verify that the Data ONTAP VSS Hardware Provider was used successfully after a Snapshot copy was made.

Step

1. Navigate to System Tools > Event Viewer > Application in MMC and look for an event with the following values:

Source	Event ID	Description
Navssprv	4089	The VSS provider has successfully completed CommitSnapshots for SnapshotSetId <i>id</i> in <i>n</i> milliseconds.

Note: VSS requires that the provider commit a Snapshot copy within 10 seconds. If this time limit is exceeded, the Data ONTAP VSS Hardware Provider logs Event ID 4364. This limit could be exceeded due to a transient problem. If this event is logged for a failed backup, retry the backup.

Verifying your VSS configuration

If the Data ONTAP VSS Hardware Provider failed to run or did not successfully create a Snapshot copy, you must take steps to ensure that VSS is configured correctly.

Steps

- 1. Verify that SnapDrive is installed and running and can communicate with the storage system by performing the following steps:
 - a) Under SnapDrive in the left MMC pane, expand the instance of SanpDrive you want to manage, and then click **Disks**.
 - b) From the menu choices at the top of MMC, navigate to Action > Refresh. No error messages should be displayed.
- 2. Verify that the lun.inquiry.mode option is set to legacy and not standard on the storage system.

This setting applies if you are using Data ONTAP 7.2 and Exchange or SQL Server for VSS-based backups. By default, the mode is set to legacy.

Attention: To change or set this option, you must first stop FC and iSCSI services on your storage system, which might temporarily disrupt any operations currently in progress. Use fcp stop and iscsi stop to stop the services. Use fcp start and iscsi start to restart the services after setting the mode.

- a) To verify the setting, at the storage system prompt, enter the following command: options lun.inguiry.mode
- b) To change the setting, enter the following command:

```
options lun.inquiry.mode legacy.
```

3. Verify that the drives for which the Data ONTAP VSS Hardware Provider failed are backed by a LUN on a storage system running Data ONTAP.

To do this, open MMC and verify that the drives appear under the Disks icon under SnapDrive.

- 4. Verify that the account used by the Data ONTAP VSS Hardware Provider is the same as the account used by SnapDrive by performing the following steps:
 - a) In the left MMC pane, select Services and Applications > Services.
 - b) Double-click the **SnapDrive** service in the main pane and click the **Log On** tab.
 - c) Note the account listed in the **This Account** field, and then click **OK** to close the **SnapDrive Properties** window.
 - d) Double-click the **Data ONTAP VSS Hardware Provider** service in the main pane and click the **Log On** tab.
 - e) Verify that the **This Account** field is selected and that it contains the same account as the SnapDrive service.

Establishing a connection to the storage system

To establish a connection to your storage system, you should understand how to manage iSCSI and Fibre Channel sessions, as well as how to manage your transport protocol settings.

Managing iSCSI sessions

SnapDrive enables you to manage iSCSI sessions on the storage system.

iSCSI Software Initiator node naming requirements

While it is possible to rename iSCSI Software Initiator nodes on a SnapDrive for Windows host, Data ONTAP requires you to use standard iSCSI Software Initiator node names.

When you install the Microsoft iSCSI Software Initiator, an applet is installed that enables you to rename the initiator node to something other than the standard iSCSI qualified name (IQN) or IEEE EUI-64 (EUI) formats. Data ONTAP, however, does not recognize nonstandard initiator node names and returns an error when you attempt to create a LUN using a node name that does not use the IQN or EUI formats.

Following is the format for IQN-type node names:

iqn.yyyy-mm.reverse_domain_name:any

The EUI-type node name format consists of the "eui." prefix, followed by 16 ASCII-encoded hexadecimal characters.

Note: A dash (-) is allowed in the IQN name; however, an underscore (_) is not allowed.

IQN-type node name example

iqn.1991-05.com.microsoft:winclient1

EUI-type node name example

eui.02004567A425678D

Establishing an iSCSI session to a target

Before creating a LUN, you need to have an iSCSI session to the target on which you will manage the LUN.

Before you begin

Verify that the iSCSI service is started.

About this task

- An iSCSI sessions can be established only between the same IP versions. The iSCSI session must be either IPv6 to IPv6 or it must be IPv4 to IPv4. The iSCSI sessions cannot be a combination of the two IP versions.
- A link-local IPv6 address can be used for iSCSI session management and for communication between a host and a target only when both are in the same subnet.

Steps

- 1. Perform the following actions to launch the Create iSCSI Session wizard:
 - a) In the left MMC pane, select the instance of SnapDrive you want to manage.
 - b) Select iSCSI Management.
 - c) From the menu choices at the top of MMC, navigate to Action > Establish Session.
- 2. In the ISCSI Session wizard, click Next.

The Provide Storage System Identification panel is displayed.

3. In the **Provide Storage System Identification** panel, enter the storage system name or IP address of the storage system management port you want to establish the iSCSI session with, and then click **Next**.

Note: The IP address you enter is the storage system management port IP address, not the target portal IP address to which SnapDrive will establish an iSCSI session. You will select the IP address for establishing an iSCSI session in Step 5.

The Provide iSCSI HBA panel is displayed.

- 4. In the upper pane of the **Provide iSCSI HBA** panel, click the radio button next to an available iSCSI HBA to specify the initiator portal you want to use.
- 5. In the lower pane of the Provide iSCSI HBA panel, perform the following actions:
 - a) Select the target portal to which SnapDrive will establish the iSCSI session by clicking the radio button next to the IP address of the target portal you want to use.
 - b) If your target requires authentication, select **Use CHAP**, and then type the user name and password that iSCSI will use to authenticate the initiator to the target.
 - c) Click Next.

The Completing the iSCSI Session Wizard panel is displayed.

- 6. In the Completing the iSCSI Session Wizard, perform the following actions:
 - a) Review the information to make sure it is accurate.
 - b) If the information is not accurate, use **Back** to go back to previous panels of the wizard to modify information.
 - c) Click Finish.

26 | SnapDrive 7.0 for Windows Administration Guide for SAN Environments

Result

An iSCSI session to the target is established.

How SnapDrive uses CHAP authentication

If it is required by your storage system, use CHAP authentication to validate the identity of the login information being sent to the storage system from an iSCSI initiator when you create an iSCSI session.

The Challenge Handshake Authentication Protocol (CHAP) enables authenticated communication between iSCSI initiators and targets. When you use CHAP authentication, you define the CHAP user names and passwords on the initiator and the storage system.

Note: SnapDrive requires that the CHAP password contains at least 12 characters.

During the initial stage of the iSCSI session, the initiator sends a login request to the storage system to begin the session. The login request includes the initiator's CHAP user name and algorithm. The storage system responds with a CHAP challenge. The initiator provides a CHAP response. The storage system verifies the response and authenticates the initiator. The CHAP password is used to compute the response.

Disconnecting an iSCSI target from a Windows host

You can disconnect an iSCSI target from a Windows host if there are no LUNs connected to it.

Before you begin

You must disconnect any LUNs connected to the target before the target can be disconnected.

Steps

- 1. Perform the following actions to disconnect an iSCSI target:
 - a) In the left MMC pane, select the instance of SnapDrive from which you want to disconnect an iSCSI target.
 - b) Double-click iSCSI Management.
 - c) Select the iSCSI session that you want to disconnect.
 - d) From the menu choices at the top of MMC, navigate to Action > Disconnect Target.

A SnapDrive dialog box is displayed prompting you to confirm your action. Additionally, if you have LUNs connected to the iSCSI target, a warning pop-up box is displayed prompting you to confirm that all access to the LUNs on the iSCSI target can be terminated.

2. Click Yes.

Result

The selected iSCSI target is disconnected from the Windows host.

Disconnecting a session to an iSCSI target

You can disconnect an iSCSI session to an iSCSI target when you have more than one session and you do not want to disconnect the target or other sessions connected to that target.

Steps

- 1. Perform the following actions to disconnect a session to an iSCSI target:
 - a) In the left MMC pane, select the instance of SnapDrive for which you want to disconnect an iSCSI session.
 - b) Double-click iSCSI Management.
 - c) Select the iSCSI target from which you want to disconnect a session.
- 2. In the center MMC pane, select the iSCSI session you want to disconnect.
- 3. From the menu choices at the top of MMC, navigate to Action > Disconnect Session.

Note: If you have only one iSCSI session connected to the iSCSI target, performing this procedure will disconnect the iSCSI target from the Windows host.

A SnapDrive pop-up box is displayed prompting you to confirm your action. Additionally, if you disconnect the last session to the iSCSI target and you have LUNs connected to the target, a warning pop-up box is displayed prompting you to confirm that all access to the LUNs on the iSCSI target can be terminated.

4. Click Yes.

Result

The selected iSCSI session is disconnected from the iSCSI target.

Examining details of an iSCSI session

You can view details for each of the iSCSI sessions in SnapDrive.

Steps

- 1. In the left MMC pane, select the instance of SnapDrive you want to examine.
- 2. Double-click iSCSI Management.
- 3. Select the connected iSCSI target whose details you want to view.

Session details are displayed in the lower pane of the center MMC panel.

SnapDrive support for ESX iSCSI initiators

SnapDrive enables you to use ESX iSCSI initiators in VMware environments to provide LUN provisioning and Snapshot copy management operations in a guest OS.

SnapDrive supports ESX iSCSI initiators to provide you the following features:

- LUN enumeration using either the SnapDrive MMC or sdcli.exe
- · LUN migration with vMotion for LUNs connected with ESX iSCSI initiators
- Creation of initiator groups using ESX iSCSI initiators
- Physical-mode RDMs using iSCSI initiators A maximum of 56 LUNs are supported in the guest OS.

Note: The iSCSI initiators must be configured on the ESX server before you can use them with SnapDrive. See your VMware documentation for more information about configuring ESX iSCSI initiators on the ESX server.

ESX iSCSI initiator limitations

SnapDrive supports the use of ESX iSCSI initiators, but there are some limitations you must keep in mind.

- SnapDrive does not support iSCSI session management using ESX iSCSI initiators. You can add targets from the initiator list during LUN creation and connection.
- Multipathing using both FC HBA and ESX iSCSI initiators is not supported.
- Multipathing using both Microsoft iSCSI Software initiators and ESX iSCSI initiators is not supported.

ESX iSCSI initiators require RDM files to provision a LUN, but Microsoft iSCSI Software initiators do not have this requirement.

- RDMs are not supported on an NFS datastore; however, you can store RDMs in any connected VMFS datastore if the virtual machine is stored on NFS.
- MPIO is not supported in a guest OS.
- Windows Server failover clustering is not supported using ESX iSCSI RDM LUNs.
- RDM LUNs larger than 2 TB are not supported in a VMFS 3.0 datastore.

Enabling storage system HTTP communication with SnapDrive

You can enable HTTP capabilities for users who want to communicate with a storage system and SnapDrive using HTTP.

About this task

A storage system administrator must assign the appropriate capabilities to users who do not have administrator access on the storage system.

If you want to restricting capabilities for users who do not have administrator access to the storage system, see the RBAC information in *Configuring and using access control.*

Steps

- 1. Create a local group on the storage system to which you want to enable HTTP communication with SnapDrive.
- 2. Associate the new group with a role that has both api and login capabilities:
- 3. Add the SnapDrive user to the new storage system group.

Example of how to enable HTTP communication with SnapDrive

This example creates on Storage1 a role called snaphttp2 with api and login capabilities, creates a group called **snapadmins** with the **snaphttp2** role capabilities, and adds the user **snapadmin1** to the **snapadmins** group.

```
Storage1> useradmin role add snaphttp2 -a api-*,login-http-admin
The role 'snaphttp2' has been added.
Role <snaphttp2> added.
Storage1> useradmin group add snapadmins -r snaphttp2
The group 'snapadmins' has been added.
Group <snapadmins> added.
Storage1> useradmin user add snapadmin1 -g snapadmins
The user 'snapadmin1' has been added.
User <snapadmin1> added.
```

Using SnapDrive in Microsoft environments

You can use SnapDrive in Microsoft environments to support cluster shared volumes, to create and manage LUNs, to manage Snapshot backup copies, and to manage space on your storage system.

Understanding the new features of Windows Server 2012

The introduction of Windows Server 2012 provides new features for use with your Microsoft products. The changes include improvements to CSV 2.0 and some enhancements to Hyper-V.

The following features are new to Windows Server 2012:

- CSV 2.0 in Windows Server 2012 on page 30
- *Hyper-V enhancements* on page 30

CSV 2.0 in Windows Server 2012

The introduction of Windows Server 2012 provides new features for Cluster Shared Volume (CSV) 2.0 that include a new file system, changes to CSV writer, changes to CSV shadow copy, and enhancements to CSV backup.

Windows Server 2012 includes the following changes to CSV 2.0:

- The CSV File System (CSVFS) is available on all nodes in the cluster as a new distributed file system.
- CSV Writer serves volume and component-level metadata from the non-requesting node for CSV volumes and acts as a proxy by including the Hyper-V writers from the remote node for the backup session.
- The CSV shadow copy provider acts as the default software provider for CSV volumes and coordinates VSS freeze and VSS thaw across all cluster nodes to provide application and crash consistency.

The CSV shadow copy provider ensures that a CSV Snapshot volume is writable on the requesting node.

• CSV now supports one application-consistent Snapshot volume across all CSVs for multiple virtual machines.

The CSV volume from the Snapshot volume is exposed to all the virtual machine owner nodes, to perform autorecovery.

CSV goes into redirected I/O mode only during Snapshot creation and not during backup.

Hyper-V VSS backup changes with Windows Server 2012

You should be aware that Windows Server 2012 introduces changes to the application transaction log and BackupComplete in Hyper-V VSS.

The following changes to Hyper-V VSS backups are introduced in Windows Server 2012:

- The application transaction log for the virtual machine is maintained in a consistent state in case of hardware snapshot failures.
- The Hyper-V VSS integration components call BackupComplete inside the virtual machine after the VSS requestor (for example, SnapManager for Hyper-V) that is running on the Hyper-V parent calls BackupComplete.

SnapDrive limitations on Windows Server 2012

You should be aware of some limitations when deploying Windows Server 2012 on SnapDrive 6.5 for Windows.

- The Windows Host Utility kit for Windows disables the Windows Server 2012 native space reclamation feature.
- Data ONTAP does not support Windows Server 2012 native thin provisioning.
- SnapDrive 6.5 for Windows does not support Remote VSS and SMB 3.0.
- You cannot perform storage replication from Hyper-V Replica.
- You should not perform any SnapDrive disk-related operations on the LUN hosting the virtual machines during live virtual machine migration and storage migration.
- Although Windows Server 2012 allows you to create a disk larger than 16 TB, you cannot create a Data ONTAP LUN larger than 16 TB.
- When you add new nodes to your cluster on Windows Server 2012, you must remap the LUNs to the new nodes.
- Windows Host Utility kit is mandatory for Windows Server 2012, on both physical and guest operating systems.

Data ONTAP DSM support for Windows MPIO

You can use Data ONTAP DSM for Windows MPIO to help storage systems to integrate with Microsoft MPIO on Windows 2008 R2 and 2012 servers and provide high availability to applications by using path-failover methods.

Microsoft MPIO is a protocol-independent feature that supports multiple data paths to a storage device with iSCSI, Fibre Channel, or SAS. Providing multiple paths that can handle failover increases the availability from a host to the storage system. Windows 2008 R2 x64 servers include support for Microsoft MPIO.

Data ONTAP DSMs for Windows MPIO help storage systems to integrate with Microsoft MPIO on Windows 2008 R2 servers and provide high availability to applications by using path-failover methods. They determine all the paths that point to the same LUN, so that MPIO can group them into the virtual disk that the Windows Server 2008 Hyper-V server mounts. The DSMs are also responsible for communicating with MPIO to identify the path on which to route I/O. This is especially important in the event of a failover. There can be multiple active paths and multiple passive paths. If all the active paths fail, the DSM automatically switches to the passive paths, maintaining the host's access to its storage.

Cluster support

SnapDrive for Windows can be deployed in a variety of cluster configurations.

SnapDrive is supported in the following cluster technologies:

Windows clusters

To protect against node failure, Windows clustering fails over applications from the host node to the surviving node. In Windows 2008 and 2012, this is called Windows failover clustering.

• Active/active storage system configurations

If a storage system fails, the partner storage system takes over the functions of the failed storage system, thus protecting data and ensuring continued storage availability.

Note: SnapDrive LUNs are supported in an active/active storage system configuration; however, during cluster takeover and giveback, SnapDrive operations fail for LUNs located on the active/active storage systems until the takeover and giveback process is completed.

Support for Microsoft Cluster Shared Volumes

You can use SnapDrive for Windows supports LUN provisioning and Snapshot copy management on Microsoft Cluster Shared Volumes (CSVs) using Hyper-V with Windows Server 2008 R2.

Cluster Shared Volumes is an option added to the Failover Clustering feature in Windows Server 2008 R2 that enables all nodes in the same Microsoft cluster concurrent access to files on each CSV-enabled shared LUN. CSV allows multiple virtual hard disks from different virtual machines to be stored on a single LUN.

Note: If you enable CSV after installing SnapDrive for Windows, you must restart the SnapDrive service and close and reopen SnapDrive MMC.

Verifying the cluster group owner

Before you use SnapDrive for Windows to add new Cluster Shared Volumes (CSV) to a Microsoft cluster, it is important to verify which Hyper-V node owns the cluster group because you must create new CSV LUNs on the node that owns the cluster group.

About this task

Perform this step at the Windows server command-line.

Step

1. Enter the following command:

```
cluster group Group Name /status
```

Example

Enter the following command for a cluster group called Available Storage:

cluster group Available Storage /status

The cluster group status is displayed, including the name of the node that owns the cluster group.

Changing the cluster group owner

You must create new CSV LUNs on the node that owns the cluster group. If you are using SnapDrive for Windows to create a new CSV LUN on a Hyper-V node that is not the cluster group owner, you can move ownership to the current node. The disk create operation might fail if either the specified cluster node is not the owner of the resource, or the node is not a possible owner of the resource.

About this task

Perform this step at the Windows server command-line.

Step

1. Enter the following command:

cluster group Group Name /move:node_name

Example

Enter the following command to make the node named HypervNode2 the owner of the cluster group called Available Storage:

cluster group Available Storage /move:HypervNode2

Troubleshooting CSVs

You should be aware of some common CSV errors and guidelines for avoiding these errors.

Enabling the Hyper-V server role

To ensure that Microsoft CSVs display properly when using SnapDrive, you must first enable the Hyper-V server role in Windows Server 2008 R2. This is necessary because SnapDrive uses the Hyper-V server role to tag the disk as a CSV, enabling live migration of virtual machines.

For file sharing, Microsoft allows CSV creation without the Hyper-V role present; however, Microsoft does not recommend this configuration.

Ensuring successful CSV disk enumeration

When you are getting information about disks on a cluster node, the enumeration hangs if SnapDrive cannot access Hyper-V VM storage details.

To ensure successful CSV disk enumeration, do not disconnect a CSV that hosts a VM in an OFF state, outside of SnapDrive. Instead, use SnapDrive to disconnect this type of CSV after you remove the VM from the Hyper-V Manager, to avoid any inconsistent behavior.

Related information

Cluster Shared Volumes Support for Hyper-V Requirements for Using Cluster Shared Volumes in a Failover Cluster in Windows Server 2008 R2

Creating LUNs

SnapDrive enables you to quickly create LUNs on a storage system for use in a Windows environment.

Rules for creating LUNs

To avoid problems creating LUNs when using SnapDrive, you must keep some rules in mind.

- Create LUN names using US-ASCII characters only, even when you are using non-ASCII operating systems.
- If you are adding the CSV LUN to a Windows Server 2008 or 2012 cluster, make sure to create the CSV LUN on the node that owns the cluster group in which you are creating a new physical disk resource.

Note:

- Shared disks on Windows Server 2008 and 2012 cluster nodes that do not own the disks often display as offline in the MMC Disk Management utility; however, the disks continue to function normally on all nodes in the cluster.
- Use the LUN path instead of the UNC path to create a LUN. When you use the UNC path, disk creation fails with the error A device attached to the system is not functioning.
- To ensure that Snapshot copies can be made, follow these guidelines:
 - Do not attempt to create a LUN on a storage system volume that holds anything other than LUNs.
 - Conversely, do not put anything other than LUNs on a storage system volume that contains LUNs.
 - All LUNs in the same storage system volume should be created using SnapDrive or, if they were created outside of SnapDrive, prepared for management in SnapDrive.

What volume mount points are

A *volume mount point* is a drive or volume in Windows that is mounted to a folder that uses NTFS. A mounted drive is assigned a drive path instead of a drive letter, enabling you to surpass the limitation of 26 letters with which to name drives.

SnapDrive supports the creation of up to 128 LUNs. By using volume mount points, you can graft, or mount, a target partition into a folder on another physical disk. After you create a volume mount point in SnapDrive, the volume mount point drive path or label displays in the Microsoft MMC Disk Management pane, as well as in SnapDrive MMC Disk List pane.

For more information about volume mount points, see Microsoft article 280297 and 205524.

Volume mount point limitations

When creating mount points on clustered Windows Servers, keep these additional limitations in mind:

- The mounted volume must be of the same type as its root; that is, if the root volume is a shared cluster resource, the mounted volume must also be shared, and if the root volume is dedicated, the mounted volume must also be dedicated.
- You cannot create mount points to the quorum disk.
- If you have a mount point from one shared disk to another, SnapDrive verifies that they are in the same cluster group and that the mounted disk resource is dependent on the root disk source.

Related information

Knowledge Base article 280297: How to configure volume mount points on a Microsoft Cluster Server Knowledge Base article 205524: How to create and manipulate NTFS junction points

Creating a dedicated LUN

You can use SnapDrive to create dedicated FC-accessed or iSCSI-accessed LUNs.

Before you begin

- Create the dedicated volumes to hold your LUNs on the storage system.
- Verify that the FC or iSCSI services have been started on the storage system.
- Before creating a LUN in a VMware guest OS and when using vMotion, you must manually create initiator groups either by using OnCommand System Manager or at the storage system console.

About this task

Keep the following considerations in mind when creating a LUN:

• Unless the LUN is shared within a Windows cluster, the LUN must not be connected to more than one host.

36 | SnapDrive 7.0 for Windows Administration Guide for SAN Environments

- LUNs should be created using SnapDrive.
- SnapDrive filters volumes, qtrees, and LUNs depending on storage system access control settings that might exist in the AccessControl.xml file on your storage system. During LUN creation, SnapDrive displays the message "Checking access control" to indicate it is checking these access control settings.
- ALUA is supported on ESX 4 and later. When you create a disk on an ESX 4 virtual machine with FC initiators and you choose **Automatic igroup management** to map the disk to the storage system, ALUA is enabled by default on the igroup.

Steps

- 1. Perform the following actions to launch the Create Disk Wizard:
 - a) Select the SnapDrive instance for which you want to create a disk.
 - b) Select Disks.
 - c) From the menu choices at the top of MMC, navigate to Action > Create Disk.

The Create Disk Wizard is launched.

2. In the Create Disk Wizard, click Next.

The Provide Storage System Name, LUN Path and Name panel is displayed.

- **3.** In the **Provide a Storage System Name, LUN Path and Name** panel, perform the following actions:
 - a) In the **Storage System Name** field, type the storage system name where the LUN will be created or select an existing storage system using the pull-down menu.
 - b) In the **LUN Path** field, type the LUN path or select the path on the storage system you added in Step a.
 - c) In the LUN Name field, enter a name for the LUN and click Next.

The Select a LUN Type panel is displayed.

- 4. In the Select a LUN Type panel, select Dedicated, and then click Next.
- 5. In the Select LUN Properties panel, either select a drive letter from the list of available drive letters or type a volume mount point for the LUN you are creating. When you create a volume mount point, type the drive path that the mounted drive will use: for example, G: \mount_drive1\.

Note: The root of the volume mount point must be owned by the node on which you are creating the new disk.

Note: You can create cascading volume mount points (one mount point mounted on another mount point); however, in the case of a cascading mount point created on an MSCS shared disk, you might receive a system event warning indicating that disk dependencies might not be correctly set. This is not the case, however, as SnapDrive will create the dependencies and the mounted disks will function as expected.
- 6. While still in the Select LUN Properties panel, complete the following actions:
 - a) Click **Limit** or **Do not limit** for the option labeled "Do you want to limit the maximum disk size to accommodate at least one snapshot?"

If you keep the default, **Limit**, which is the recommended option, the disk size limits displayed are accurate only when they first appear on the Select LUN Properties panel. When this option is selected, the following actions might interfere with the creation of at least one Snapshot copy:

- Changing the option to **Do not limit** and using SnapDrive to create an additional LUN in the same storage system volume.
- Creating a LUN in the same storage system volume without using SnapDrive
- Storing data objects other than LUNs on this storage system volume.
- b) Select a LUN size, which must fall within the minimum and maximum values displayed in the panel.
- c) Click Next.

If the settings on the storage system volume or qtree on which you are creating the LUN do not allow SnapDrive to proceed with the create operation, the Important Properties of the Storage System Volume panel is displayed, as described in Step 7. Otherwise, Step 7 is skipped.

7. The **Important Properties of the Storage System Volume** panel displays the settings that will be used for the volume or qtree you specified in Step 4 of this procedure.

SnapDrive requires the storage system volume containing LUNs to have the following properties:

- create_ucode = on
- convert_ucode = on
- snapshot_schedule = off

Note: SnapDrive cannot proceed to create a LUN unless these settings are configured as shown. Therefore, you must accept these settings.

Note: The create_ucode and convert_ucode volume options are no longer used, but they are set to maintain backwards compatibility with earlier versions of SnapDrive.

Click Next.

The Select Initiators panel is displayed.

8. In the Initiator List pane, select an initiator for the LUN you are creating.

If you select an iSCSI initiator, and an iSCSI connection to the storage system on which you are creating the LUN does not exist, SnapDrive launches the Create iSCSI Session wizard, and you are prompted to select a target portal and initiator. Also, if your target requires authentication of hosts that connect to it, you can type that information here. After you click OK, the iSCSI connection from the Windows host to the storage system is established, even if you do not complete the Create Disk wizard.

If you have MPIO installed and you are using iSCSI and FC, you have the option to select an iSCSI initiator and several FC initiators.

9. Click Next.

The Select Initiator Group Management panel is displayed.

10. In the Select Initiator Group Management panel, specify whether you will use automatic or manual igroup management. If you select automatic igroup management, SnapDrive uses existing igroups or, when necessary, creates new igroups for the initiator you specified in Step 8. If you select manual igroup management, you manually choose existing igroups or create new ones as needed.

If you specify	Then
Automatic igroup management	Click Next.
	SnapDrive uses existing igroups, one igroup per initiator, or, when necessary, creates new igroups for the initiators you specified in Step 8.
Manual igroup management	Click Next, and then perform the following actions:
	a. In the Select Initiator Groups panel, select from the list the igroups to which you want the new LUN to belong.
	Note: A LUN can be mapped to an initiator only once.
	OR Click Manage Igroups and for each new igroup you want to create, type a name in the Igroup Name text box, select initiators, click Create , and then click Finish to return to the Select Initiator Groups panel.
	b. Click Next.
	Note: The Next button will remain unavailable until the collection of selected igroups contains all the initiators you selected in Step 8.

You are done with igroup management.

11. In the Completing the Create Disk Wizard panel, perform the following actions:

a) Verify all the settings.

If you need to change any settings, click **Back** to go back to the previous Wizard panels.

b) Click Finish.

Disk creation might take several seconds to complete. SnapDrive displays disk creation status in the lower panel of the center MMC pane.

Creating a shared LUN

You can use SnapDrive to create FC-accessed or iSCSI-accessed LUNs that are shared between clustered Windows servers.

Before you begin

• FC or iSCSI services have been started on the storage system.

About this task

You should keep the following considerations in mind when creating a LUN:

- You can create LUNs using SnapDrive.
- SnapDrive filters volumes, qtrees, and LUNs depending on storage system access control settings might exist in the AccessControl.xml file on your storage system, which means that during LUN creation, SnapDrive displays the message "Checking access control" to indicate that it is checking these access control settings.

Steps

- 1. Perform the following actions to launch the Create Disk wizard:
 - a) Select the SnapDrive instance for which you want to create a disk.
 - b) Select Disks.
 - c) From the menu choices at the top of MMC, navigate to Action > Create Disk.

The Create Disk wizard is launched.

2. In the Create Disk wizard, click Next.

The Provide Storage System Name, LUN Path and Name panel is displayed.

- **3.** In the **Provide a Storage System Name, LUN Path and Name** panel, perform the following actions:
 - a) In the **Storage System Name** field, type the name of the storage system on which the LUN will be created or select an existing storage system using the pull-down menu.
 - b) In the LUN Path field, type the LUN path or select the path on the storage system that you added in Step 3a.
 - c) In the LUN Name field, type a name for the LUN and click Next.

The Select a LUN Type panel is displayed.

- 4. In the Select a LUN Type panel, select Shared, and then click Next.
- 5. In the **Information About the Microsoft Cluster Services System** panel, verify that you want the disk to be shared by the nodes listed, and then click **Next**.

The Specify Microsoft Cluster Services Group panel is displayed.

- 6. In the **Specify Microsoft Cluster Services Group** panel, perform one of the following actions and then click **Next**:
 - Select a cluster group from the Group Name drop-down list.
 - Select Create a new cluster group to create a new cluster group.

Note: When selecting a cluster group for your LUNs, choose the cluster group your application will use. If you are creating a volume mount point, the cluster group is already selected. This is because the cluster group owns your root volume physical disk cluster resources. It is recommended that you create new shared LUNs outside of the cluster group.

• Select Add to cluster shared volumes.

7. In the Select LUN Properties panel, either select a drive letter from the list of available drive letters or enter a volume mount point for the LUN you are creating. When you create a volume mount point, enter the drive path that the mounted drive will use: for example, G: \mount_drivel\.

Note: The root of the volume mount point must be owned by the node on which you are creating the new disk.

Note: You can create cascading volume mount points (one mount point mounted on another mount point); however, in the case of a cascading mount point created on an MSCS shared disk, you might receive a system event warning indicating that disk dependencies might not be correctly set. This is not the case, however, as SnapDrive will create the dependencies and the mounted disks will function as expected.

- 8. While still in the Select LUN Properties panel, complete the following actions:
 - a) Click **Limit** or **Do not limit** for the option labeled "Do you want to limit the maximum disk size to accommodate at least one snapshot?"

If you select **Limit**, the disk size limits displayed are accurate only when they first appear on the Select LUN Properties panel. When this option is selected, the following actions might interfere with the creation of at least one Snapshot copy:

- The option is changed to **Do not limit** and SnapDrive is used to create an additional LUN in the same storage system volume.
- A LUN is created in the same storage system volume without using SnapDrive
- Data objects other than LUNs are stored on this storage system volume.
- b) Select a LUN size. The size must fall within the minimum and maximum values displayed in the panel.

Note: If the volume on which you are creating the LUN is thin provisioning enabled, a note displays in the LUN Size pane.

- c) Select whether you want to allow the maximum LUN size for this LUN. No is selected by default, which means that the LUN cannot exceed the size available in the volume.
- d) Click Next.

If the settings on the storage system volume or qtree on which you are creating the LUN do not allow SnapDrive to proceed with the create operation, the Important Properties of the Storage System Volume panel is displayed, as described in Step 8. Otherwise, Step 8 is skipped.

9. The **Important Properties of the Storage System Volume** panel displays the settings that will be used for the volume or qtree you specified earlier in this procedure.

SnapDrive requires the storage system volume containing LUNs to have the following properties:

- create_ucode = on
- convert_ucode = on
- snapshot_schedule = off

Note: SnapDrive cannot proceed to create a LUN unless these settings are configured as shown. Therefore, you must accept these settings.

Note: The create_ucode and convert_ucode volume options are no longer used, but they are set to maintain backwards compatibility with earlier versions of SnapDrive.

10. Click Next.

11. In the Select Initiators panel, perform the following actions:

- a) Double-click the cluster group name to display the hosts that belong to the cluster.
- b) Click the name of a host to select it.

The available initiators for that host are displayed in the Initiator List in the lower half of the pane.

12. In the Initiator List pane, select an initiator for the LUN you are creating.

If you select an iSCSI initiator, and an iSCSI connection to the storage system on which you are creating the LUN does not exist, SnapDrive launches the Create iSCSI Session wizard, and you are prompted to select a target portal and initiator. Also, if your target requires authentication of hosts that connect to it, you can type that information here. After you click OK, the iSCSI connection from the Windows host to the storage system is established, even if you do not complete the Create Disk wizard.

If you have MPIO installed and you are using FC, you have the option to select several FC initiators.

13. Repeat Step 10 and Step 11 for all hosts, and then click Next.

Note: The Next button remains unavailable until initiators for all hosts of a cluster have been selected.

The Select Initiator Group management panel is displayed.

14. In the Select Initiator Group management panel, specify whether you will use automatic or manual igroup management. If you select automatic igroup management, SnapDrive uses existing igroups or, when necessary, creates new igroups for the initiators you specified in Step 10 through Step 12. If you select manual igroup management, you manually choose existing igroups or create new ones as needed.

If you specify	Then
Automatic igroup	Click Next.
management	You are done with igroup management.

If you specify	Then
Manual igroup management	Click Next, and then perform the following actions:
	a. In the Select igroups panel, select from the list the igroups to which you want the new LUN to belong. Repeat this action for each node in the cluster.
	Note: A LUN can be mapped to an initiator only once.
	OR
	Click Manage igroups and for each new igroup you want to create, type a name in the igroup Name text box, select initiators, click Create , and then click Finish to return to the Select igroups panel.
	b. Click Next.
	Note: The Next button will remain unavailable until the collection of selected igroups contains all the initiators you selected in Step 11.

The Completing the Create Disk Wizard panel is displayed.

15. In the Completing the Create Disk Wizard panel, perform the following actions:

a) Verify all the settings.

If you need to change any settings, click **Back** to go back to the previous wizard panels.

b) Click Finish.

Disk creation might take several seconds to complete. SnapDrive displays disk creation status in the lower panel of the center MMC pane.

Windows Server 2008 and 2012 failover cluster support

SnapDrive supports the use of shared LUNs in a Windows Server 2008 and 2012 failover cluster for all cluster configuration models using up to eight nodes. However, you cannot perform enumeration and management of offline disk resource tasks in a failover cluster.

For more information about failover clusters and cluster configuration models, see the Windows Server 2008 documentation and online help.

Related information

Failover Clusters in Windows Server 2008

Configuring a Windows Server 2008 and 2012 failover cluster witness disk

SnapDrive for Windows and Windows Server 2008 and 2012 provides a simpler way of configuring a shared disk as a witness disk than in earlier versions of SnapDrive and Windows Server. While SnapDrive still supports the previous method of configuring a quorum, it is no longer necessary to

create a shared disk before creating a cluster, nor is it necessary to connect that shared disk to the node before that node can be added to the cluster.

Before you begin

- Install the Windows Server 2008 or 2012 failover cluster feature. For more information, see Windows Server 2008 or 2012 online Help.
- Create the failover cluster using the Windows Server 2008 or 2012 MMC snap-in, Failover Cluster Management.

For more information, see Windows Server 2008 or 2012 online Help.

• Create a shared LUN, ensuring that you select the Microsoft Cluster Services Group named "Cluster Group" to own that disk resource.

Note: The shared LUN must be created on the node that owns "Cluster Group." To determine which node owns "Cluster Group," type the cluster group command at a Windows command prompt.

Steps

- 1. Navigate to Start > Administrative Tools > Failover Cluster Management to launch the Windows Server 2008 or 2012 Failover Cluster Management snap-in.
- 2. Click the name of the failover cluster for which you want to configure the witness disk.
- 3. From the menu choices at the top of the snap-in, navigate to Action > More Actions > Configure Cluster Quorum Settings.

The Configure Cluster Quorum Wizard is launched.

4. In the Configure Cluster Quorum Wizard, click Next.

The Select Quorum Configuration panel is displayed.

5. In the Select Quorum Configuration panel, select Node and Disk Majority, and then click Next.

The Configure Storage Witness panel is displayed.

6. In the **Configure Storage Witness** panel, select the shared LUN you created in SnapDrive to be the witness disk, and then click **Next**.

The Confirmation panel is displayed.

7. In the **Confirmation** panel, click **Next** to configure the cluster quorum settings.

The quorum settings are configured and the Summary panel is displayed.

8. In the Summary panel, click Finish to close the wizard.

Creating a highly available Hyper-V virtual machine using SnapDrive

You can make a Hyper-V virtual machine highly available by creating the VM on a SnapDrive shared LUN in a Windows 2008 failover cluster.

Before you begin

- You must have already installed the Windows Server 2008 failover cluster feature. For more information, see Windows Server 2008 online Help.
- You must have already created a shared LUN owned by the Microsoft Cluster Services Group named "Available Storage."

Note: You can determine which node owns "Available Storage" by typing the cluster group command at a Windows command prompt.

Steps

- 1. Open Hyper-V Manager and use Virtual Network Manager to create and configure your virtual networks on all physical hosts.
- 2. From the Actions menu, click Virtual Network Manager to create or add a virtual network.

See your Hyper-V documentation for more information.

Note: All nodes in the same cluster must use the same name for the virtual network that provides external networking for the virtual machines. The virtual network adapters must be named the same on all Hyper-V hosts to enable quick migration.

- 3. From the Hyper-V Manager Action menu, click New > Virtual Machine.
- 4. In the Virtual Machine wizard, choose the option to store the virtual machine in a new folder and specify the location of your shared storage.
- 5. Navigate to Start > Administrative Tools > Failover Cluster Management to launch the Windows Server 2008 Failover Cluster Management snap-in.

Note: Make sure the virtual machine is not running before you make it highly available.

6. Using the High Availability (HA) wizard, navigate to Failover Cluster > Services and Applications > Configure a Server or Application and select the virtual machine you created in Steps 1 through 3 to add it to a cluster group.

You have created a virtual machine cluster group.

7. Use SnapDrive to create additional shared disks as virtual machines to add to your virtual machine cluster group, as needed.

Support for creating disks on a virtual storage server operating in clustered Data ONTAP

You can use SnapDrive to create disks on a virtual storage server that is operating in clustered Data ONTAP. You must have virtual storage server administrator privileges to create a disk. You must also meet certain requirements before you create the disk.

You can create a disk using the Create Disk wizard.

Requirements for creating disks on a virtual storage server in clustered Data ONTAP

You must be aware of the configuration requirements before you create a disk on a virtual storage server in clustered Data ONTAP.

- The virtual storage server administrator (vsadmin) account must be unlocked. For more information about how to unlock a vsadmin account, see the *Clustered Data ONTAP Commands: Manual Page Reference.*
- SAN data protocol must be enabled and configured on the virtual storage server.
- You must use HTTP or HTTPS protocol to communicate with the virtual storage server. You must not use RPC protocol, which is not supported in clustered Data ONTAP.
- The virtual storage server management LIF's firewall policy must be set to management, the role set to data, and the data protocol set to none.
- Before you provision LUNs on a virtual storage server, you must add your storage system name and IP address to the Windows Server etc\host file.

You must be able to resolve the virtual storage server management LIF IP address to the virtual storage server name.

For more information about virtual storage servers and virtual storage server administrator capabilities, see the *Clustered Data ONTAP System Administration Guide for Vserver Administrators*.

GPT partition support

SnapDrive supports the GUID partition table (GPT) partitioning style on new LUNs created by SnapDrive when you have Data ONTAP installed on your storage system.

Neither SnapDrive nor Data ONTAP support MBR LUNs that are converted to GPT-style LUNs. If you have an existing MBR-style LUN, rather than converting, you must create a new GPT LUN by using SnapDrive, and then copy all the data from the MBR LUN to the GPT LUN.

GPT LUNS have a Microsoft reserved partition (MSR), which is invisible to applications like Disk Management and Windows Explorer. When you create a LUN that has the GPT partition style, the LUN size appears smaller than the size you specified when you created it. This is due to the space used by the MSR. To create a GPT LUN that is less than 16 GB, you must have at least 32 MB of space available for the MSR. For GPT LUNs greater than or equal to 16 GB, you must have at least 128 MB for MSR space.

Configuring SnapDrive to create LUNs with Virtual Fibre Channel

Windows Server 2012 supports direct connectivity of virtual Fibre Channel (vFC) ports within guest operating systems, which allows your guest operating system to rapidly connect to your storage system array and create LUNs.

About this task

You can create a maximum of four Fibre Channel Adapters per VM.

Steps

- 1. In Hyper-V Manager, select the VM for which you want to configure vFC, and select Virtual San Manager... in the Actions pane.
- 2. Follow the steps in the Virtual SAN Manager wizard to create and configure a virtual SAN.
- **3.** In Hyper-V Manager, turn off the VM for which you are configuring vFC, and then select settings from the **Actions** pane for your VM.
- 4. In the Add Hardware drop-down of the Settings for VM window, select Virtual Fibre Channel and click Add.
- 5. In the Hardware pane of the Settings for VM window, select the virtual SAN you just created.
- 6. In the Virtual SAN field Fibre Chanel Adapter pane, select Virtual Fibre Channel SAN, and then click Apply, then OK.

Managing LUNs

SnapDrive for Windows enables you to manage LUNs from a Windows environment.

How LUNs work

The following section describes how LUNs work by interacting with Windows hosts and with storage systems.

How the storage system interacts with the LUN

To the storage system, a LUN is a logical representation of a physical unit of storage.

The storage system handles each LUN as a single storage object. The size of this LUN is slightly larger than the raw disk size reported to the Windows host. SnapDrive must be used to expand the disk, because SnapDrive expands both the LUN and the Windows partition.

How Windows hosts interact with a LUN

You manage LUNs on the storage system just as you manage other Windows disks that store application data.

LUNs on the storage systems are automatically formatted by SnapDrive the same way that you format other Windows disks. Moreover, a Windows host interacts with all user data files on the LUN as if they were NTFS files distributed among the disks of a locally attached RAID array.

LUN capabilities and limitations

A LUN managed by SnapDrive can be used for data storage and can be a boot disk. A LUN cannot be a dynamic disk.

SnapDrive can also make a Snapshot copy of LUNs when they are used for data storage, and it can work with SnapMirror at the volume level and SnapVault at the qtree level for data protection.

Protocols for LUN access

You can access the SnapDrive-created LUNs using either FC or iSCSI protocol, or both.

You must have the appropriate hardware and firmware, if any, and software installed on your host and the storage system before you can use these protocols to access LUNs.

How data is accessed from LUNs

In a SAN environment, an initiator (on the Windows host) initiates a SCSI I/O operation to a target (storage system). The operation can be initiated using either the FC or the iSCSI protocol, depending on the type of initiator installed on your Windows host and the setup on the target. A target can receive SCSI requests using FC if a supported HBA is installed and FC is licensed. Similarly, a target can receive SCSI requests using iSCSI if a supported HBA or the Microsoft iSCSI Software Initiator is installed, and if iSCSI is licensed.

After a target receives a SCSI I/O request, the appropriate operation is performed by writing data to or fetching data from the LUN.

List of guidelines for connecting LUNs

Properly connecting a LUN in SnapDrive for Windows enables you to save, delete, modify, and manage the files it contains from the Windows host. You can also make Snapshot copies of the entire disk and restore the disk, along with its contents, to the state captured by a previous Snapshot copy.

- LUNs should be created using SnapDrive. If you want to connect a LUN that was not created in SnapDrive, you must take some steps to prepare the LUNs for SnapDrive management.
- Unless the LUN is shared within a Windows cluster, the LUN must not be connected to more than one host.

You should not try to connect to a LUN if it is already connected to another machine; SnapDrive does not support such simultaneous use.

- Offline LUNs are visible from the SnapDrive MMC, and if you attempt to connect to these LUNs using the Connect Disk wizard, you receive an error.
 To ensure that the LUN you are attempting to connect is online, unmap the LUN, bring it online from the storage system, and refresh the SnapDrive MMC.
- On operating systems earlier than Windows Server 2012, after you have added a new node to a cluster, you must run the Connect Disk wizard to connect a shared disk or a CSV disk to the new node.

Connecting to a LUN

You can connect your SnapDrive for Windows host to a LUN by using the Connect Disk wizard in the SnapDrive MMC snap-in.

Before you begin

You must have manually created initiator groups by using OnCommand System Manager or the storage system console.

About this task

SnapDrive filters volumes, qtrees, and LUNs depending on storage system access control settings that might exist in the AccessControl.xml file on your storage system. During LUN connection, SnapDrive displays the message "Checking access control" to indicate that it is checking these access control settings.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage. Then, select **Disks**.
- 2. From the menu choices at the top of MMC, navigate to Action > Connect Disk.
- 3. In the Connect Disk wizard, click Next.
- 4. In the **Provide a Storage System Name, LUN Path and Name** panel, perform the following actions:
 - a) In the "Storage System Name" field, type the name of the storage system where the LUN will be connected, or choose a storage system from the drop-down list.
 - b) In the LUN Path field, type the path to the LUN, or click **Browse** and navigate to the LUN you want to connect.
 - c) Click Next.
- 5. If the LUN is a dedicated disk, go to the next step; otherwise, if the LUN is a Windows cluster resource, perform the following steps in the **Specify Microsoft Cluster Services Group** panel:
- 6. In the **Specify Microsoft Cluster Services Group** panel, perform one of the following actions and then click **Next**:
 - Select a cluster group from the Group Name drop-down list.

• Select Create a new cluster group to create a new cluster group.

Note: When selecting a cluster group for your LUNs, choose the cluster group your application will use. If you are connecting to a volume mount point, the cluster group is already selected, because the cluster group owns your root volume physical disk cluster resources. You should create new shared LUNs outside of the cluster group.

- Select Add to cluster shared volumes.
- 7. In the Select LUN Properties panel, perform the following actions:
 - a) Either select a drive from the list of available drive letters, or enter a mount point for the LUN you are connecting.

When you connect a volume mount point, enter the drive path that the mounted drive will use: for example, G:\mount_drive1\.

Note: The root volume of a new mount point must be owned by the node on which you are creating the new disk.

Note: You can connect cascading volume mount points (by mounting one mount point on another mount point); however, in the case of a cascading mount point connected on a MSCS shared disk, you might receive a system event warning indicating that disk dependencies might not be correctly set. This is not the case, however, and the mounted disks function as expected.

- b) Click Next.
- 8. In the Select Initiators panel, choose an initiator for the LUN.

Note: If MPIO is installed on the system, you can select multiple FC initiator ports or one iSCSI session.

If the LUN will	Then do this
Belong to a single system	Select at least one initiator for the LUN you are creating from the list of available initiators, and then click Next .
Be a Windows cluster resource	 a. Double-click the cluster name to display the hosts that belong to the cluster. b. Click the name of a host to select it. c. Select at least one initiator for the LUN you are creating from the list of available initiators on that host. d. Repeat steps b and c for all hosts in the cluster. e. Click Next. Note: Next remains unavailable until initiators for all hosts of a cluster are selected.

9. In the Select Initiator Group Management panel, specify whether to use automatic or manual igroup management.

If you select automatic igroup management, SnapDrive uses existing igroups or, when necessary, creates new igroups for the initiators you have specified. If you select manual igroup management, you manually choose existing igroups or create new ones as needed.

If you specify	Then
Automatic igroup management	Select Automatic, and then click Next.
Manual igroup management	Select Manual , click Next , and then, in the Select Initiator Groups panel, perform one of the following actions:
	• Select from the list the igroups to which you want the LUN to belong. Repeat this action for each node in the cluster, then click Next .
	Note: A LUN can be mapped to an initiator only once.
	 Click Manage igroups, and for each new igroup you want to create, type a name in the Igroup Name text box, select initiators, and click Create. Then, click Finish to return to the Select Initiator Groups panel, and click Next.
	Note: The Next button in the Select Initiator Groups panel remains unavailable until the collection of selected igroups contains all the initiators you previously selected for use.

10. In the Completing the Connect Disk Wizard panel, perform the following actions:

- a) Verify all the settings.
- b) If you need to change any settings, click **Back** to go back to the previous wizard panels.
- c) Click Finish.

The newly connected LUN appears under **SnapDrive > Disks** in the left MMC panel.

Making drive letter or path modifications to a LUN

SnapDrive for Windows enables you to add, change, or remove a drive letter or mount point path for an existing LUN.

Adding, removing, or changing a drive letter or path for an existing LUN

You can add, remove, or change a drive letter or mount point path for an existing LUN using the SnapDrive for Windows MMC snap-in.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then expand **Disks** and select the disk you want to manage.
- From the menu choices at the top of MMC, navigate to Action > Change Drive Letter and Paths.
- **3.** In the **Change Drive Letter and Paths** window, click **Add**, **Remove**, or **Change**, depending on the action you want to take.

If	Then
You are removing a drive letter or path	Click OK to proceed with the operation.
You are adding or changing a drive letter or path	In the Add or Change Drive Letter or Path window, select a drive letter or enter path in the space provided, then click OK .

Note: The Change option is unavailable for mount points.

By removing the last volume mount point on a shared disk, SnapDrive for Windows removes the resource dependency from the root disk. If you are creating a mount point from one shared disk to another, SnapDrive verifies they are in the same cluster group and creates a dependency to the root disk resource if it is the first volume mount point to that root disk.

Note: When you create the first volume mount point to a root disk that is shared and is being used by MSCS, SnapDrive, as part of the resource dependency process, takes the physical disk resource offline, presenting the mounting volume. As a result, any other cluster resources that depend on the physical disk resource will also be taken offline. An example of this is the Exchange System Attendant cluster resource. SnapDrive automatically brings the physical disk resource back online but will not bring the Exchange resources back online. Exchange resources should be brought back online manually using the Cluster Administrator.

Moving a mount point with Windows Explorer

Complete these steps to move an existing LUN mount point using Windows Explorer.

Steps

- 1. Identify the folder that represents the volume mount point.
- **2.** Using Windows Explorer, cut and paste the mount point folder to another folder on the same drive.

Note: You cannot cut and paste a volume mount point folder to a different drive.

Guidelines for disconnecting or deleting LUNs

You can disconnect a LUN from a host in SnapDrive for Windows without affecting the contents of the LUN, or you can permanently delete a LUN. You can also use SnapDrive for Windows to disconnect a LUN in a Snapshot copy or FlexClone volume.

- When the host is disconnected from a LUN, you cannot see or modify the LUN's contents, make Snapshot copies of the LUN, or use Snapshot technology to restore the LUN to a previous Snapshot copy; however, the LUN still exists on the storage system volume.
- You must make sure that the LUN you are disconnecting or deleting is not monitored with the Windows Performance Monitor (perfmon).
- Make sure that the LUN you want to disconnect or delete is not being used by a host.
- You can only disconnect or delete a shared LUN (that is, a non-quorum disk) after removing the cluster resource dependencies from the LUN and verifying that all nodes in the cluster are powered on and functioning properly.

Note: SnapDrive manages the dependencies to the root disk for volume mount points.

- When disconnecting or deleting LUNs on a Microsoft cluster, you must make sure that all hosts in the cluster are available to SnapDrive for the disconnect or delete operation to succeed.
- You can disconnect a quorum disk only after replacing it with another disk that takes over as a quorum disk for the cluster.
- Use the Delete Disk feature cautiously, because after you delete a LUN, you cannot use SnapDrive to undelete it.
- If you disconnect a LUN in a FlexClone volume that SnapDrive for Windows created and it is the last LUN connected on the volume, SnapDrive, deletes that volume resulting in the deletion of all LUNs in the FlexClone volume. SnapDrive displays a message warning you that the FlexClone volume might be deleted.

To avoid automatic deletion of the FlexClone volume, rename the volume before disconnecting the last LUN. When you rename the volume, be sure to change more than just the last integers in the name. For instance, if the FlexClone volume is named sdw_cl_myvol_0, rename it to new_sdwvol_0, and not to sdw_cl_myvol_20. If you rename only the last integers in the volume name, SnapDrive still recognizes that it created that volume and it will delete the volume when you disconnect the last LUN. The renamed FlexClone volume is visible after you refresh SnapDrive MMC.

• If you unmap or delete a LUN from the storage system console, you must also remove any stale RDMs that result from unmapping or removing the LUN.

Disconnecting a LUN

You can use the SnapDrive for Windows MMC snap-in to disconnect a dedicated or shared LUN, or a LUN in a Snapshot copy or in a FlexClone volume.

Before you begin

- Make sure that neither Windows Explorer nor any other Windows application is using or displaying any file on the LUN you intend to disconnect. If any files on the LUN are in use, you will not be able to disconnect the LUN except by forcing the disconnect.
- If you are disconnecting a disk that contains volume mount points, change, move, or delete the volume mount points on the disk first before disconnecting the disk containing the mount points; otherwise, you will not be able to disconnect the root disk. For example, disconnect G: \mount_disk1\, then disconnect G: \.
- Before you decide to force a disconnect of a SnapDrive LUN, be aware of the following consequences:
 - Any cached data intended for the LUN at the time of forced disconnection is not committed to disk.
 - Any mount points associated with the LUN are also removed.
 - A pop-up message announcing that the disk has undergone "surprise removal" appears in the console session.

About this task

Under ordinary circumstances, you cannot disconnect a LUN that contains a file being used by an application such as Windows Explorer or the Windows operating system. However, you can force a disconnect to override this protection. When you force a disk to disconnect, it results in the disk being unexpectedly disconnected from the Windows host.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then expand **Disks** and select the disk you want to manage.
- From the menu choices at the top of MMC, navigate to either Action > Disconnect Disk to disconnect normally, or Action > Force Disconnect Disk to force a disconnect.
- 3. When prompted, click Yes to proceed with the operation.

Note: This procedure will not delete the folder that was created at the time the volume mount point was added. After you remove a mount point, an empty folder will remain with the same name as the mount point you removed.

The icons representing the disconnected LUN disappear from both the left and right MMC panels.

Deleting a LUN

You can delete a LUN using the SnapDrive for Windows MMC snap-in.

Before you begin

If you are deleting a disk that contains volume mount points, disconnect the mounted volumes on the disk first before deleting the disk. For example, disconnect G:\mount_disk1\, then disconnect G: \. If you do not disconnect the mounted volume before you delete it, Windows keeps the volume mount point information in the Recycle Bin and both Windows and SnapDrive continue to see the mount point as valid. If your volume mount point contains data, remember that SnapDrive will not warn you that data is present when you delete the mount point.

About this task

Use the Delete Disk feature cautiously, because after you delete a LUN, you cannot use SnapDrive to undelete it.

Steps

- 1. Make sure that neither Windows Explorer nor any other Windows application is using or displaying any file on the LUN you intend to delete.
- 2. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then expand **Disks** and select the disk you want to manage.
- 3. From the menu choices at the top of MMC, navigate to Action > Delete Disk.

4. When prompted, click Yes to proceed with the operation.

Note: This procedure will not delete the folder that was created at the time the volume mount point was added. After you remove a mount point, an empty folder will remain with the same name as the mount point you removed.

The icons representing the deleted LUN disappear from MMC.

Deleting folders within volume mount points

You can delete a folder within a volume mount point by bypassing the Windows Recycle Bin.

About this task

When you use the Windows Explorer to delete a folder that you have created under a volume mount point, you might receive an error message similar to the following, where *Foldername* is the name of the folder you want to delete:

Cannot delete Foldername: Access Denied. The source file may be in use.

This happens because the Windows Recycle Bin does not understand volume mount points and tries to delete the drive on which the mount point resides rather than the folder on the mount point.

For more information about deleting folders within volume mount points, see Microsoft article 243514.

Steps

- 1. Using Windows Explorer, click the folder you want to delete.
- 2. Click Shift and Delete simultaneously to bypass the Recycle Bin.

Guidelines for resizing disks

As your storage needs change, you might need to resize a disk to hold more data or shrink the disk to make space available on the storage system volume.

- The ability to shrink a disk is supported only on Windows Server 2008 and later.
- A good time to expand a disk is right after you have expanded your storage system volumes.
- A LUN cannot be expanded to more than 10 times its original size.
- LUNs with MBR-style partitions have a size limit of 2 TB, and LUNs with GPT-style partitions have a storage system size limit of 16 TB.
- Understand the storage management implications of resizing the LUN volume size.
- If it is necessary to restore a LUN from a Snapshot copy made before the LUN was resized, SnapDrive for Windows automatically resizes the LUN to the size of the Snapshot copy and performs the restore operation.

When the disk is restored, SnapDrive reconnects the disk. If you restore a LUN from a Snapshot copy made before the LUN was resized, the LUN returns to its former size before it was reduced or enlarged. After the restore operation, data added to the LUN after it was resized must be restored from a Snapshot copy made after it was resized.

- When creating a quorum disk, make sure that it is the size recommended by Microsoft for your Windows cluster setup.
- For extra data protection during disk resizing operations, you can make a Snapshot copy prior to resizing.
- SnapDrive for Windows cannot shrink LUNs by more than half. This is a Microsoft Support limitation. In Microsoft Windows Server 2008 SP2 and Windows Server 2008 R2, you might not be able to shrink a volume by more than half the original size.

Related information

Microsoft support article 2020591: Unable to shrink a volume more than half the original size

Resizing a disk

You can resize a disk using the SnapDrive for Windows MMC snap-in to either increase or decrease the amount of space it uses.

Before you begin

Take a Snapshot copy of your disk before you resize it. If necessary, you can use the Snapshot copy to restore the disk to its original size.

About this task

If the disk you want to resize is a quorum disk in a Microsoft cluster configuration, instead of performing the following steps, you need to follow the procedure to resize a quorum disk.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then expand **Disks** and select the disk you want to manage.
- 2. From the menu choices at the top of MMC, navigate to Action > Resize Disk.
- **3.** Next to "Maximum size" in the **Resize Disk** window, leave "Reserve space for at least one Snapshot copy" selected.

Note: When you select this option, the disk size limits displayed are accurate only when they first appear on the Select LUN Properties panel.

- 4. In the "New size" box, either type a value, or use the slider bar to increase or decrease the amount of space the disk uses.
- 5. Select the "Take a Snapshot before resizing the lun" check box, to take a Snapshot copy before you resize your disk.
- 6. Click OK.
- 7. Create a new Snapshot copy of the resized disk.

After you finish

If you change the size of your disk, you might need to close and reopen MMC before the resized disk size becomes visible in the Disk Management snap-in.

Resizing a quorum disk

You cannot resize a disk while it is serving as a quorum disk, so a few special steps are required when using SnapDrive for Windows to resize a quorum disk in a Microsoft cluster.

Before you begin

Decide whether you would like to keep the disk as quorum or designate a new disk as quorum.

About this task

If you decide to create a new LUN and designate that disk as a quorum, you can simply create a new disk, designate it as the quorum using the Cluster Administrator on the owning node of your Windows cluster, and then delete the original quorum disk. Otherwise, follow this procedure to keep the original quorum disk and resize it.

For information about how to set a disk as a quorum, see your Windows documentation.

Steps

- 1. Create a new disk.
- 2. Designate the newly created disk as the quorum using the Cluster Administrator on the owning node of your Windows cluster.
- 3. Resize the original quorum disk (which is now a regular LUN).
- 4. Designate the expanded disk as the quorum using the Cluster Administrator on the owning node of your Windows cluster.
- 5. Delete the disk you created in Step 1.

Managing LUNs not created in SnapDrive

You can use SnapDrive for Windows to manage LUNs not created by SnapDrive by preparing the LUNs for SnapDrive management.

Before you begin

- The names of LUNs, qtrees, igroups, volumes, and storage systems must use US-ASCII characters only.
- The LUN should have a single partition on an NTFS file system.
- The ostype of the LUN must match the partition style and the type of OS accessing it, as follows:

Windows OS and partition style	ostype
Windows Server 2008 (MBR or GPT partitioned)	windows_2008
Windows Server 2008 Hyper-V (MBR or GPT partitioned)	hyper_v
Windows Server 2012 (MBR or GPT partitioned)	hyper_v

If the LUN was created with ostype set to a different value than you need, you can create a new LUN using SnapDrive, copy the data to it, and delete the original LUN. You can use OnCommand System Manager to check the ostype of the LUN. For more information about ostype, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

Note: In earlier versions of SnapDrive, LUNs were required to have the .lun extension to be managed by SnapDrive; however, .lun extensions are no longer required as of SnapDrive 4.2.

Steps

- 1. If you have a clustered Windows configuration, perform the following steps:
 - a) In SnapDrive, create a shared disk on the storage system to temporarily designate as the quorum disk.
 - b) Right-click the resource and select **Properties > Dependencies** to record all dependencies for each resource in this cluster group, using the Windows cluster management console.
 - c) Designate the newly created disk as the quorum on the owning node of your Windows cluster using the Windows cluster management console.

For information about how to set a disk as a quorum, see your Windows documentation.

- d) Check that space reservation is enabled or that there is enough space available for space reservation to be enabled.
- 2. Shut down the stand-alone Windows host, or all nodes in a cluster.

Shutting down your Windows hosts ensures that all data has been flushed and that Snapshot copies are consistent.

- **3.** Using OnCommand System Manager or the storage system console, complete the following steps:
 - a) Unmap the LUN from the initiator group.
 - b) Make a Snapshot copy of the volume on which the LUNs reside.
- 4. Restart the stand-alone Windows host, or all nodes in a cluster.
- **5.** If you have a clustered Windows configuration, delete the shared disk resource using the Windows cluster management console.
- 6. Connect to the LUN using SnapDrive.

- 58 | SnapDrive 7.0 for Windows Administration Guide for SAN Environments
 - 7. If you are working in a Windows cluster, perform the following substeps:
 - a) Designate the newly connected LUN as the quorum using the Windows cluster management console on the owning node of your cluster.
 - b) Re-create any dependencies you recorded in Step 1.
 - c) Delete the temporary disk you created in Step 1.
 - 8. Using SnapDrive, make a Snapshot copy of the newly connected LUN.

Guidelines for renaming LUNs

You can rename an existing LUN to conform to SnapDrive for Windows naming requirements.

- LUN names must use US-ASCII characters only, even when you use non-ASCII operating systems.
- You can rename a LUN while it is connected to a host, but rename only the LUN or volume while connected.

If any other objects need to be renamed, such as qtrees or igroups, disconnect the LUN, rename the objects, and then reconnect the LUN.

Requirements for dynamically adding and removing pass-through disks on Hyper-V virtual machines

SnapDrive supports the Windows Server 2008 R2 feature that enables you to add or remove a passthrough disk on a Hyper-V virtual machine without shutting down that virtual machine. This is sometimes called "the hot add and remove feature," and you must ensure that your configuration meets certain requirements before you can use it.

A *pass-through* disk is a disk that is physically connected to a Hyper-V parent host and is assigned to a Hyper-V virtual machine as a SCSI hard disk for use by that virtual machine.

The following configuration requirements must be met to use this feature:

- You must be using Windows Server 2008 R2 or later.
- .NET 3.5 SP1 or higher must be installed on the Hyper-V parent hosts.
- A version of SnapDrive that supports this feature must be installed on all Hyper-V parent hosts and on all targeted Hyper-V virtual machines.
- If iSCSI is used for the pass-through disk, the iSCSI groups must already be configured and the iSCSI session between the storage system and the Hyper-V parent host must already exist.
- The SCSI controller must be added to target virtual machines in advance.
- The Hyper-V parent host and the virtual machine must have TCP/IP network communication to all parent hosts and vice versa.
- A virtual machine must not have MPIO enabled.
- On an HA VM on a shared disk, the virtual network names for both the nodes must be the same.
- A pass-through disk must be created from an HA VM by providing the drive letter or mount point.
- All the nodes in the cluster must be running the same version of Data ONTAP DSM.
- The VM host name and the VM name must be the same.

- If you uninstall SnapDrive, all SnapDrive registry key entries are removed from your Windows host. If you reinstall SnapDrive, you must reconfigure pass-through disks.
- When creating pass-through disks in the MMC, SnapDrive shows all the FCP initiators (connected and unconnected) if it fails to find connected FCP initiators.

Hyper-V pass-through disk support limitations

SnapDrive supports Hyper-V pass-through disks; however, there are some limitations related to this feature.

- SnapDrive does not support direct-attached storage to a Hyper-V parent host. Only Data ONTAP iSCSI and FC LUNs are supported.
- SnapDrive does not support LUNs already mapped to a Hyper-V parent host.
 A pass-through disk must be freshly provisioned from a storage system to be added dynamically.
- Windows 2008 R2 does not support pass-through disks using IDE. Windows 2008 R2 supports only SCSI disks.
- Microsoft cluster shared volume (CSV) disks are not supported.
- Hyper-V does not support the SnapDrive space reclaimer feature.
- Pass-through LUN creation fails when Data ONTAP DSM is installed on Hyper-V virtual machines.

Note: Removing MPIO does not remove MPIO iSCSI sessions that SnapDrive might have been using. After removing MPIO, you must remove existing iSCSI sessions and, if needed, create new iSCSI sessions.

• Due to a Microsoft limitation, if a pass-through disk is down, the virtual machine to which the disk is assigned cannot be rebooted.

To reboot the virtual machine, you must bring the pass-through disk back up, or remove it from the virtual machine using Microsoft Hyper-V Manager or another Microsoft remote server management tool.

- If OnCommand Unified Manager Core Package role-based access control is enabled on a Hyper-V parent host, you cannot dynamically add or remove pass-through disks with SnapDrive in a Hyper-V virtual machine.
- Windows Server 2008 Server Core does not support the ability to dynamically add and remove Hyper-V pass-through disks, because Server Core does not support Windows Communication Foundation (WCF) Web services that are required to dynamically add and remove pass-through disks.
- In a clustered environment with multiple HA VMs residing in different clustered disks, where the owner-node is not same, parallel LUN management operations can lead to an inconsistent SnapDrive state.

Managing Snapshot copies

You can use SnapDrive for Windows to create, schedule, restore, and delete Snapshot copies as well as some other Snapshot copy management tasks.

Reasons for creating Snapshot copies

You use SnapDrive to ensure that you create consistent Snapshot copies in the event you need to restore a LUN from that copy.

Snapshot operations on a single LUN actually make a Snapshot copy of all the LUNs on the volume. Because a storage system volume can contain LUNs from multiple hosts, the only consistent Snapshot copies are those of LUNs connected to the host that created the SnapDrive Snapshot copy. In other words, within a Snapshot copy, a LUN is not consistent if it is connected to any host other than the one that initiated the Snapshot copy. (This is why you are advised to dedicate your storage system volumes to individual hosts.) Therefore, it is important to back up a LUN using a SnapDrive Snapshot copy rather than using other means, such as creating Snapshot copies from the storage system console.

Additionally, as part of the SnapDrive Snapshot copy process, the file system (NTFS) is flushed to disk and the disk image in the Snapshot copy is in a consistent state. This consistency cannot be ensured if the Snapshot copy was created outside the control of SnapDrive (that is, at the storage system console, or using either On Command System Manager, rsh, or by backing up the LUN file in the active file system.)

Restrictions on Snapshot copy creation

You must keep in mind some restrictions for creating Snapshot copies.

- You can keep a maximum of 255 Snapshot copies with Data ONTAP. After the number of Snapshot copies has reached the limit, the Snapshot Create operation fails, and you must delete some of the old Snapshot copies before you can create any more.
- SnapDrive does not support Snapshot copies that are created from the storage system console, because such a practice can lead to inconsistencies within the NTFS file system. Therefore, use only SnapDrive to create Snapshot copies of LUNs.
- You cannot create a Snapshot copy of a LUN connected to a Snapshot copy.
- SnapDrive automatically turns off the Snapshot copy schedule on a storage system volume that stores LUNs, so that the storage system does not create automatic Snapshot copies.
- When Snapshot copies follow the Data ONTAP schedules Snapshot copy naming convention, they are unavailable in SnapDrive for Windows. To enable access to these Snapshot copies, enable inconsistent Snapshot copies.

Creating a Snapshot copy

You should always use SnapDrive to create Snapshot copies of LUNs to ensure that Snapshot copies are consistent.

Before you begin

The following requirements must be met in order to successfully create Snapshot copies using SnapDrive:

- You must create Snapshot copies through the SnapDrive MMC snap-in or through sdcli.exe, rather than the storage system console or the volume Snapshot copy schedule on the storage system. This is because SnapDrive must first flush NTFS so that the LUN is consistent at the moment the Snapshot copy is made. This ensures the usability of the LUN in the Snapshot copy directory.
- Snapshot names must be created using US-ASCII characters only, even when using non-ASCII operating systems.

Note: The SnapDrive service can perform only one task at a time. If you schedule multiple tasks to start at exactly the same time, the first will proceed, and SnapDrive will queue the others until the first task either succeeds or times out.

Steps

- 1. Perform the following actions to get to the Create Snapshot menu item:
 - a) In the left MMC pane, select the instance of SnapDrive you want to manage.
 - b) Double-click Disks.
 - c) Double-click the disk for which you want to create a Snapshot copy.
 - d) Select Snapshots.
 - e) From the menu choices at the top of MMC, navigate to Action > Create Snapshot.

The Create Snapshot text box is displayed.

- 2. In the Create Snapshot text box, perform the following actions:
 - a) Type an easy-to-interpret name for the Snapshot copy.

Example

expenses_db_15Jan07_4pm

Note: Snapshot copy names must be created using US-ASCII characters only, even when using non-ASCII operating systems.

b) Click OK.

Note: All LUNs in the volume attached to this host are included in the Snapshot copy, as well as LUNs from other hosts that are also in the volume. The only LUNs that will be consistent, however, are those attached to the host initiating the Snapshot copy request.

Result

Your Snapshot copy is created.

Details about the Snapshot copy appear at the bottom panel of the center MMC pane.

Scheduling Snapshot copies

You can create a Snapshot copy schedule to ensure that SnapDrive creates Snapshot copies at intervals appropriate to your environment.

About this task

All steps except Step 1 in the following procedure are performed using the Scheduled Task Wizard, a Windows task scheduling tool available on your Windows server.

Steps

1. Create a .bat file containing the following command on the Windows host on which you are scheduling Snapshot copies:

```
sdcli snap create [-m MachineName] -s SnapshotName -D DriveLetterList
[. . .] [-x]
```

Example

The following example creates a Snapshot copy named Jun_13_07 for each volume containing one or more of the LUNs mapped to the specified drives (that is, J:, K:, and L:). The Snapshot copies created are consistent for all LUNs contained by those volumes.

sdcli snap create -s Jun_13_07 -D j k l

- 2. Click Start Menu > Settings > Control Panel > Scheduled Tasks.
- 3. Double-click Add Scheduled Task.

The Scheduled Task Wizard is launched.

- 4. In the Scheduled Task Wizard, click Next.
- 5. Click Browse and locate the batch (.bat) file you created in Step 1.
- 6. Select the batch file and click **Open**.
- 7. In the wizard page, select from the list of frequencies, and then click Next.
- **8.** In the next wizard page, type a start time and complete the detailed frequency parameters. The option details displayed on this page vary depending on the Snapshot copy frequency you picked on the previous page.
- **9.** In the next wizard page, type the user name (the administrator account name and password, repeated for confirmation), and then click **Next**.

10. Click Finish.

Support for FlexClone volumes in SnapDrive

By default, if the prerequisites are met, SnapDrive uses FlexClone technology to connect to LUNs in a Snapshot copy. The use of FlexClone technology by SnapDrive is helpful for conducting tests or for verifying data on a live SnapMirror destination.

SnapDrive connects a host to a LUN in a Snapshot copy in read/write mode by mounting a LUN that is stored in a Snapshot copy, or by connecting to a clone of a FlexVol volume using a FlexVol volume clone (FlexClone).

Note: FlexClone operations might fail if the virtual storage server configuration is incomplete. In clustered Data ONTAP, the aggregate must be assigned to the virtual storage server for successful SnapDrive operations.

Prerequisites for using FlexClone volumes with SnapDrive

There are several prerequisites that must be met in order for SnapDrive to use FlexClone volumes.

- Your storage system must be running Data ONTAP 7.2.7 or later.
- Only FlexVol volumes can be used to create FlexClones.
- There must be enough space available on the aggregate to create a non-space-reserved FlexVol volume (volume guarantee=none).
- FlexClone must be licensed on your storage system.

About read/write connections

If FlexClone volumes are not available because the prerequisites for their use have not been met, SnapDrive uses a read/write connection to a LUN in a Snapshot copy that is actually a connection to a special type of LUN.

Read/write connection to LUNs in a Snapshot copy have the following properties:

- It is backed by a LUN in a Snapshot copy.
- It resides in the active file system and always has an .rws extension.
- When the host reads data from this LUN, it receives data that is in the LUN that is in the Snapshot copy.
- When the host writes data to this LUN, the data is written to the LUN with the .rws extension.
- When the host reads data that has been written to the LUN with the .rws extension, that data is received from the LUN with the .rws extension.

For details, see your Data ONTAP documentation.

Snapshot copy cautions

Keep the following points in mind when working with Snapshot copies and LUNs that are backed up by a Snapshot copy:

• Information written to the .rws file is temporary; SnapDrive deletes the .rws file when you disconnect.

- You cannot merge the data written to the .rws file with the data in the Snapshot copy referenced by the .rws file.
- You cannot delete a Snapshot copy that is in use by a LUN backed by a Snapshot copy.
- You can connect to the LUN Snapshot copy only by using read/write mode and a LUN that is backed by a Snapshot copy.
- You should avoid creating a Snapshot copy of a LUN backed by a Snapshot copy. Doing so will lock the Snapshot copy backing the LUN until the newer Snapshot copy—and all Snapshot copies of that LUN—are deleted.

Connecting to a LUN in a Snapshot copy

You can connect to a LUN in a Snapshot copy using either a FlexClone volume or a read/write connection to a LUN in a Snapshot copy, depending on which version of Data ONTAP you have installed on your storage system.

Before you begin

- You FlexClone license is enabled.
- You have set cifs.show_snapshot to on and vol options nosnapdir is set to off on your storage system.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then expand **Disks** and select the disk you want to manage.
- 2. Expand the LUN whose Snapshot copy you want to connect to and then click **Snapshot Copies** and the name of the Snapshot copy you want to connect to.
- **3.** From the menu choices at the top of MMC, navigate to Action > Connect Disk to launch the Connect Disk wizard.
- 4. In the Connect Disk wizard, click Next.
- 5. In the Provide a Storage System Name, LUN Path and Name panel, click Next.
- 6. In the Select a LUN Type panel, Dedicated is automatically selected; click Next.
- 7. In the **Select LUN Properties** panel, either select a drive letter from the list of available drive letters or type a volume mount point for the LUN you are connecting to and then click **Next**.

When you create a volume mount point, type the drive path that the mounted drive will use: for example, G:\mount_drivel\.

- 8. In the Select Initiators panel, select the FC or iSCSI initiator for the LUN you are connecting to and click Next.
- **9.** In the **Select Initiator Group management** panel, specify whether you will use automatic or manual igroup management:

If you specify	Then do this
Automatic igroup management	Click Next.
	SnapDrive uses existing igroups, one igroup per initiator, or, when necessary, creates new igroups for the initiators you specified in the Select Initiators panel.
Manual igroup management	Click Next and then perform one of the following actions:
	a. In the Select Initiator Groups panel, select from the list the igroups to which you want the new LUN to belong.
	Note: A LUN can be mapped to an initiator only once.
	Click Manage Igroups and for each new igroup you want to create, type a name in the Igroup Name text box, select initiators from the initiator list, click Create , and then click Finish to return to the Select Initiator Groups panel.
	b. Click Next.

10. In the Completing the Connect Disk Wizard panel, perform the following actions.

- a) Verify all the settings
- b) If you need to change any settings, click **Back** to go back to the previous wizard pages.
- c) Click Finish.

Result

The newly connected LUN appears under Disks in the left MMC pane.

Data protection through archiving and restoring Snapshot backup copies

A good way to protect and retain data is to archive the SnapDrive Snapshot copies of the LUNs to offline, offsite media, such as near-line technology or alternate storage methods.

The practice of archiving Snapshot copies is particularly beneficial for disaster recovery.

What to back up

When archiving backups, it is important that you select the LUNs that are not in the active file system. The disks in the active file system are not consistent and, therefore, do not result in reliable backups. You must also select the Snapshot copies of the LUNs when creating backups.

Ways to archive SnapDrive backups

You can use the Data ONTAP dump command or an NDMP-based backup application to archive the Snapshot backup copies of your LUNs.

Process for restoring LUNs from archival media

First, restore the LUN file from your archive media to the active file system. After the file is restored, use the SnapDrive management interface to connect to the LUN file using its original drive letter.

Note: You cannot use CIFS-based or NFS-based backup products to archive the Snapshot copies of your LUNs.

For more information about LUN backups, see the *Data ONTAP SAN Administration Guide for 7-Mode.* For more information about how to perform a recovery from an offline archive, see your backup application software documentation.

Note: Further steps might be required to bring online data recovered in LUN files. This is true for all SnapManager products. For more information about recovering LUNs using SnapManager, see your SnapManager product documentation.

How LUN restoration from Snapshot copies works

When you restore a LUN from a Snapshot copy, the LUN reverts to the state it was in when the Snapshot copy was made: the restore operation overwrites all data written to the LUN since the Snapshot copy was made.

A LUN restore recalls a selected Snapshot copy. During a restore, the entire LUN drive is restored from the Snapshot copy. For a restore to succeed, no open connections can exist between the host machine (or any other application) and the files in the LUN.

Note: If it is necessary to restore a LUN from a Snapshot copy made before the LUN was resized, SnapDrive for Windows automatically resizes the LUN to the size of the Snapshot copy and performs the restore. Such a restore causes the loss of any data added to the LUN after it was resized, and it can damage virtual machines (such as Hyper-V VMs) or applications if they are running on the LUN during the restore.

About the Data ONTAP LUN clone split (rapid LUN restore) feature

SnapDrive uses the LUN clone split (rapid LUN restore) feature of Data ONTAP when restoring a LUN.

A LUN clone is a point-in-time, writable copy of a LUN in a Snapshot copy. After the clone is created, all read/write operations are made on the clone and read/write operations are no longer made on the original LUN.

A LUN clone shares space with the original LUN in the backing Snapshot copy. The clone does not require additional disk space until changes are made to it. When Data ONTAP splits the clone from the backing Snapshot copy, Data ONTAP copies the data from the Snapshot copy, and copies the blocks from the original LUN, to the clone. After the splitting operation, the clone becomes a regular LUN, and the original LUN is deleted by Data ONTAP.

Note: If you do not have enough disk space for both the clone and the original LUN, the split will not be initiated and no LUN restoration can occur.

Benefit of using rapid LUN restore

When rapid LUN restore, or LUN cloning, is used by SnapDrive, the clone is split from the backing Snapshot copy in the background, and the restored LUN is available to the Windows host for I/O operations within a few seconds.

Note: You might not be able to delete the Snapshot copy after a restore operation due to the clone split operation. You can delete the Snapshot copies after the clone split operation is complete.

Restoring a LUN from a Snapshot copy

SnapDrive restores a LUN using the rapid LUN restore feature.

Before you begin

• Shut down all resources directly or indirectly dependent on the LUN.

Make sure that the LUN is not being used by the Windows file system or any other process, and that no user has the LUN open in Windows Explorer. Shut down any application that is using the LUN.

Attention: Make sure that the Windows Performance Monitor (perfmon) is not monitoring the LUN.

Note: Make sure that virtual machines, Microsoft Exchange, or any other applications are no longer running on a LUN before you restore that LUN from a Snapshot copy.

Steps

- 1. Perform the following actions:
 - a) In the left MMC pane, select the instance of SnapDrive you want to manage.
 - b) Double-click **Disks** to display all available disks.
 - c) Select the LUN that you want to restore and double-click it to display all the Snapshot copies list.
 - d) Select the Snapshot copy from which to restore the LUN.
- 2. In the menu choices at top of MMC, click Action > Restore Disk.

Note: You can only restore a Snapshot copy that is consistent with the active file system. Inconsistent Snapshot copies are not available for restoration and are grayed out.

The Restore Snapshot Copy panel is displayed.

3. In the **Restore Snapshot Copy** panel, click **Yes** to restore the LUN from the Snapshot copy you selected.

Attention: Do not attempt to manage any Windows cluster resources while the restore is in progress.

Checking LUN restore status

Check whether LUN restoration has completed by viewing the Restore Status field in the SnapDrive MMC.

Steps

- 1. Perform the following actions:
 - a) In the left MMC pane, select the instance of SnapDrive you want to manage.
 - b) Double-click Disks.
- 2. In the center MMC pane, locate the name of the disk you are restoring. The status is displayed in the lower panel of the center MMC pane.

Note: You can also check the status of a LUN restore using the disk list command of the sdcli.exe utility.

Result

If a restore is in progress, SnapDrive will display the percentage completed, otherwise; the status will display Normal.

Command-line volume-based Snapshot backup restoration with SnapDrive

You can perform volume-based Snapshot copy restoration (VBSR) using the SnapDrive sdcli utility.

SnapDrive supports volume-based Snapshot backup restoration on volumes containing only unconnected LUNs, and breaks the SnapMirror connection if the operation is performed on a live SnapMirror destination.

Volume restore functions are currently available only through the sdcli.exe utility.

You can perform VBSR when you have a Snapshot backup of a read/write volume, and the Snapshot backup was created with all LUNs connected to the same host. When you perform the VBSR operation, you must ensure that all your LUNs are disconnected.

Support for restoring data at the file level

You can use the SnapDrive command-line interface to restore one or more files from a Snapshot copy.

You can use the file-level restore feature in the following cases:

- To restore generic files on a LUN from its corresponding Snapshot copy A file being restored on a LUN can be a virtual hard disk, a database file, or any large, generic file, provided that no applications or systems are affected by the restoration of those files.
- To restore the individual database files, provided that the administrator ensures that the files being restored do not cause the database to become inconsistent

Attention: SnapDrive guarantees data consistency of files restored from a consistent Snapshot copy; however, application consistency is outside the function of SnapDrive for Windows. Files restored using the file-level restore operation might result in application inconsistency. Use file-level restoration with caution, following the recommended practices for the operating system or applications using the files. The file-level restore operation must be used by experienced administrators with full knowledge of the operating system and applications using the files.

Deleting a Snapshot copy

You should delete older SnapDrive Snapshot copies to keep the number of stored Snapshot copies less than the limit of 255 for Data ONTAP and to free space on the storage system volume. Be sure to delete old Snapshot copies before the hard limit is reached; otherwise, subsequent Snapshot copies could fail. Even before the Snapshot copy limit is reached, a Snapshot copy fails if insufficient reserved space for it remains on the disk.

Steps

- 1. Perform the following actions:
 - a) In the left MMC pane, select the instance of SnapDrive you want to manage.
 - b) Double-click Disks.
 - c) Select the LUN with the Snapshot copy you want to delete.
- 2. In the right MMC pane, select the Snapshot copy you want to delete.

Note: You can only delete a Snapshot copy that is consistent with the LUN. Inconsistent Snapshot copies are not available for deletion.

3. From the menu choices on top of MMC, click Action > Delete.

The Delete Snapshot panel is displayed.

4. In the Delete Snapshot panel, click Yes to delete the Snapshot copy you selected.

Note: If you get an error message stating that the Snapshot copy is busy or cannot be deleted, it is likely that the Snapshot copy is in use by a LUN that is backed by a Snapshot copy.

Problems deleting Snapshot copies due to busy snapshot error

If you attempt to delete a Snapshot copy and you get an error message saying that the Snapshot copy is busy and cannot be deleted, you either have a Snapshot copy that was taken of a LUN backed by another Snapshot copy or the Snapshot copy backed LUN is still connected.

If you have a Snapshot copy that was taken of a LUN backed by another Snapshot copy, you need to delete the newer Snapshot copy before the older Snapshot copy, the Snapshot copy backing the LUN, can be deleted.

If the LUN backed by a Snapshot copy is still connected, disconnect it.

Attention: Data on the LUN is no longer available when you disconnect it. If you need the data on the LUN, back it up or copy it to another LUN before you disconnect it.

To see if you have busy Snapshot copies, you can view your application event log in the Event Viewer to check for messages related to busy Snapshot copies. For more information about deleting busy Snapshot copies, see the *Data ONTAP SAN Administration Guide for 7-Mode* for your version of Data ONTAP.

Managing space on storage system volumes

SnapDrive for Windows enables you to manage space on your storage system volumes.

What SnapDrive fractional space reservation monitoring does

Fractional space reservation monitoring in SnapDrive for Windows enables you to monitor fractional space reserved for LUNs on a storage system volume.

To monitor the fractional space reserved on your storage system from your Windows host, SnapDrive lets you perform the following tasks:

- Set fractional space reservation thresholds for volumes containing LUNs
- Set rate-of-change percentage between two Snapshot copies or between a Snapshot copy and the active file system of the storage system volume
- Monitor space that can be reclaimed by deleting a Snapshot copy
- Set monitor polling interval
- Enable and disable e-mail notification

For more information about fractional space reservation, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

Configuring space reservation monitoring

You can configure how SnapDrive for Windows monitors the fractional space reserved for LUNs on a storage system volume.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then select **Disks**.
- 2. From the menu choices at the top of MMC, navigate to Action > Properties.
- 3. In the Disks Properties window, select the Space Reservation Monitor tab.
- 4. In the Space Reservation Monitor panel, perform the following actions:
 - a) Click to deselect the Disable Space Reservation Monitoring check box.
 - b) Type a value in the Monitor Time Interval field, in minutes.

Values can be between 0 (disable) and 10,080 minutes (7 days).

- c) Under the Space Reservation Monitor Settings tree, select the storage system and volume name.
- d) Type a value for the Reserve Available percentage threshold.

- e) Type a value for the Rate of Change threshold and choose MB, GB, or TB for the Unit.
- f) Select the **Alert** check box if you want to be notified if this condition occurs.
- 5. Click OK or Apply to save your changes.
- 6. Click OK.

Using the storage access control tool to enable thinly provisioned LUNs

You can use storacl.exe to set the space reservation option for any volume on a storage system. The space reservation option determines whether LUNs are fully provisioned or thinly provisioned. Thin provisioning using storacl.exe is not supported with clustered Data ONTAP 8.1.

Before you begin

- You have run storacl.exe from a Windows host and created the ThinProvision.xml file on your storage system in the /etc directory of the root volume.
- You have enabled HTTPS using options ssl.enable and secureadmin setup ssl.
- You are logged in to the storage system as root.

About this task

You can run spacereserve help at the storacl.exe prompt to view a list of additional commands and how to use them.

Setting space reservation to True indicates that LUNs on the volume are fully provisioned, and setting space reservation to False indicates that LUNs are thinly provisioned. The default value is set to True on storage system volumes that do not have space reserve set in the ThinProvision.xml file.

Steps

- 1. Run storacl.exe from your Windows host.
- 2. Run spacereserve help and determine which additional commands you must run to enable thinly provisioned LUNs.

Setting space reservation to True indicates that LUNs on the volume are fully provisioned, and setting space reservation to False indicates that LUNs are thinly provisioned.

A file named ThinProvision.xml is created on the storage system to store the space reservation settings you specify.

What Space Reclaimer does

Space Reclaimer is a SnapDrive for Windows feature that optimizes LUN space by marking newly freed space that is visible to NTFS so that it is also seen as available by Data ONTAP. Using Space Reclaimer lessens the disparity in available space that is reported by the two systems.

When files are deleted or modified on a LUN, the space is tracked by NTFS, but since this information is not communicated to the Data ONTAP file system, a disparity can grow between the

available space reported by a SnapDrive host and a storage system. Space Reclaimer ensures that newly freed blocks are marked as available on the storage system.

You can use Space Reclaimer on traditional LUNs and on VMDK files attached as virtual disks using NFS datastores.

Space Reclaimer requirements and restrictions

Before you can use Space Reclaimer, it is important that you understand some restrictions on how you can use it and information about how to run it for optimum storage performance.

- For optimum storage performance, run Space Reclaimer as often as possible and until the entire NTFS file system has been scanned.
- Space reclamation is a time-consuming operation; therefore, it is best to run Space Reclaimer on your NTFS volume when there is a large amount of unused deleted space.
- The space reclamation process is CPU intensive, so run Space Reclaimer when storage system and Windows host usage is low: for instance, at night.
- Do not run disk defragmentation at the same time that Space Reclaimer is running, because doing so can slow the disk reclamation process.
- In Microsoft Cluster Server (MSCS) configurations, you can start Space Reclaimer from the owner node only.
- Although Space Reclaimer reclaims nearly all space from newly freed blocks, it does not reclaim 100 percent of the space.
- When you are running Space Reclamation in Windows Server 2012 environments, ensure that you are using the recommended version of Windows Host Utilities (WHU) for the your version of Data ONTAP.

See the N series interoperability matrix website (accessed and navigated as described in *Websites* on page 12) for the latest supported configurations.

Starting Space Reclaimer

You can start space reclamation using SnapDrive for Windows MMC snap-in.

Before you begin

To use this feature, you must have Data ONTAP 7.2.7 or later installed on your storage system.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then expand **Disks** and select the disk you want to manage.
- 2. From the menu choices at the top of MMC, navigate to Action > Start Space Reclaimer.
| If SnapDrive
detects that there
is | Then a window appears |
|--|--|
| Space to reclaim | Confirming that the LUN space can be optimized. |
| | Continue to the next step. |
| No space to
reclaim | Notifying you that you do not need to run Space Reclaimer on the selected disk. |
| | Click Cancel to exit the Confirm Space Reclamation on Disk window, or go to the next step to continue with space reclamation. |
| | Note: When you run Space Reclaimer on a disk that SnapDrive has determined has no space to reclaim, the space reclamation process can still take as long to complete as it would for a disk that does have reclaimable space. This is because SnapDrive performs NTFS block comparisons and analyzes disk infrastructure regardless of whether there is space to reclaim. |

3. In the **Confirm Space Reclamation on Disk** window, limit the amount of time Space Reclaimer runs on a LUN by selecting the "Limit (in minutes) Space Reclamation operation" check box.

In the space provided by the check box, type the number of minutes you want Space Reclaimer to run on the LUN. By default, Space Reclaimer runs until the LUN is optimized.

4. Click **OK** to continue.

The space reclamation process runs in the background. You can monitor the Space Reclaimer progress for the selected LUN by watching the status bar in the Details pane in MMC.

Note: SnapDrive might indicate that it is necessary to run Space Reclaimer again after it has successfully completed a space reclamation process. This can happen if data is deleted from a LUN while Space Reclaimer is running. It is strongly recommended that the space reclamation process is performed when there is little or no activity on both the storage system and the Windows host.

Stopping Space Reclaimer manually

You can stop space reclamation using SnapDrive for Windows MMC snap-in.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then expand **Disks** and select the disk you want to manage.
- 2. From the menu choices at the top of MMC, navigate to Action > Stop Space Reclaimer.

Reasons for SnapDrive to automatically stop Space Reclaimer

SnapDrive for Windows automatically stops Space Reclaimer in several instances.

• During any LUN management operation on a LUN running Space Reclaimer, including LUN disconnect and LUN delete operations

- · During any Snapshot copy management operation except Snapshot copy rename and delete
- · On all LUNs of the same storage system volume during Snapshot copy creation
- If the SnapDrive service is stopped
- During LUN restore operations for any volume mount points directly or indirectly mounted (cascading) from the disk being restored
- During Windows host cluster failover If a host cluster failover operation occurs on a Windows Server 2008 failover cluster while Space Reclaimer is running on a shared LUN, space reclamation will stop running on that LUN.
- During any MPIO path management operations, including adding or removing an initiator or active path selection

Enabling space reclamation on pass-through LUNs

Before you run Space Reclaimer on a pass-through LUN, you must disable SCSI Command Descriptor Block filtering on your parent host.

Steps

- 1. Create a new Hyper-V VM.
- 2. From the PowerShell command prompt, run the following commands to disable SCSI Command Descriptor Block filtering:

```
PS C:\Users\administrator> $HyperVGuestName = "Vm Name";
PS C:\Users\administrator>
PS C:\Users\administrator> $VMManagementService = qwmi -Class
"Msvm VirtualSystemManagementService" -Namespace "root\virtualization"
PS C: Users \administrator>
PS C:\Users\administrator> $Vm = gwmi -Namespace "root
\virtualization"
-Query "Select * from Msvm ComputerSystem Where
ElementName='$HyperVGuestName'"
PS C:\Users\administrator>
PS C:\Users\administrator> $SettingData = gwmi -Namespace "root
\virtualization"
-Query "Associators of {$Vm} Where
ResultClass=Msvm VirtualSystemGlobalSettingData
AssocClass=Msvm ElementSettingData"
PS C:\Users\administrator>
PS C:\Users\administrator> $SettingData.AllowFullSCSICommandSet =
$true
PS C:\Users\administrator>
PS C:\Users\administrator>
$VMManagementService.ModifyVirtualSystem($Vm, $SettingData.GetText(1))
PS C:\Users\administrator>
```

- 3. Restart your VM.
- 4. Create your pass-through LUN.
- 5. Run Space Reclaimer on your new pass-through LUN.

Requirements for running Space Reclaimer on CSVs

Before you run Space Reclaimer on CSVs, you should be aware of the requirements for doing so.

- When you run Space Reclaimer on CSVs, you must do so from the owner node.
- If the owner node is changed during Space Reclaimer operations, the Space Reclaimer operation is aborted.

Using SnapDrive in VMware environments

You can use SnapDrive in VMware environments to create and manage LUNs, manage Snapshot backups, manage space on the storage system, and manage Microsoft clusters on ESX.

VMware support

SnapDrive for Windows provides LUN provisioning and Snapshot copy management support with VMware ESX 3.0.2 or later guest OS on x86, and x64 platforms when using either the Microsoft iSCSI Software Initiator 2.04 or later, FC HBAs, or ESX iSCSI software initiators or iSCSI HBAs.

SnapDrive supports the following VMware guest OS configurations:

- Windows Server 2008 and 2012, on x86 and x64 platforms
- Microsoft cluster configurations up to a maximum of eight nodes supported on VMware when using the Microsoft iSCSI Software Initiator, or up to two nodes using FC
- A maximum of 56 RDM LUNs with four LSI Logic SCSI controllers for normal RDMS; 48 RDM LUNs with three LSI Logic SCSI controllers on VMware VM MSCS box-to-box SnapDrive for Windows configuration
- Paravirtual SCSI (PVSCSI) adapters, with some additional requirements:
 - PVSCSI adapters require ESX/ESXi 4.0 or later.
 - The PVSCSI controller must exist before the LUN is created.

VMware ESX server-related limitations

SnapDrive is supported on VMware ESX server; however, there are some limitations you must keep in mind.

• Installing SnapDrive on a Microsoft cluster on virtual machines using ESX credentials is not supported.

Use the vCenter credentials when installing SnapDrive on clustered virtual machines.

- RDM LUNs greater than 2 TB are not supported either in a VMFS 3.0 datastore or if the ESX or ESXi server version is earlier than 5.0.
- MPIO is present on the ESX and is not required on the VMware guest OS.
- iSCSI and FC initiators are not supported together on a VMware guest OS.

Enabling and disabling vCenter or ESX logon from SnapDrive MMC

You can use SnapDrive MMC to enable and disable VMware vCenter or ESX logon settings after SnapDrive is installed on your Windows host.

About this task

Keep the following information in mind when setting vCenter or ESX logon from SnapDrive MMC:

- You cannot disable vCenter or ESX logon from SnapDrive MMC when RDM LUNs are present. You must disconnect or delete RDM LUNs before you can disable ESX logon.
- If you migrate a virtual machine from one ESX server to another, you must configure SnapDrive with vCenter account information.

Steps

- 1. In the left MMC pane, select the instance of SnapDrive for which you want to enable or disable vCenter or ESX logon.
- From the menu choices at the top of MMC, navigate to Action > vCenter Server or ESX Server Login Settings.

The vCenter Server or ESX Server Log On window is displayed.

- **3.** To enable vCenter or ESX logon, in the vCenter Server or ESX Server Log On window, select the "Enable vCenter Server or ESX server settings" check box.
- **4.** Type the IP address or hostname, user name, and password for the vCenter or ESX to which you want to log in.
- 5. Click OK.
- 6. To disable vCenter or ESX logon, complete Steps 1 and 2 and then clear the "Enable vCenter Server or ESX server settings" check box.

The vCenter or ESX settings are unavailable (grayed out).

Note: When you disable vCenter or ESX settings, SnapDrive cannot display WWPNs for FC HBAs on the ESX server.

Minimum vCenter privileges required for SnapDrive operations

To perform RDM operations in a guest OS, you must have minimum vCenter privileges.

You must have the following minimum privileges set on the host:

- Datastore: Remove File
- Host: Configuration -> Storage Partition Configuration
- Virtual Machine: Configuration

You must assign these privileges to a role at the Virtual Center Server level. The role to which you assign these privileges cannot be assigned to any user without root privileges.

After you assign these privileges, you can install SnapDrive on the guest OS.

Requirements for VMware vMotion support

Before SnapDrive can support VMware vMotion, which enables the live migration of running virtual machines from one physical machine to another without interrupting service to those machines, you must ensure that vMotion requirements have been met.

The following vMotion requirements must be met to use vMotion with SnapDrive:

- You must use VMware vCenter instead of the ESX during SnapDrive installation.
- If SnapDrive was installed to communicate directly with the ESX, you must modify settings using vCenter Server or ESX Server login Settings in the SnapDrive MMC.
- You must manually create an igroup that has all WWPNs from the source and the destination ESX.

You must use the same igroup for all RDM LUN create and connect operations.

Note: When you perform a vMotion operation, the RDM LUN validation might fail. Perform an HBA rescan from the virtual infrastructure client and retry the operation.

Creating LUNs in VMware environments

You can use SnapDrive to create LUNs in VMware environments.

Creating an RDM LUN on a guest OS

You can use SnapDrive to create FC, iSCSI, or ESX iSCSI accessed RDM LUNs on a guest OS.

Before you begin

- Create the dedicated volumes to hold your LUNs on the storage system.
- Verify that the FC or iSCSI service has been started on the storage system.
- Before creating a LUN in a VMware guest OS, you must manually create initiator groups by using OnCommand System Manager or at the storage system console.
 When you create a shared FC or iSCSI RDM LUN, you must choose an initiator from each ESX server to create a single initiator group automatically using SnapDrive.

Steps

- 1. Perform the following actions to launch the Create Disk wizard:
 - a) Select the SnapDrive instance for which you want to create a disk.
 - b) Select Disks.
 - c) From the menu choices at the top of MMC, navigate to Action > Create Disk.

The Create Disk Wizard is launched.

2. In the Create Disk Wizard, click Next.

The Provide Storage System Name, LUN Path and Name panel is displayed.

- **3.** In the **Provide a Storage System Name, LUN Path and Name** panel, perform the following actions:
 - a) In the **Storage System Name** field, type the storage system name where the LUN will be created or select an existing storage system using the pull-down menu.
 - b) In the **LUN Path** field, type the LUN path or select the path on the storage system you added in Step a.

c) In the LUN Name field, enter a name for the LUN and click Next.

The Select a LUN Type panel is displayed.

- 4. In the Select a LUN Type panel, select Dedicated, and then click Next.
- 5. In the Select LUN Properties panel, either select a drive letter from the list of available drive letters or type a volume mount point for the LUN you are creating. When you create a volume mount point, type the drive path that the mounted drive will use: for example, G: \mount_drive1\.

Note: The root of the volume mount point must be owned by the node on which you are creating the new disk.

Note: You can create cascading volume mount points (one mount point mounted on another mount point); however, in the case of a cascading mount point created on an MSCS shared disk, you might receive a system event warning indicating that disk dependencies might not be correctly set. This is not the case, however, as SnapDrive will create the dependencies and the mounted disks will function as expected.

- 6. While still in the Select LUN Properties panel, complete the following actions:
 - a) Click Limit or Do not limit for the option labeled "Do you want to limit the maximum disk size to accommodate at least one snapshot?"

If you keep the default, **Limit**, which is the recommended option, the disk size limits displayed are accurate only when they first appear on the Select LUN Properties panel. When this option is selected, the following actions might interfere with the creation of at least one Snapshot copy:

- Changing the option to **Do not limit** and using SnapDrive to create an additional LUN in the same storage system volume
- Creating a LUN in the same stroage system volume without using SnapDrive
- Storing data objects other than LUNs on this storage system volume
- b) Select a LUN size, which must fall within the minimum and maximum values displayed in the panel.
- c) Click Next.

If the settings on the storage system volume or qtree on which you are creating the LUN do not allow SnapDrive to proceed with the create operation, the Important Properties of the Storage System Volume panel is displayed, as described in Step 7. Otherwise, Step 7 is skipped.

7. The **Important Properties of the Storage System Volume** panel displays the settings that will be used for the volume or qtree you specified in Step 4 of this procedure.

SnapDrive requires the storage system volume containing LUNs to have the following properties:

- create_ucode = on
- convert_ucode = on
- snapshot_schedule = off

Note: SnapDrive cannot proceed to create a LUN unless these settings are configured as shown. Therefore, you must accept these settings.

Note: The create_ucode and convert_ucode volume options are no longer used, but they are set to maintain backwards compatibility with earlier versions of SnapDrive.

Click Next.

The Select Initiators panel is displayed.

8. In the Initiator List pane, select an initiator for the LUN you are creating.

If you have MPIO installed, you have the option to select several FC initiators.

Note: You cannot select both iSCSI and FC initiators when creating a LUN on a guest OS.

9. Click Next.

The Select Initiator Group Management panel is displayed.

10. In the Select Initiator Group Management panel, specify whether you will use automatic or manual igroup management. If you select automatic igroup management, SnapDrive uses existing igroups or, when necessary, creates new igroups for the initiator you specified in Step 8. If you select manual igroup management, you manually choose existing igroups or create new ones as needed.

If you specify	Then
Automatic igroup management	Click Next.
	SnapDrive uses existing igroups, one igroup per initiator, or, when necessary, creates new igroups for the initiators you specified in Step 8.
Manual igroup management	Click Next, and then perform the following actions:
	a. In the Select Initiator Groups panel, select from the list the igroups to which you want the new LUN to belong.
	Note: A LUN can be mapped to an initiator only once.
	OR
	Click Manage Igroups and for each new igroup you want to create, type a name in the Igroup Name text box, select initiators, click Create , and then click Finish to return to the Select Initiator Groups panel.
	b. Click Next.
	Note: The Next button will remain unavailable until the collection of selected igroups contains all the initiators you selected in Step 8.

You are done with igroup management.

11. In the Select a Datastore panel, perform the following steps.

If your virtual machine resides on a	Th	en
VMFS datastore	a.	Choose either Store with Virtual Machine , which is the default, or choose Specify datastore to select a different VMFS datastore for your virtual machine disk format (VMDK) file to reside.
	b.	Click Next.
NFS datastore	a.	Choose Specify datastore to select a different VMFS datastore for your VMDK file to reside. The Store with Virtual Machine option is unavailable because a VMDK file cannot be stored on an NFS datastore.
	b.	Click Next.

12. In the Completing the Create Disk Wizard panel, perform the following actions:

a) Verify all the settings.

If you need to change any settings, click **Back** to go back to the previous Wizard panels.

b) Click Finish.

Disk creation might take several seconds to complete. SnapDrive displays disk creation status in the lower panel of the center MMC pane.

Troubleshooting RDM LUN creation

If you experience errors creating RDM LUNs, you should be aware of some of the common errors and workarounds.

Error message

Failed to create disk in virtual machine, Failed to Map virtual disk: File [datastore] *path_name*.vmdk was not found.

Problem

You might encounter this error when you attempt to create an RDM LUN with ESX Software Initiator on a VM with the VM name greater than 33 characters.

You have several options to work around this issue.

Workaround 1

Manually create the same directory inside the datastore.

Workaround 2

Rather than selecting your datastore with the Store with Virtual machine option, select the datastore in which you intend to create the RDM LUN. When you create your RDM LUN, use the same datastore you just selected.

Workaround 3

Configure SnapDrive VirtualCenter or ESX Server login settings with the VirtualCenter credentials.

Using FC RDM LUNs in a Microsoft cluster

You can use SnapDrive to manage a Microsoft cluster using RDM LUNs, but you must first create the shared RDM quorum and shared storage outside of SnapDrive, and add the disks to the virtual machines in the cluster.

Requirements for using FC RDM LUNs in a Microsoft cluster

SnapDrive provides support for Microsoft clusters using FC RDM LUNs on two different virtual machines that belong to two different ESX servers, also known as "cluster access boxes", when you meet specific configuration requirements.

The following configuration requirements must be met to use FC RDM LUNs on virtual machines in a Microsoft cluster:

- The virtual machines must be running the same Windows server version. The virtual machines must both be running either Windows Server 2008 or Windows Server 2012.
- ESX server versions must be the same for each VMware parent host.
- Each parent host must have at least two network adapters.
- There must be at least one VMFS datastore shared between the two ESX servers.
- VMware recommends that the shared datastore be created on an FC SAN; however, the shared datastore can also be created over iSCSI.
- The shared RDM LUN must be in physical compatibility mode.
- The shared RDM LUN must be created manually outside of SnapDrive. You cannot use virtual disks for shared storage.
- A SCSI controller must be configured on each virtual machine in the cluster in physical compatibility mode.

Windows Server 2008 requires you to configure the LSI Logic SAS SCSI controller on each virtual machine.

Shared LUNs cannot use the existing LSI Logic SAS controller if only one of its type exists and it is already attached to the C: drive.

SCSI controllers of type paravirtual are not supported on VMware Microsoft clusters.

Note: When you add a SCSI controller to a shared LUN on a virtual machine in physical compatibility mode, you must select the **Raw Device Mappings** option and not the **Create a new disk** option in the VMware Infrastructure Client.

- Microsoft virtual machine clusters cannot be part of a VMware cluster.
- You must use VCenter credentials and not ESX credentials when you install SnapDrive for Windows on virtual machines that will belong to a Microsoft cluster.
- A single initiator group must be created for both nodes in the cluster. You can create the initiator group automatically during disk creation or connection using the SnapDrive MMC. SnapDrive automatically selects an FC initiator from each of the ESX servers in the cluster. You can also create the initiator groups manually. If initiator groups do not already exist, you must create one manually on the storage system.
- You can create a RDM LUN on ESXi 5.0 using an FC initiator. When you create RDM LUN, an initiator group is created with ALUA.

Microsoft cluster support limitations when using FC RDM LUNs

SnapDrive supports Microsoft clusters using FC RDM LUNs on different virtual machines belonging to different ESX servers, but in some instances this feature is not supported.

- SnapDrive does not support clusters on ESX iSCSI and NFS datastores.
- SnapDrive does not support mixed initiators in a clustered environment. Initiators must be either FC or Microsoft iSCSI, but not both.

Note: ESX iSCSI initiators and HBAs are not supported on shared disks in a Microsoft cluster.

- SnapDrive does not support virtual machine migration with vMotion if the virtual machine is part of a Microsoft cluster.
- SnapDrive does not support MPIO on virtual machines in a Microsoft cluster.
- SnapDrive does not support ALUA in a Microsoft cluster using shared RDM LUNs.

Creating a shared FC RDM LUN

Before you can use FC RDM LUNs to share storage between nodes in a Microsoft cluster, you must first create the shared quorum disk and shared storage disk, and add them to both virtual machines in the cluster.

About this task

The shared disk is not created using SnapDrive for Windows.

Step

1. Create and then add the shared LUN to each virtual machine in the cluster using the procedure in the VMware *Setup for Failover Clustering and Microsoft Cluster Service* documentation. See the section that describes how to cluster virtual machines across physical hosts.

Managing LUNs in VMware environments

You can use SnapDrive to manage LUNs in VMware environments.

Connecting to an RDM LUN on a guest OS

You can connect your SnapDrive for Windows host to an RDM LUN on a guest OS using the Connect Disk wizard in the SnapDrive MMC snap-in.

Before you begin

Before connecting a LUN in a VMware guest OS, you must manually create initiator groups by using either OnCommand System Manager or at the storage system console.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage. Then, select **Disks**.
- 2. From the menu choices at the top of MMC, navigate to Action > Connect Disk.
- 3. In the Connect Disk wizard, click Next.
- 4. In the **Provide a Storage System Name, LUN Path and Name** panel, perform the following actions:
 - a) In the "Storage System Name" field, type the name of the storage system where the LUN will be connected, or choose a storage system from the drop-down list.
 - b) In the "LUN Path" field, type the path to the LUN. Alternatively, click Browse and navigate to the LUN you want to connect.
 - c) Click Next.
- 5. In the Select a LUN Type panel, select Dedicated, and then click Next.
- 6. In the Select LUN Properties panel, perform the following actions:
 - a) Either select a drive from the list of available drive letters, or enter a mount point for the LUN you are connecting. When you create a volume mount point, enter the drive path that the mounted drive will use: for example, G:\mount drive1\.

Note: The root volume of a new mount point must be owned by the node on which you are connecting the new disk.

- b) Click Next.
- 7. In the **Select Initiators** panel, choose at least one initiator for the LUN to which you are connecting from the list of available initiators, and then click **Next**.

Note: If MPIO is installed on the system, you can select multiple FC initiator ports. You cannot select both iSCSI and FC initiators when creating a LUN on a guest OS.

8. In the Select Initiator Group Management panel, specify whether you will use automatic or manual igroup management.

If you select automatic igroup management, SnapDrive uses existing igroups or, when necessary, creates new igroups for the initiators you have specified. If you select manual igroup management, you manually choose existing igroups or create new ones as needed.

If you specify	Then
Automatic igroup management	Select Automatic, and then click Next.
Manual igroup management	Select Manual , click Next , and then, in the Select Initiator Groups panel, perform ONE of the following actions:
	• Select from the list the igroups to which you want the LUN to belong, then click Next .
	Note: A LUN can be mapped to an initiator only once.
	• Click Manage igroups , and for each new igroup you want to create, type a name in the Igroup Name text box, select initiators, and click Create . Then, click Finish to return to the Select Initiator Groups panel, and click Next .
	Note: The Next button in the Select Initiator Groups panel remains unavailable until the collection of selected igroups contains all the initiators you previously selected for use.

9. If you selected an FC initiator on an ESX server in the Select Initiator Group Management panel, then the Select a Datastore panel is displayed.

If your virtual machine resides on a	ien	
VMFS datastore	Choose either Store with Virtual Machine , which is the default, or choose Specify datastore to select a different VMFS datastore for your virtual machine disk format (VMDK) file to reside.	;
	Click Next.	
NFS datastore	Choose Specify datastore to select a different VMFS datastore for your VMDK file to reside. The Store with Virtual Machine option is unavailab because a VMDK file cannot be stored on an NFS datastore.	ole
	Click Next.	

10. In the Completing the Connect Disk Wizard panel, perform the following actions:

- a) Verify all the settings.
- b) If you need to change any settings, click **Back** to go back to the previous wizard panels.
- c) Click Finish.

The newly connected LUN now appears under **SnapDrive > Disks** in the left MMC panel.

Guideline for managing RDM LUNs not created in SnapDrive

When you create an RDM LUN outside of SnapDrive, you must restart the SnapDrive for Windows service to enable SnapDrive to properly recognize the newly created disk. For the best results, use SnapDrive to create RDM LUNs.

Managing Snapshot backups in VMware environments

You can use SnapDrive to manage Snapshot backups in VMware environments.

Support requirements for performing Snapshot copy operations in VMDKs on NFS and VMFS datastores

You can use SnapDrive with Virtual Storage Console Backup and Recovery when you want to create or delete Snapshot copies in VMDKs on NFS and VMFS datastores.

The following conditions must exist before you can use SnapDrive to create and delete Snapshot copies in VMDKs on an NFS or VMFS datastore:

• Virtual Storage Console is installed and the appropriate IP address and port number are configured in SnapDrive.

SnapDrive provides a Virtual Storage Console configuration panel in the Installation wizard.

- Virtual Storage Console is installed on either the same system running the vCenter Server or another 32-bit or 64-bit Windows computer.
- Virtual Storage Console is available and reachable from the virtual machine on which SnapDrive is installed.
- The storage system IP address is configured in the Virtual Storage Console to enable identification of VMDKs created across NFS and VMFS datastores.

For more information about configuring Virtual Storage Console, see the *Virtual Storage Console for VMware vSphere Installation and Administration Guide*.

Snapshot copy support limitations on VMDKs

When you use SnapDrive to perform Snapshot copy operations on VMDKs, you might encounter known interoperability issues of which you should be aware in advance.

The following VMDK and Snapshot functionality interoperability issues exist in SnapDrive environments:

- You cannot create VMDKs on Virtual Storage Console mounted datastores.
- You cannot create a Snapshot copy on a VMDK that contains VMware snapshots. If you try to create a Snapshot copy using SnapDrive when a VMware snapshot exists, SnapDrive displays an error indicating that backup creation has failed, and Snapshot copy mount and restore operations cannot occur.
- When you migrate a VMDK disk from one VM to another, the Snapshots copies made earlier cannot be listed in the destination VM.

• If you are running VMDKs on NFS and VMFS datastores, clustered Data ONTAP requires Virtual Storage Console 4.1 and SnapDrive 6.4.2 for Windows or later.

Troubleshooting VMDKs

You should be aware of some common VMDK problems and workarounds.

Inconsistent VMDK disk enumeration

Problem

SnapDrive does not consistently enumerate VMDK disks.

Cause

This issue occurs because VMDK disks do not support the RPC transport protocol, which is the default transport protocol for SnapDrive.

Workaround

Change the transport protocol setting from RPC to HTTP.

Incorrect VMDK disk partition style

Problem

VMDK disk partition style incorrectly displays as Unknown in SnapDrive CLI and GUI.

Cause

This problem occurs with VMs that have Windows Server 2012 installed on ESXi 5.0 U2. There is an entry missing from the VMX file.

Workaround

1. Manually edit the VMX file by adding the following:

disk.EnableUUID="TRUE"

2. Save your edited VMX file.

Independent disks are excluded from backups

Problem

Independent disks are excluded during backup when VMware snapshots are used.

Cause

When a VMware snapshot of a VMDK disk already exists, you cannot use a SnapDrive Snapshot copy created on an independent disk (persistent or nonpersistent) to mount or restore a VMDK disk.

Workaround

Remove the VMware snapshot and then create a Snapshot copy of VMDK disks.

VMDK unable to discover all VMDK disks

Problem

VMDK not discovering all the VMDK disks.

Cause

This problem occurs with VMs that have Windows Server 2012 installed on ESXi 5.0 U2. There is an entry missing from the VMX file.

Workaround

1. Manually edit the VMX file by adding the following:

disk.EnableUUID="TRUE"

2. Save your edited VMX file.

Support requirements for space reclamation in VMDK files in NFS datastores

You can use the Space Reclaimer feature to reclaim space left by recently freed blocks in VMDK files located in NFS datastores when you use SnapDrive with Virtual Storage Console in an ESX environment.

You can run Space Reclaimer on VMDK files using SnapDrive MMC and sdcli.exe.

SnapDrive supports VMDK file space reclamation in NFS datastores in the following cases:

- When VMDK files have the simple extent type FLAT Space reclamation is not supported on sparse extents and VMFS datastores.
- When Virtual Storage Console is present If Virtual Storage Console is not available, SnapDrive is unable to discover VMDKs and, therefore, cannot reclaim space on VMDK files.
- When SnapDrive is configured with either ESX server or vCenter credentials

Additional Microsoft clusters on ESX documentation resources

You can find additional information about configuring your Microsoft cluster on an ESX server by reading these documents and knowledge base articles.

The following documentation is located on the VMware website.

- KB article 1009287, *ESX machines hosting passive MSCS nodes report reservation conflicts during storage operations*
- KB article 1004617, Microsoft Cluster Service (MSCS) Support on ESX
- Setup for Microsoft Cluster Service
- Setup for Failover Clustering and Microsoft Cluster Service
- VMware View 5.1 Release Notes

Performing SnapVault and SnapMirror operations

You can integrate SnapDrive with SnapVault and SnapMirror, to perform vaulting and mirroring operations for data protection and recovery.

Using SnapVault with SnapDrive

SnapVault is a Data ONTAP feature that enables you to back up Snapshot copies to a secondary storage system quickly, efficiently, and in a centralized manner.

Considerations when using SnapVault

You should be aware of certain considerations when you are using SnapVault with SnapDrive, including installation and configuration requirements.

System requirements for performing SnapVault operations

Before you perform SnapVault operations, you must have met the following requirements:

- You must have Data ONTAP 7.2.6 or later installed on your storage system.
- SnapVault must be licensed on the primary and secondary storage systems.
- You must have a license for the secondary SnapVault instance when you are using SnapVault with SnapDrive without SnapManager or N series Management Console.
- -mo relationships must be configured and initialized.

Backup set considerations when performing SnapVault operations

You should be aware of the following backup set considerations when you are performing SnapVault operations using SnapDrive:

- A backup set might contain multiple primary storage systems and volumes, but only one secondary volume and storage system.
- Each backup set can span only one volume on a secondary storage system; if multiple volumes are required, the backup fails.

SnapVault operation limitations in 7-Mode environments

You should be aware of the following additional limitations when you are operating SnapVault in 7-Mode environments:

- Only qtree SnapVault configurations are supported. SnapDrive does not support volume-based SnapVault.
- SnapVault cascaded configurations are not supported.
- There is no SnapVault restore feature.

Initiating SnapVault backup jobs from SnapDrive in 7-Mode SAN environments

You can initiate SnapVault backup jobs in 7-Mode SAN environments from MMC or by using the SnapDrive for Windows CLI.

About this task

These steps describe how to initiate a backup job using the Update SnapVault option in MMC.

Alternatively, you can also initiate a backup using the sdcli snapvault archive command.

For information about initiating SnapVault backup jobs in clustered Data ONTAP in SMB 3.0 environments, see the *SnapDrive for Windows PowerShell Cmdlet Reference Guide*.

Steps

- 1. Under SnapDrive in the left MMC pane, expand the instance of SnapDrive you want to manage, then select **Disks**.
- 2. Double-click the LUN for which you want to perform a SnapVault update.
- 3. Select Primary Snapshots to display the Snapshot copies on the primary system.
- 4. In the right MMC pane, right-click the Snapshot copy from which you want the SnapVault update to be initiated and select **SnapVault** from the menu.

A Snapshot copy with the same name as the Snapshot copy you selected on the primary system is created on the secondary storage system after the SnapVault update.

SnapVault operations supported in a clustered Data ONTAP

SnapDrive for Windows supports SnapVault operations in clustered Data ONTAP environments through SnapManager for SQL Server-initiated operations. SnapVault support is on the volume level. You should be aware of SnapVault support SnapDrive for Windows configuration requirements.

To ensure that SnapDrive supports SnapVault operations, you must configure both the primary and secondary storage system login details in TPS.

Supported SnapVault operations

You can initiate the following SnapVault operations using SnapManager for SQL Server:

- Connecting to a Snapshot copy from SnapVault secondary storage
- · Disconnecting from a Snapshot copy from SnapVault secondary storage
- Creating a Snapshot copy on a primary storage system and transferring it to a secondary storage system
- Updating SnapVault with a Snapshot created on your primary storage system
- Restoring from your secondary SnapVault storage system

Using SnapMirror with SnapDrive for Windows

You can use SnapMirror with SnapDrive for Windows to replicate data.

SnapMirror overview

SnapMirror creates either asynchronous or synchronous replicas of volumes that host LUNs.

With asynchronous SnapMirror, data is replicated from a source volume to a partner destination volume at regular intervals.

With synchronous SnapMirror, data from a source volume or qtree is replicated on a destination volume or qtree at, or near, the same time it is written to the first storage system.

When the LUN data on your source volume is offline or no longer valid, you can connect to and use the copy of the LUN on the SnapMirror destination volume. Unless otherwise indicated, the information discussed in this chapter applies to volumes that host SnapMirror LUNs, whether they are asynchronous or synchronous.

If a storage system volume or storage system holding one or more LUNs suffers a catastrophic failure, you can use a mirrored destination volume to recover the LUNs.

Data consistency is ensured in SnapDrive with up-to-date SnapMirror synchronous operation. If the synchronous operation is not updated, SnapDrive reverts to previous Snapshot copies. However, if the SnapMirror synchronous operation is up-to-date, the destination system contains Snapshot copies and SnapDrive connects to the current active file system, ensuring that the data is consistent.

Types of SnapMirror replication

SnapMirror replicas are initiated upon normal Snapshot copy creation or when using special "rolling" Snapshot copies.

Replication upon Snapshot copy creation

Each time a Snapshot copy of a LUN is created—manually or because of a Snapshot copy schedule —SnapDrive determines whether the LUN from which the copy was made resides on a SnapMirror source volume. If so, then after the Snapshot copy is made, SnapDrive might send a SnapMirror update request to all the destination volumes associated with the source volume for that LUN.

When you initiate a Snapshot copy of a LUN on a SnapMirror source through SnapDrive, a window with a check box labeled "Initiate SnapMirror Update" is displayed. The check box is selected by default.

Replication using rolling Snapshot copies

You can also create a special type of Snapshot copy called a "rolling" Snapshot copy, using the Update SnapMirror operation in SnapDrive. These Snapshot copies are used exclusively to facilitate

frequent SnapMirror volume replication. Like regular Snapshot copies, rolling Snapshot copies are replicated to the SnapMirror destination volume as soon as they are created.

SnapDrive creates a new rolling Snapshot copy every time you initiate a mirror update operation (using the **Update Mirror** option in the **Action** menu or the sdcli snap update_mirror command) for a specific LUN drive residing on a SnapMirror source volume.

To guarantee that at least one rolling Snapshot copy for each LUN is always available on the destination volume, SnapDrive maintains a maximum of two rolling Snapshot copies on the source volume.

How SnapDrive manages rolling Snapshot copies

When you initiate an Update Mirror operation, SnapDrive checks for any existing rolling Snapshot copies of the LUN containing the specified LUN drive. What happens next depends on whether SnapDrive detects none, one, or two Snapshot copies.

- If SnapDrive does not find any rolling Snapshot copies containing the LUN image, it creates a rolling Snapshot copy on the SnapMirror source volume. SnapDrive then initiates a SnapMirror update operation, which replicates the rolling Snapshot copy on the destination volume.
- If SnapDrive finds one rolling Snapshot copy, it creates a second rolling Snapshot copy and initiates a SnapMirror update.
- If SnapDrive detects two rolling Snapshot copies for the LUN, it deletes the older rolling Snapshot copy and creates a new one to replace it. Then SnapDrive initiates a SnapMirror update. When you connect to a LUN in a Snapshot copy that is located on a traditional volume, SnapDrive creates a LUN backed by a Snapshot copy on the active file system. When a new Snapshot copy is created as part of a synchronous SnapMirror update, that new Snapshot copy locks the Snapshot copy from which the LUN was connected. While the original Snapshot copy is locked, you are unable to delete it until the next SnapMirror update, when the first Snapshot copy is deleted automatically.

How rolling Snapshot copies are named

Rolling Snapshot copies can be identified by the unique names they are given.

The following format is used to name the rolling Snapshot copies:

@snapmir@{GUID}

GUID (Globally Unique Identifier) is a unique 128-bit number generated by SnapDrive to uniquely identify each rolling Snapshot copy.

Examples of rolling Snapshot copies

@snapmir@{58e499a5-d287-4052-8e23-8947e11b520e}

@snapmir@{8434ac53-ecbc-4e9b-b80b-74c5c501a379}

Requirements for using SnapMirror with SnapDrive

Before you can use SnapMirror with SnapDrive for Windows, your system must meet several requirements.

• SnapMirror must be licensed on the source and destination storage systems. For how to license and set up SnapMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode.*

Note: You must provide cluster credentials in the transport protocol settings before performing SnapMirror operations in a clustered Data ONTAP environment.

- Depending on the LUN protocols you are using, enable the iSCSI and FC licenses on the destination storage systems to enable LUN connect and LUN management operations.
- You must manually create and initialize a plex between the source and destination volumes, but you must not create a SnapMirror replication schedule.
 When setting up SnapMirror on your storage system, you can avoid schedule conflicts with SnapDrive by setting the replication schedule on the storage system to "- - ", which disables any scheduled transfers. When you set the replication schedule, ensure that the destination volume is in a restricted state. See the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode* for additional details.
- You must create your SnapMirror relationship using storage system names (either the fully qualified DNS name or the storage system name alone), and the name of the network interface to be used for SnapMirror transfers (for example, storage1-e0), not IP addresses.
- If you are using the optional MultiStore feature of the Data ONTAP software to create virtual storage systems (vFiler units), you must create your SnapMirror relationship on the vFiler unit, not on the physical storage system.
- The system must contain one or more SnapMirror source volumes hosting LUNs.
- The system must contain one or more SnapMirror destination volumes for each source volume.

Note: SnapDrive supports the use of SnapMirror at the volume level only; it does not support qtree-level SnapMirror operations.

- The destination volume must be at least as large as the source volume.
- The Windows domain account used by the SnapDrive service must be a member of the local BUILTIN\administrators group and must have management access to both the source and destination storage systems.
- The Windows domain account used to administer SnapDrive must have full access to the Windows domain to which both the source and destination storage systems belong.
- The source and destination storage systems must be configured to grant root access to the Windows domain account used by the SnapDrive service.
 That is, the wafl.map_nt_admin_priv_to_root option must be set to On. For information

about enabling storage system options, see your Data ONTAP documentation.

• If you want to use a Windows host to access the replicated LUNs on the destination volume, the destination storage system must have at least one LUN access protocol licensed (iSCSI or FC).

- A TCP/IP connection must exist between the source storage system and the destination storage system.
- The SnapDrive service can perform only one task at a time; therefore, if you are scheduling multiple tasks on a host, ensure that you do not schedule these tasks to start at exactly the same time.

If multiple tasks are scheduled at the same time, the first proceeds, and SnapDrive queues the others until the first task either succeeds or times out.

Note: SnapMirror cascaded configurations are not supported.

Initiating replication manually

SnapDrive initiates SnapMirror replication automatically when a Snapshot copy is created, but you can also initiate SnapMirror replication manually.

Before you begin

Because SnapDrive automatically initiates SnapMirror replication after a Snapshot copy for a LUN on a SnapMirror source volume has been created, to initiate replication after a Snapshot copy has been created, you must either manually create a Snapshot backup copy or set up a schedule for automatic Snapshot backup copy creation.

About this task

Manual replication is not monitored by SnapDrive, so you do not know if replication succeeded.

Steps

- 1. Perform the following actions to find the LUN that you want to replicate:
 - a) In the left MMC pane, select the instance of SnapDrive you want to manage.
 - b) Double-click Disks.
- 2. In the right MMC panel, select the LUN that you want to replicate.
- 3. Click Action (from the menu choices at the top of MMC window).
- 4. Select Update Mirror from the drop-down menu.

Note: The Update Mirror option is not available if no live mirror copies are configured.

Result

The **Update Mirror** operation is initiated and a rolling Snapshot copy of the LUN is created. After the Snapshot copy has been created on the mirrored source volume, SnapDrive automatically updates the mirrored destination volume.

Connecting to a LUN in a mirrored destination volume

You can connect to a LUN on a SnapMirror destination volume when you want to continue to serve data but a LUN on the source volume is inaccessible.

Before you begin

- SnapDrive supports the use of FlexClone volumes, which enable you to clone an independent volume from a parent FlexVol volume so that the mirror can remain unbroken.
- The LUN on an asynchronous SnapMirror destination must be restored from the most recent SnapDrive-created Snapshot copy containing a valid image of that LUN. The restoration is performed by SnapDrive as part of the LUN connect operation on an active file system or on a SnapMirror destination volume.

Note: The most recent Snapshot copy must be one created by SnapDrive to ensure data consistency. Data ONTAP creates a Snapshot copy that is more recent than the Snapshot copy created by SnapDrive; however, the Data ONTAP Snapshot copy cannot be used by SnapDrive because it is not consistent.

Steps

- 1. Connect to the mirrored LUN on the SnapMirror destination storage system.
- 2. If you want to break the SnapMirror relationship and connect to a SnapMirror destination volume that is online and, in the case of an asynchronous SnapMirror volume, perform a single file SnapRestore operation or rapid LUN restore, click **Yes** in the **Connect Disk** dialog box.

Note:

- You must perform this step only if the destination volume is not "broken". The mirror does not need to be broken if you connect to a LUN inside a Snapshot copy.
- SnapDrive identifies a volume as mirrored even after the SnapMirror relationship is broken.
- In clustered Data ONTAP, you cannot connect a qtree LUN in a Snapshot copy from a mirrored destination volume.

Restoring a volume on a SnapMirror destination

The volume restore feature in SnapDrive enables you to restore all the LUNs on a volume automatically from a single Snapshot copy when you establish a connection to the first LUN on a SnapMirror destination.

Before you begin

The following prerequisites must be met before SnapDrive can initiate a volume-based Snapshot copy restoration:

• All LUNs on the active file system must be consistent in the Snapshot copy you intend to use to restore.

- LUNs on the active file system must be of the same size and have the same name as the selected Snapshot copy.
- A SnapMirror relationship must exist.
- LUNs on the volume being restored must be disconnected from the host before they can be restored from the Snapshot copy.

About this task

Volume restore functions are currently available through the sdcli utility.

Step

1. Enter the following command from a Windows command prompt:

```
sdcli snap restore_volume [-f StorageSystemName] -volume
StorageSystemVolumeName -s SnapshotCopyName [-force] [-m MachineName]
```

-f StorageSystemName is the name of the storage system on which the volume resides.

-volumeStorageSystemVolumeName indicates name of the volume on which the restore operation will be performed.

-s SnapshotCopyName indicates the name of the Snapshot copy from which the volume will be restored.

-force is an optional switch that you use to ensure the volume restoration is performed even when non-LUN files or newer Snapshot copies are found on the volume.

-mMachineName is the host on which the operation is executed. You can use an IP address or a machine name to identify the host.

Result

The restoration is performed on the volume indicated.

Example

```
sdcli snap restore_volume -f clpubs-filer1 -volume vol3 -s my_snap
```

The preceding example restores a volume from the Snapshot copy named my_snap on a volume called vol3 that resides on a storage system called clpubs-filer1.

Recovering a cluster from shared LUNs on a SnapMirror destination

Connect to shared LUNs on a SnapMirror destination in order to recover your MSCS cluster.

Before you begin

The following prerequisites must be met before you can successfully use the procedure described in this section to connect to shared LUNs on a SnapMirror destination and thus recover your MSCS cluster:

- A SnapMirror replica of the source volume must exist on the destination volume prior to the failure of the physical disk resource.
- You must know the original drive letters and paths to the shared LUNs on the SnapMirror source volume.
- You must know the MSCS cluster name.

Steps

- 1. Configuring the cluster service to start manually on page 98
- 2. Creating a temporary quorum disk on page 98
- 3. Starting the cluster service with the -fixquorum option on page 99
- 4. Connecting to the new quorum disk on page 99
- 5. Connecting to a shared LUN on the SnapMirror destination volume on page 99

Configuring the cluster service to start manually

In order to recover a cluster from shared LUNs on a SnapMirror destination, you must first configure the cluster service to start manually.

Steps

- 1. Configure the cluster service to start manually on all nodes of the cluster by performing the following actions on each node of the cluster:
 - a) In the left MMC pane, expand the Services and Applications option, if it is not expanded already.
 - b) Click Services.
 - c) Double-click Cluster Service.
 - d) Select Manual from the Startup Type list.
- 2. Reboot all nodes of the cluster.

Note: The reboot is required so the existing LUNs fail to mount and, therefore, the drive letters that were in use will be released.

Creating a temporary quorum disk

After you configure the cluster service to start manually, create a temporary quorum disk.

Steps

1. Create a shared disk on the SnapMirror destination storage system to be used as a temporary quorum disk.

After you have successfully completed the Create Disk wizard, you see the following message. This message is expected and does not indicate a problem.

You have successfully configured a disk on this system with the intention of it being a shared resource in MSCS. As MSCS does not appear to be installed on this system, please install MSCS.

- 2. Click **OK** to ignore the message.
- 3. Disconnect the shared disk you just created.

Starting the cluster service with the -fixquorum option

Complete this procedure after you create a temporary quorum disk.

Steps

- 1. In the left MMC panel, click Services.
- 2. In the Start Parameters field, enter -fixquorum.
- 3. In the Service Status field, click Start, then click OK.

Connecting to the new quorum disk

Complete this procedure after you start the cluster service with the -fixquorum option.

Steps

- 1. Reconnect the shared disk you created.
- 2. Using the Cluster Administrator, make the newly connected shared disk the quorum disk.
- 3. Stop the cluster service, then restart the cluster service on all nodes in the cluster.
- **4.** Remove dependencies on all failed physical disk resources, then remove the physical disk resources.

Connecting to a shared LUN on the SnapMirror destination volume

After you have connected to the new quorum disk, you can connect to a shared LUN on the SnapMirror destination volume.

Steps

- 1. On the cluster node you used earlier, follow the steps to connect to a LUN, keeping in mind the following information to connect to a LUN:
 - a) When prompted for the LUN path in the **Provide Storage System, LUN Path, and Name** panel, specify or browse to the LUN file in the active file system (not the one in the Snapshot copy) on the SnapMirror destination volume.
 - b) After you specify the LUN path and click **Next**, you see a message that a single file SnapRestore or rapid LUN restore will be performed. Click **Yes** to continue.
 - c) When prompted for disk type in the Select a LUN Type panel, select Shared.

- d) When prompted for a drive letter in the **Select LUN Properties** panel, select the same drive letter that was being used for the LUN on the SnapMirror source volume.
- 2. After you have successfully completed the Connect Disk wizard, you see one of the following two error messages. These error messages are expected and do not indicate a problem.

Error message 1:

Unable to connect disk. Failure in Mounting volume on the disk. Error: Could not find the volume mounted for the LUN as there does not seem to be any new volumes mounted by the Mount Manager

This error might also appear in the following form:

Unable to connect disk. Failure in connecting to the LUN. Error: Timeout has occurred while waiting for disk arrival notification from the operating system.

Error message 2:

```
Unable to retrieve a list of LUN snapshots. Error: The device is not ready.
```

Note: Error message 2 is displayed instead of error message 1 when McAfee NetShield is installed on your Windows server.

Click **OK** to ignore the error message.

- 3. Repeat Step 1 and Step 2 for each shared LUN on the cluster.
- **4.** Configure the cluster service to start automatically on the system to which you connected shared LUNs by performing the following actions:
 - a) In the left **MMC** pane, expand the Services and Applications option, if it is not expanded already.
 - b) Click Services.
 - c) Double-click Cluster Service.
 - d) Select Manual from the Startup Type list.
- 5. Restore any resource dependencies you removed earlier.

After you finish

Use the Cluster Administrator to verify that the cluster is functioning correctly as follows:

- 1. Ensure that all resources are online.
- 2. Perform a "move group" operation from one node to the other and then back to the original node.
- 3. Move the quorum disk from the temporary disk you created in Step 3 back to the original disk.
- 4. Delete the temporary disk.

SnapDrive integration with OnCommand Unified Manager Core Package data protection capabilities

SnapDrive integrates with OnCommand Unified Manager Core Package data protection capabilities to provide management of SnapMirror and SnapVault deployments using datasets.

How SnapDrive integrates with OnCommand Unified Manager Core Package data protection capabilities

SnapDrive makes it easy for you to manage very large SnapMirror and SnapVault deployments by supporting OnCommand Unified Manager Core Package data protection capabilities through the SnapManager products.

The OnCommand Unified Manager Core Package data protection capabilities ease the management of very large deployments by grouping data and storage systems into *datasets* and *resource pools*, enabling automation of many routine data protection tasks. You can configure SnapDrive with a set of OnCommand Unified Manager Core Package credentials so that it can authenticate to a DataFabric Manager server. This allows SnapManager to use SnapDrive as a conduit to support OnCommand Unified Manager Core Package retention policies and schedules.

To take advantage of OnCommand Unified Manager Core Package data protection capabilities through SnapDrive, see the interoperability matrix at *www.ibm.com/systems/storage/network/ interophome.html* for supported Data ONTAP and OnCommand Unified Manager Core Package versions.

For more information, see your SnapManager documentation.

Dataset concepts

Associating data protection, disaster recovery, a provisioning policy, or a storage service with a dataset enables storage administrators to automate tasks, such as assigning consistent policies to primary data, propagating policy changes, and provisioning new volumes, qtrees, or LUNs on primary and secondary dataset nodes.

Configuring a dataset combines the following objects:

Dataset of
physical storageFor protection purposes, a collection of physical resources on a primary node,
such as volumes, FlexVol volumes, and qtrees, and the physical resources for
copies of backed-up data on secondary and tertiary nodes.

For provisioning purposes, a collection of physical resources, such as volumes, FlexVol volumes, qtrees, and LUNs, that are assigned to a dataset node. If the protection policy establishes a primary and one or more nonprimary nodes, each node of the dataset is a collection of physical resources that might or might not be provisioned from the same resource pool.

Dataset of virtual objects	A collection of supported VMware virtual objects that reside on storage systems. These virtual objects can also be backed up locally and backed up or mirrored on secondary and tertiary nodes.
Resource pool	A collection of physical resources from which storage is provisioned. Resource pools can be used to group storage systems and aggregates by attributes, such as performance, cost, physical location, or availability. Resource pools can be assigned directly to the primary, secondary, or tertiary nodes of datasets of physical storage objects.
	They can be assigned indirectly both to datasets of virtual objects and to datasets of physical storage objects through a storage service.
Data protection policy	A set of rules that define how to protect primary data on primary, secondary or tertiary storage, as well as when to create copies of data and how many copies to keep.
	Protection policies can be assigned directly to datasets of physical storage objects. They can be assigned indirectly to both datasets of virtual objects and to datasets of physical storage objects through a storage service.
Provisioning policy	A set of rules that define how to provision storage for the primary or secondary dataset nodes, and provides rules for monitoring and managing storage space and for allocating storage space from available resource pools.
	Provisioning policies can be assigned directly to the primary, secondary, or tertiary nodes of datasets of physical storage objects. They can be assigned indirectly to both datasets of virtual objects and datasets of physical storage objects through a storage service.
Storage service	A single dataset configuration package that consists of a protection policy, provisioning policies, resource pools, and an optional vFiler template (for vFiler unit creation). You can assign a single uniform storage service to datasets with common configuration requirements as an alternative to separately assigning the same protection policy, provisioning policies, and resource pools, and to setting up similar vFiler unit attachments to each of them.
	The only way to configure a dataset of virtual objects with secondary or tertiary backup and mirror protection and provisioning is by assignment of a storage service. You cannot configure secondary storage vFiler attachments for datasets of virtual objects.
Protection or provisioning related objects	Snapshot copies, primary volumes, secondary volumes, or secondary qtrees that are generated as a result of protection jobs or provisioning jobs.
Naming settings	Character strings and naming formats that are applied when naming related objects that are generated as a result of protection jobs or provisioning jobs.

Configuring and managing access control

You can configure access control by enabling RBAC with SnapDrive, and you can use the AccessControl.xml file to manage storage system access control.

Support for storage system access control

SnapDrive provides support for storage system access control to separate server administrator and storage administrator functions, and to limit SnapDrive actions and operations that depend on the user. This feature is not supported with clustered Data ONTAP.

SnapDrive enables you to control storage system access by reading a file called AccessControl.xml that is created by the storage system administrator. The file is created in the /etc directory of the storage system root volume and lists the operations and storage resources that are allowed access by users who are assigned specific roles. The access control file associates access rights with specific storage resources. A tool called storacl.exe is used to edit the access control file.

Using storage system access control

You can use the AccessControl.xml file on the storage system to determine what roles are assigned to a user and what operations are allowed for specific roles.

Before you begin

- The storage access control (storacl.exe) tool is installed in the SnapDrive directory on any Windows host.
- HTTPS is enabled using the options ssl.enable command and secureadmin setup ssl command on the storage system.

You can support HTTP use with vFiler units when using the MultiStore feature of Data ONTAP software.

• You are logged into the storage system as root.

About this task

Always use storacl.exe to modify the AccessControl.xml file. Using a standard XML editor might corrupt the file and cause SnapDrive operations to fail.

Steps

- 1. Run storacl.exe from your Windows host.
- 2. Edit the AccessControl.xml file to add or change users and to grant or deny access rights to the storage system resources.

You can also use the default entries.

Storage system access control reference

Your storage system administrator can use the AccessControl.xml file to allow or deny access to resources on a storage system and the ThinProvision.xml file to enable the creation of thinly provisioned LUNs when using SnapDrive.

Note: Storage system access control is not supported in clustered Data ONTAP.

Storage system access control operations descriptions

You can use several operations in the AccessControl.xml file to allow or restrict access to users, depending on their assigned roles.

You can use the following operations with the default roles provided in the AccessControl.xml file, or you can create new roles using one or more of the available operations.

Operation	Descriptions
SD.Config	Allows users to perform all configuration operations.
SD.Config.Read	Allows users to only read configuration information on the storage system.
SD.Config.Write	Allows users to create and modify configuration information on the storage system, such as creating and modifying igroups, creating mappings, and setting options.
SD.Config.Delete	Allows users to delete any configuration information.
SD.Storage	Allows users to perform all storage operations.
SD.Storage.Read	Allows users to list storage system objects, such as enumerating and reading aggregates, volumes, qtrees, LUNs, and files.
SD.Storage.Write	Allows users to create and modify storage system objects.
SD.Storage.Delete	Allows users to delete storage objects.
SD.Snapshot	Allows users to perform all Snapshot copy operations.
SD.Snapshot.Read	Allows users to view Snapshot copies, including archived copies and objects inside copies.

Operation	Descriptions
SD.Snapshot.Delete	Allows users to delete Snapshot copies.
SD.Snapshot.Write	Allows users to create, rename, and modify Snapshot copies.
SD.Snapshot.Restore	Allows users to restore from a Snapshot copy.
SD.Snapshot.Clone	Allows users to clone related operations, such as FlexClone volume operations and LUN clone operations.
SD.Access.None	Denies all access to a storage system.

Storage system access control roles

You can assign roles to a user in the AccessControl.xml file to allow or restrict only certain operations to SnapDrive and storage system resources.

You can use these default roles to restrict access or you can create new roles using the available operation types.

Role	Operations	Descriptions
SD.Admin	All operations	Allows all SnapDrive operations.

Role	Operations	Descriptions
SD.Provision	SD.Storage.Write SD.Storage.Read SD.Config.Read SD.Config.Write	Allows all LUN provisioning and Snapshot copy operations, including create, connect, and map, if the operations are set on the storage system. If a LUN is disconnected by a user to whom you have assigned the default SD.Provision role, but the volume on which the LUN resides does not have Storage.Read permission, that user cannot reconnect the LUN from SnapDrive MMC using manual igroup management. This is because without the Storage.Read permission on the storage system, no igroups are listed. In this case, the user can reconnect the LUN using automatic igroup management or sdcli.exe.
SD.Discovery	SD.Config.Read SD.Storage.Read SD.Snapshot.Read	Allows all operations for discovering volumes, qtrees, igroups, and Snapshot copies.
SDBackup	SD.Snapshot.Read SD.Snapshot.Write	Allows create, replicate, and archive backup operations.
SDRestore	SD.Snapshot.Restore	Allows restore operations from a Snapshot copy or archive.
SDNoAccess	SD.Access.None	Denies all access. When you use this role, the storage access control tool does not allow other roles to exist.

Storage system access control commands

The storage system access control tool provides several commands that enable you to control what actions a user can perform on a storage system.

The following commands are available in the storage system access control tool:

create	Creates the AccessControl.xml file that contains default operations and roles. The storacl.exe tool places the AccessControl.xml file in the storage system /etc directory.	
delete	Deletes the AccessControl.xml file after user confirmation. Deleting the AccessControl.xml files disables access control on the storage system.	
operation	Lists the SnapDrive operations.	
roles	Lists, adds, removes, and modifies both default and user-created roles on the storage system.	
user	Enables you to perform the following actions:	
	List users for whom access rights have been setAdd, remove, and modify access rights for users	
dfmrbac	Sets the DFM-RBAC value that determines which RBAC method to use: either DFM or file-based.	
storage	Lists the storage system resources for which access rights have been configured.	
spacereserve	Enables and disables thin provisioning of LUNs using the value true or false. The true value enables fully provisioned LUNs. The false value enables thinly provisioned LUNs.	
host	Enables you to list storage system volumes and to add or remove host access.	
hvol	Lists the volumes a host can access.	
help	Provides Help on any storacl.exe command or operation.	
exit	Exits storacl.exe.	

Storage system access control command examples

You might find it useful to see examples showing some of the commands you can run when you use storacl.exe to manage storage access and to determine space reservation status on volumes.

Launch storacl.exe This example launches the storacl tool:

storacl

Launch storacl.exe and connect to storage system	This example launches the storacl tool and connects to the storage system called SYSTEM1:
	storacl -stor SYSTEM1
Disable access control on a storage system	This example disables access control to the root user on the storage system called System1:
	delete -stor System1 -user root
Create a new role	This example creates a role called TESTROLE with the ability to list storage system objects:
	role add -rn TESTROLE -OPN SD.Storage.Read
Add operations to an	This example adds storage write access to the role called TESTROLE:
existing role	role add -rn TESTROLE -OPN SD.Storage.Write
Remove operations from an existing role	This example removes storage read access from the role called TESTROLE:
	role remove -rn TESTROLE -OPN SD.Storage.Read
Add user access rights	This example adds the access rights defined by the SDProvision and SDBackup role to the domain user usr1 on the volume volTest on storage system System1:
	user add -rsn System1:/vol/volTest -rtype vol -un mydomain\usr1 -RN SDProvision SDBackup
Remove roles assigned to a user for a resource	This example removes only the SDBackup role for the domain user usr1 on the volume volTest on the storage system System1:
	user remove -rsn System1:/vol/volvpn -rtype vol -un mydomain\usr1 -RN SDBackup
List resources accessible to a user	This example lists the resources that are accessible by the domain user usr1:
	user list -un mydomain\user1
List resources	This example lists the resources that are accessible to all users:
accessible to all users	user list
Remove all access	This example removes all access rights to the domain user usr1:
rights to a user	user remove -un mydomain\usr1
List storage system	This example lists all volumes on the storage system:
resources	storage list -rtype vol
This example lists all aggregates in the storage system:

storage list -rtype aggr

List volumes accessible to a host	This example lists storage system volumes accessible by HOST1: hvol list -h HOST1
Remove volume access from a host	This example makes the volume volTest inaccessible to the host HOST1: hvol remove -h HOST1 -vol volTest
Remove a host entry from a storage system	This example removes the entry for HOST1 from the storage system: host remove -h HOST1
Check the current LUN space reservation policy	This example checks the current LUN space reservation policy on the volume vol1 on System1 using the root user to log in to the storage system:
	spacereserve get -vol voll -stor System1 -user root
	When no protocol type is specified, storacl uses HTTPS as the default. The default HTTPS port is 443. The default HTTP port is 80.
Enable space reservation	This example enables space reservation on the volume vol1 on System1 by logging on to the storage system as the user root:
Enable space reservation	This example enables space reservation on the volume vol1 on System1 by logging on to the storage system as the user root: spacereserve set -vol vol1 -val true -stor System1 - user root
Enable space reservation Disable space reservation	This example enables space reservation on the volume vol1 on System1 by logging on to the storage system as the user root: spacereserve set -vol vol1 -val true -stor System1 - user root This example disables space reservation on the volume vol1 on the storage system System1:

Enabling RBAC for use with SnapDrive

Before you can use OnCommand Unified Manager Core Package role-based access control (RBAC) with SnapDrive, you must make this feature available by first enabling the service on your storage system, then configuring your SnapDrive for Windows system with the appropriate OnCommand

Unified Manager Core Package server credentials, and finally adding and assigning roles to SnapDrive users.

Using RBAC with the OnCommand Unified Manager Core Package server

SnapDrive provides support for role-based access control (RBAC) with OnCommand Unified Manager Core Package server to separate server administrator and storage administrator functions and to limit SnapDrive actions and operations, depending on the role or job function of the user.

Role-based access control enables an application or Windows server administrator to provision and manage storage for their resources without needing the storage system root password.

Note: RBAC is not supported in clustered Data ONTAP environments.

When configured for role-based access control, SnapDrive serves as the policy decision point. For every SnapDrive operation, an RBAC access check to DataFabric Manager server is performed. The RBAC access check verifies whether a SnapDrive user can perform a SnapDrive operation on a storage system resource. The SnapDrive MMC displays only the object types to which a user has been granted access.

If DataFabric Manager server encounters a short service outage, SnapDrive is not affected. If a prolonged outage occurs, however, SnapDrive cannot operate.

Note: If DataFabric Manager server is down, the SnapDrive administrator can request the storage system administrator to disable role-based access control.

The following requirements must be met for SnapDrive to use role-based access control:

- OnCommand Unified Manager Core Package server must be present and configured in the IP network in which the storage system and the SnapDrive hosts exist.
- OnCommand Unified Manager Core Package server must be version 3.7 or later. For the most up-to-date supported configuration information, see N series interoperability matrix website (accessed and navigated as described in *Websites* on page 12).
- The DataFabric Manager server must have predefined roles, operations, and capabilities for SnapDrive.
- The storage system administrator must enable the DataFabric Manager server RBAC flag in the AccessControl.xml file located in the root volume of the storage system.
- SnapDrive must be configured by the host administrator with the DataFabric Manager server name and credentials to use role-based access control; otherwise, SnapDrive continues to use the original access control method, if one was configured.

The DataFabric Manager server name and credentials can be entered during SnapDrive installation or by using sdcli.exe.

• To enable backup operations to proceed, a storage administrator must create and apply a role with SD.Config.Read and SD.Config.Write capabilities at the storage system level. For example, at the STORACL prompt, use the following command to create a new role:

role add -rn SDConfigRole -OPN SD.Config.Read SD.Config.Write To apply the new role to the domain user on the specified storage system use the following command: user add -rsn StorageSystemName -rtype stor -un domain\user -RN SDConfigRole

- For disk create and connect operations to succeed using either the automatic igroup creation option or the manual igroup creation option with SnapDrive, the following permissions must be set:
 - If you choose the automatic igroup creation option, SD. Config. Write
 - If you choose the manual igroup creation option, SD.Config.Read

See the OnCommand Unified Manager Operations Manager Administration Guide for your version of DataFabric Manager server for more information about roles and capabilities and how they are used. See the Storage Access Control Tool User's Guide for more information about the storacl.exe tool.

Enabling RBAC on the storage system

The first step toward enabling RBAC for use with SnapDrive is having the storage administrator enable the feature on your storage system.

Steps

- 1. Download storacl.exe from the N series support website (accessed and navigated as described in *Websites* on page 12) to your Windows host.
- 2. Run storacl.exe from the location to which you downloaded it on your Windows host.
- **3.** At the STORACL> prompt, enter the following command:

```
create -stor StorageSystem -user UserName -pwd password
```

The AccessControl.xml file is created in the /etc directory of your storage system root volume.

4. At the STORACL prompt, enter the following command:

dfmrbac set -val true

See the *Operations Manager Administration Guide* for your version of DataFabric Manager server for more information about RBAC.

Note: SnapDrive does not support file-based RBAC in clustered Data ONTAP.

After you finish

You next configure SnapDrive to use the OnCommand Unified Manager Core Package RBAC feature.

Configuring SnapDrive for Windows to use RBAC

After your storage administrator enables RBAC on the storage system, you must configure your SnapDrive for Windows host with the proper DataFabric Manager server credentials.

Before you begin

- The SnapDrive service account has administrative privileges on both the storage system and the host.
- The user you specify from SnapDrive is configured with DFM. Core. AccessCheck capability in the global scope in DataFabric Manager server.

Step

1. If you did not configure your SnapDrive for Windows host with your DataFabric Manager server information when you installed SnapDrive, use the sdcli dfm_config command to enter the hostname, username, and password of the user accessing DataFabric Manager.

After you finish

You next create SnapDrive user roles using DataFabric Manager.

Creating SnapDrive user roles on DataFabric Manager server

After you configure SnapDrive for Windows to use RBAC, you must create roles on the DataFabric Manager server to grant access to SnapDrive for Windows users according to the SnapDrive operations you want them to perform.

Before you begin

The domain or workgroup users to whom you want to assign roles must already be created.

About this task

Perform these steps on DataFabric Manager using the Operations Manager interface. See the *Operations Manager Administration Guide* for your version of DataFabric Manager server for more information about creating roles using the Operations Manager interface.

Steps

- 1. Select Roles from the Setup menu.
- 2. Click Add Capabilities..., and from the Capabilities window, select a resource from the resource tree.
- 3. Select the capabilities you want for the new role you are creating in one of two ways:
 - Select the operations that you want to allow for the resource.

- To inherit roles, select that role from the Inherit Capabilities list on the left, and click >> to move it to the list on the right.
- 4. Click Add Role.

After you finish

You will next assign roles to SnapDrive users.

Assigning roles to SnapDrive users on DataFabric Manager server

After you create roles on DataFabric Manager server, you must assign the roles to SnapDrive users.

About this task

Perform these steps on DataFabric Manager using the Operations Manager interface. See the *Operations Manager Administration Guide* for your version of DataFabric Manager server for more information about assigning roles using the Operations Manager interface.

Steps

- 1. Select Administrative Users from the Setup menu.
- 2. Type the name of the administrator or Windows group to which you want to assign a role.
- 3. Select the role from the list on the left and click >> to move it to the list on the right.
- 4. Click Add.

SnapDrive for Windows to DataFabric Manager role mappings

To enable RBAC to work between SnapDrive for Windows and DataFabric Manager server, the DataFabric Manager server must be configured with predefined SnapDrive roles, capabilities and operations.

SnapDrive operation	Role and capability	Description
disk create	SD.Storage.Write	Enables LUN creation when set at the volume or qtree resource
disk expand	SD.Storage.Write	Enables LUN expansion when set at the disk resource
disk delete	SD.Storage.Delete	Enables LUN deletion when set at the disk resource

SnapDrive for Windows operations map to the DataFabric Manager roles as follows:

SnapDrive operation	Role and capability	Description
disk connect disk disconnect disk disconnect -f	SD.Storage.Write	Enables LUN connect, disconnect, and forced disconnect when set at the disk resource
disk list	SD.Storage.Read	Enables the ability to list all disks on a volume when set at the storage system volume resource
snap list	SD.Snapshot.Read	Enables the ability to list all Snapshot copies when set on a storage system volume resource
snap create	SD.Snapshot.Write	Enables Snapshot copy creation when set on a storage system volume resource
snap restore	SD.Snapshot.Restore	Enables Snapshot copy restoration when set on a storage system volume resource
snap delete	SD.Snapshot.Delete	Enables Snapshot copy deletion when set on volume when set at the storage system volume resource
snap rename	SD.Snapshot.Write	Enables Snapshot copy renaming on a volume when set at the storage system volume resource
snapvault archive	SD.Snapshot.Write	Enables SnapVault archiving when set on a secondary SnapVault volume resource
snapvault snapshot_delete	SD.Snapshot.Delete	Enables deletion of SnapVault Snapshot copies when set on a secondary SnapVault volume resource
snapvault snap_list	SD.Snapshot.Read	Enables the ability to list SnapVault Snapshot copies when set on a secondary SnapVault volume resource

SnapDrive operation	Role and capability	Description
Snapvault verify_configuration	SD.Snapshot.Read	Enables SnapVault configuration verification when set on primary and secondary SnapVault volumes resource
snapvault rename	SD.Snapshot.Write	Enables renaming of SnapVault Snapshot copies when set on the secondary SnapVault volume resource
igroup create	SD.Config.Write	Enables initiator group creation when enables on the storage system resource
igroup rename	SD.Config.Write	Enables renaming of initiator groups when set on the storage system resource
igroup delete	SD.Config.Delete	Enables deletion of initiator groups when set on the storage system resource
igroup list	SD.Config.Read	Enables the ability to list initiator groups when set on the storage system resource
snap mount	Set on the volume for LUN clones: SD.Snapshot.Clone Set on the qtree for LUN clones: SD.Snapshot.Clone Set on the storage system for traditional volume clones: SD.Snapshot.Clone Set on the parent volume for FlexVol volume clones: SD.Snapshot.Clone Set on the parent for FlexVol volume clones that are split: SD.Snapshot.Clone SD.Snapshot.Clone SD.Snapshot.Clone SD.Snapshot.Clone	Enables Snapshot copy mounting when set on the specified resources

SnapDrive operation	Role and capability	Description
snapshot unmount	For LUN clones in a volume or qtree: SD. Snapshot.Clone on the volume	Enables unmounting of LUN clones or volume clones on the indicated resources
	For volume clones: SD.Snapshot.Clone on the parent volume	
volume create	SD.Storage.Write	Enables volume creation when set on the aggregate resource
volume rename	SD.Storage.Write	Enables renaming of a volume when set on the volume resource
volume delete	SD.Storage.Delete	Enables deletion of a volume when set on the volume resource
volume list	SD.Storage.Read	Enables the ability to list volumes on a storage system aggregate when set on an aggregate resource
aggregate list	SD.Storage.Read	Enables the ability to list aggregates when set on the storage system resource

SnapDrive command-line reference

The SnapDrive command-line utility, sdcli.exe, enables you to perform many of the tasks available in the SnapDrive MMC, as well as some that are only available using the command-line.

About sdcli commands

Use the SnapDrive for Windows sdcli command-line utility to execute SnapDrive commands individually or through automation scripts.

Executing sdcli commands

The sdcli commands consist of three input parameters, which must be specified in the correct order, followed by one or more command-line switches. You can specify the command-line switches in any order.

Before you begin

Your system must meet the following requirements when you use the sdcli command-line utility on a Windows 2008 server:

- You must be logged in as Administrator, or as a user with administrative rights.
- If you logged in as a user other than the one used to install SnapDrive, you must have updated the SnapDrive service credentials with the new user information and restarted the SnapDrive service.

About this task

Command-line switches are case-sensitive. For instance, the -d switch refers to a single drive letter, while the -D switch refers to one or more drive letters, separated by spaces.

Steps

- 1. Using a host that has SnapDrive installed, select Start > Run.
- 2. In the dialog box entry field, type the following:

cmd

- 3. Click OK.
- **4.** After the Windows command prompt window opens, navigate to the directory on your host where SnapDrive is installed.

Example

cd \Program Files\IBM\SnapDrive\

5. Enter the individual command you want to run.

Make sure to include all input parameters in the proper order and to specify both required and desired command-line switches in any order:

Example

sdcli disk disconnect -d R

Alternatively, enter the name and path of the automation script you want to run:

Example

```
C:\SnapDrive Scripts\disconnect_R_from_host4.bat
```

Common command switches

Many of the sdcli commands share command-line switches. Common command-line switches are listed in the following table.

Switch	Description
-d	The drive letter, mount point, volume name, or CSV reparse point of the LUN. If sdcli cannot find the drive letter specified through the -d switch, it displays a list of all LUNs connected to the host.
-D	A list of drive letters or mount points separated by spaces.
-dtype	The drive type (shared or dedicated).
-e	The name of an existing MSCS resource group, which is required only if the LUN is shared among MSCS nodes.
-i	 The initiator name. For FC, the initiator name is the WWPN (World Wide Port Name) for the initiator, which takes the form hh:hh:hh:hh:hh:hh:hh. For iSCSI, the initiator name takes the form iqn.iSCSI qualified name. For more information on iSCSI node names, see the <i>Data ONTAP SAN Administration Guide for 7-Mode</i>.

Switch	Description
- I	The list of hosts and initiators.
	Separate the character strings that specify hosts and initiators with spaces.
	To specify the host, you can use either an IP address (<i>nnn.nnn.nnn.nnn</i>) or a machine name recognized by the domain controller.
	To specify the initiator, type the appropriate WWPN, which you can determine through the lputilnt.exe utility supplied with your FC HBA Attach Kit. After you launch lputilnt.exe, navigate to Main Menu > Adapter > Configuration Data and select "16 - World-Wide Name" in the Region field. The available WWPNs appear in the list box directly beneath the Region field.
	When MPIO is running, you can specify up to four node-initiator pairs. The first NodeMachineName in the cluster applies to two of the available initiator WWPNs; the other NodeMachineName applies to the remaining pair of initiator WWPNs.
-IG	The list of node machine names and existing igroup names, in pairs. One pair is required for dedicated disks. Two pairs are required for shared disks (at least one pair for each cluster node).
- m	The host on which the operation is executed. You can use an IP address or a machine name to identify the host.
	Note: When you use an IP address to identify a host, you should use only those IP addresses displayed in the output from the ipconfig /all command rather than those displayed in the output from the sdcli sysconfig list command.
	Note: Do not specify the -m switch when running an sdcli command on the local host.
-n	The name and description of an MSCS cluster resource group to be created as part of the associated command.
	This switch is required only if you need to create an MSCS cluster resource group to facilitate the sharing of a LUN among MSCS cluster nodes.
-np	The IP address and port of the network portal on the iSCSI connection target.
-p	The storage system path to the location of the LUN on the storage system. This string takes the following form: storagesystemname:/vol/volname/ [qtree]/lun for a storage system path.
- Z	Specifies the size (in megabytes) of a new LUN—or the number of megabytes by which an existing LUN is to be expanded. The minimum size for MBR partition- style LUNs is 32 MB and the minimum for GPT style LUNs is 64 MB. The maximum sizes vary according to the remaining available space in your volume.

Note: Switches that apply to just one command appear with those commands in the sections that follow.

Configuration commands

The sdcli utility provides command-line support for viewing SnapDrive configuration information.

The sysconfig list command

The sysconfig list command displays the SnapDrive configuration information for your host.

Syntax for this command is:

sdcli sysconfig list

Dataset management commands

You can use dataset commands to perform operations on and get information about datasets.

dataset add_members

You can use dataset add_members to add objects to your specified dataset.

Syntax

sdcli dataset add_members-dn DatasetName-AP AccessPointList

Description

You can use dataset add_members to add objects to your specified dataset.

Parameters

-dn

Indicates the dataset to which you want to add members.

-AP

Specifies a list of type-value pairs. For example, in the type-value pair

"G:\mp I:\mp"

mp indicates the mount point and

G:\

is the specifier.

dataset backup_add

You can use dataset backup_add to add a backup version to your dataset.

Syntax

```
sdcli dataset backup_add-dn DatasetName-bv BackupVersion-S
PrimarySnapshotNameList-AP AccessPointList-nn NodenameorId
```

Description

You can use dataset backup_add to add a backup version to your dataset.

Parameters

-dn Indicates the name of the dataset to which you want to add a backup version. -bv Indicates the backup version you want to add to your dataset. -s Indicates the Snapshot copy list for the dataset. When you are using the N series Management Console data protection capabilities in OnCommand Unified Manager Core Package, this list is mandatory.

-AP

Provides a list of type-value pairs: for example, G:\ mp I:\ mp.

-nn

Specifies the storage system node name or ID. If you do not specify a node or ID, the default is the primary storage system.

dataset backup_change_retention_type

You can use dataset backup_change_retention_type to change your dataset retention policy associated with a dataset backup.

Syntax

```
sdcli dataset backup_change_retention_type-dn DatasetName-bv
BackupVersion-i BackupId-rt RetentionType-nn NodenameorId
```

Description

You can use dataset backup_change_retention_type to change your dataset retention policy associated with a dataset backup.

Parameters	
-dn	
	Specifies the name of the dataset for which you want to change the backup retention policy.
-bv	
	Indicates the version of the backup for which you want to change the retention policy. Specify either the backup version or the backup ID.
-i	
	Indicates the ID of the backup for which you want to change the retention policy. Specify either the backup version or the backup ID.
-rt	
	Indicates the new retention type you want to set for your dataset backup.
-nn	
	Indicates a storage system node name or ID. If you do not specify a node name or ID, the retention type is applied to all backups.

dataset backup_delete

You can use dataset backup_delete to delete a dataset backup.

```
sdcli dataset backup_delete-dn DatasetName-bv BackupVersion-all | -local
| -remote
```

Description

You can use dataset backup_delete to delete a remote dataset backup.

Parameters

-dn

Specifies the name of the dataset containing backups you want to delete.

-bv

Specifies the version of the dataset backup you want to delete.

-all | -local | -remote

Indicates whether you want to delete local backups, remote backups, or all backups. If no option is specified, the default is -all.

dataset backup_end

You can use dataset backup_end to end a dataset backup and provide cleanup for VSS backups by providing Snapshot copy locators.

Syntax

sdcli dataset backup_end-dn DatasetName-bv NewBackupVersion

Description

Provides cleanup for VSS backups by providing Snapshot copy locators.

Parameters

-dn

Indicates the name of the dataset containing the backup you want to clean up.

-bv

Indicates the version of the new dataset backup.

dataset backup_get_metadata

You can use dataset backup_get_metadata to obtain dataset backup metadata.

Syntax

sdcli dataset backup_get_metadata-dn DatasetName-bv BackupVersion

Description

You can use dataset backup_get_metadata to obtain dataset backup metadata.

Parameters

-dn

Specifies the dataset name for which you want to view backup metadata.

-bv

Specifies the version of the dataset backup for which you want to view metadata.

dataset backup_list

You can use dataset backup_list to list your dataset backup copies. You can also list backup copies of a specified version or ID.

Syntax

sdcli dataset backup_list-dn DatasetName-version

Description

You can use dataset backup_list to list your dataset backup copies. You can also list backup copies of a specified version or ID.

Parameters

-dn

Specifies the name of the dataset for which you want to view the backup copies.

-version

Indicates the dataset version for which you want to view backup copies.

dataset backup_set_metadata

You can use dataset backup_set_metadata to initially set up and then later change the dataset backup metadata.

Syntax

```
sdcli dataset backup_set_metadata-dn DatasetName-bv BackupVersion-i
BackupId-filename-md FieldName FieldValue ...
```

Description

You can use dataset backup_set_metadata to initially set up and then later change the dataset backup metadata.

Parameters

-dn

Specifies the name of the dataset for which you want to add or change backup set metadata.

-bv

Specifies the backup version for which you want to set or change metadata. You can specify the backup version or the backup ID.

Specifies the backup ID for which you want to set or change metadata. You can specify the backup version or the backup ID.

-filename

The -filename flag indicates that the following string is a file name for the *FieldName* metadata value.

-md

Specifies the backup set metadata. Metadata is specified in *FieldName-FieldValue*pairs: for example, -md name1 value1 name2 value 2

When the -filename flag is present, you only need to specify *FieldValue*.

dataset backup_start

You can use the dataset backup_start command to start a dataset backup.

Syntax

```
sdcli dataset backup_start-dn DatasetName-desc BackupDescription-rt
RetentionType
```

Description

You can use the dataset backup_start command to start a dataset backup.

Parameters

-dn

Specifies the name of the dataset you want to back up.

-desc

Optional: used to write a description of the dataset backup you want to begin.

-rt

Specifies the retention type you want to apply to the dataset backup. Your options are hourly, daily, weekly, monthly, and unlimited. The default is unlimited. If no retention type is specified, the default is applied.

dataset backup_status

You can use dataset backup_status to view the status of a scheduled backup.

sdcli dataset backup_status-dn DatasetName-BV BackupVersionList

-i

Description

You can use dataset backup_status to view the status of a scheduled backup.

Parameters

-dn

Specifies the name of the dataset for which you want to view the backup status.

-BV

Specifies the list of backups for which you want to view a status.

dataset backup_version_convert

You can use dataset backup_version_convert to convert an existing backup version into a user-visible timestamp so you better manage backups with multiple versions.

sdcli dataset backup_version_convert-bv BackupVersion

Description

You can use dataset backup_version_convert to convert an existing backup version into a user-visible timestamp.

Parameters

-bv

Indicates the dataset backup version you want to convert to a timestamp.

dataset create

You can use dataset create to create a new dataset.

Syntax

sdcli dataset create -dn DatasetName-as ApplicationServerName-an ApplicationName-av ApplicationVersion-dd DatasetDescription

Description

You can use dataset create to create a new dataset.

Parameters

-dn

Indicates the name of the dataset you want to create.

-as

Specifies the application server on which you want to create the new dataset.

Specifies the application on which you want to create the new dataset.

-av

-an

Specifies the version of the application on which you want to create the new dataset.

-dd

Optional: adds a description of the dataset you want to create.

dataset create_local_backup

You can use dataset create_local_backup to create a local dataset backup.

Syntax

```
sdcli dataset create_local_backup-Standalone-dn DatasetName-bv
BackupVersion-SL SnapshotLocatorList
[SnapshotName1 AccessPoint1, SnapshotName2 AccessPoint2, ...] -nn
NodenameorId
```

Description

You can use dataset create_local_backup to create a local dataset backup.

Parameters

-Standalone

Specifies that you want to create a stand-alone dataset backup.

-dn

Specifies the name of the dataset for which you want to create a local backup.

-SL

The Snapshot Locator List is composed of space-separated pairs of *SnapshotName* and *AccessPoint*.

-nn

Specifies the node name or ID of the system on which you want to create a local backup. If you do not specify a node name or ID, the default is your primary system.

dataset delete

You can use dataset delete to delete a specified dataset.

sdcli dataset delete-dn DatasetName

Description

You can use dataset delete to delete a specified dataset.

Parameters

-dn

Specifies the dataset that you want to delete.

dataset dfm_request

You can use dataset dfm_request to send requests to Protection Manager for backup version information.

sdcli dataset dfm_request-filename FileNameAndPath

Description

You can use dataset dfm_request to send requests to Protection Manager for backup version information. If you do not specify the version of the backup copy, information about all versions of a dataset backup is retrieved.

Parameters

-filename

Specifies the name and path of the file.

dataset get_available_policies

You can use dataset get_available_policies to view information about the storage policies available for your dataset.

sdcli dataset get_available_policies-dn DatasetName

Description

You can use dataset get_available_policies to view information about the storage policies available for your dataset.

Parameters

-dn

Specifies the name of the dataset for which you want available storage policy information.

dataset get_backup_location

You can use the dataset get_backup_location command to locate your dataset backup copy.

Syntax

sdcli dataset get_backup_location-dn DatasetName-i BackupId

Description

You can use dataset get_backup_location to locate your dataset backup copy.

Parameters

-dn

Indicates the name of the dataset backup copy you want to locate.

-i

Indicates the backup ID of the dataset backup copy you want to locate. You must have the backup ID to locate the backup copy.

dataset get_backup_version_info

You can use dataset get_backup_version_info to obtain backup version information for your dataset.

Syntax

sdcli dataset get_backup_version_info-dn DatasetName-bv BackupVersion

Description

You can use dataset get_backup_version_info to obtain backup version information for your dataset.

Parameters

-dn

Specifies the name of the dataset for which you want to get backup version information.

-bv

Specifies the dataset backup version about which you want information. If you do not indicate a backup version, you receive information about all available backup versions.

dataset get_metadata

You can use dataset get_metadata to obtain information about specified datasets.

Syntax

sdcli dataset get_metadata-dn DatasetName

Description

You can use dataset get_metadata to obtain information about specified datasets.

Parameters

-dn

Indicates the name of the dataset for which you want to retrieve metadata.

dataset get_policy

You can use dataset get_policy to get the name of the dataset storage policy currently in use for your dataset.

Syntax

sdcli dataset get_policy-dn DatasetName

Description

You can use dataset get_policy to get the name of the dataset storage policy currently in use for your dataset.

Parameters

-dn

Indicates the name of the dataset for which you want to view storage policy information.

dataset get_retention_info

You can use dataset get_retention_info to get dataset backup retention policy information.

Syntax

sdcli dataset get_retention_info-dn DatasetName-rt RetentionType

Description

You can use dataset get_retention_info to get dataset backup retention policy information.

Parameters

-dn

Indicates the name of the dataset for which you want backup retention policy information.

-rt

Indicates the type of backup retention policy you want to see information about. Your options are daily, weekly, monthly, and all. If no retention type is provided, you retrieve information about monthly backup retention policies.

dataset info

You can use dataset info to display information about a specified dataset.

Syntax

```
sdcli dataset info-dn DatasetName-an ApplicationName-av
ApplicationVersion
```

Description

You can use dataset info to display information about your specified dataset. You can also request information about datasets managed by a specific application.

Parameters

-dn

Indicates the dataset about which you are requesting information.

-an

Indicates the name of the application for which you want to retrieve dataset information. You can specify either the application name or the application version.

av

Indicates the version of the application for which you want to retrieve dataset information. You can specify either the application version or the application name.

dataset initiate_conformance

You can use dataset initiate_conformance to initiate a conformance check.

sdcli dataset initiate_conformance-dn DatasetName

Description

You can use the dataset initiate_conformance command to initiate a conformance check.

Parameters

-dn

Indicates the name of the dataset for which you want to start a conformance check.

dataset list_members

dataset list_members displays the members of the specified dataset.

Syntax

sdcli dataset list_members[-dn DatasetName

Description

dataset list_members displays the members of the specified dataset.

Parameters

-dn

Specifies the name of the dataset for which you want to list the members.

dataset mount_backup

You can use dataset mount_backup to mount a LUN in a Snapshot copy.

Syntax

```
sdcli dataset mount_backup-dn DatasetName-i BackupID-bv BackupVersion-
SAP SourceAccessPointList-DAP DestinationAccessPointList-S
PrimarySnapshotList-dtype {shared | dedicated}-e ResourceGroupName
```

Description

You can use dataset mount_backup to mount a LUN in a Snapshot copy.

Parameters

-dn

Specifies the name and path of the dataset backup you want to mount.

-i

Specifies the ID of the dataset backup you want to mount.

-bv

	Specifies the version of the dataset backup you want to mount.
-SAP	
	Specifies the source access point of the dataset backup you want to mount.
-DAP	
	Specifies the destination access point to which you want to mount the dataset backup.
-s	
	Provides a list of the primary Snapshot copies in the dataset you want to mount.
-dtype	
	Specifies whether you want to mount your dataset backup to a shared or dedicated LUN.
-e	

Specifies the name of the resource group for the dataset backup you want to mount.

dataset protect

You can use dataset protect to secure your dataset.

Syntax

```
sdcli dataset protect-dn DatasetName-pt Protect {ON, OFF}
```

Description

You can use dataset protect to secure your dataset.

Parameters

-dn

Indicates the name of the dataset you want to protect.

-pt

Indicates whether you want to protect your dataset. If no -pt value is specified, your dataset is protected by default.

dataset remove_members

You can use dataset remove_members to remove objects from your specified dataset.

Syntax

sdcli dataset remove_members-dn DatasetName-AP AccessPointList

Description

You can use dataset remove_members to remove objects from your specified dataset.

Parameters

-dn

Indicates the dataset from which you want to remove members.

-AP

Specifies a list of type-value pairs: for example, $G: \ mp$ I: $\ mp$, where the type mp denotes mount point and $G: \$ is the specifier.

dataset restore

You can use the dataset restore command to restore data from a dataset backup.

Syntax

```
sdcli dataset restore-dn DatasetName-bv BackupVersion-i BackupId-SAP SourceAppointedList-S PrimarySnapshotList
```

Description

You can use the dataset restore command to restore data from a dataset backup.

Parameters

-dnSpecifies the name of the dataset you want to restore.-bvSpecifies the version of the dataset backup from which you want to restore. You can specify either the backup version or the backup ID.-iSpecifies the ID of the dataset backup from which you want to restore. You can specify either the backup version or the backup ID.SAPSpecifies the source-appointed list for the dataset backup from which you want to restore.-sSpecifies the primary Snapshot copy list for the dataset backup from which you want to restore.

dataset restore_status

You can use dataset restore_status to view the status of a backup restore operation.

Syntax

```
sdcli dataset restore status-job JobID
```

Description

You can use dataset restore_status to view the status of a backup restore operation.

Parameters

-job

Specifies the restore operation ID for which you want to view status.

dataset set_metadata

dataset set_metadata adds metadata to a specified dataset.

Syntax

```
sdcli dataset set_metadata[-dn DatasetName][-filename][-md Field1 Value1 Field2 Value2 ...]
```

Description

The dataset set_metadata command adds metadata to the dataset you specify.

Parameters

-dn DatasetName

Indicates the name of the dataset to which you want to add metadata.

-filename

When the -filename flag is specified, indicates *FieldName* value for the dataset metadata you are setting. When filename is not set, you must include *FieldName* for every *FieldValue* metadata element you specify.

-md

-md enables you to specify a *FieldName FieldValue* metadata pair, directly on the command line when -filename is absent: for example, -md *name1 value1 name2 value2*.

dataset set_policy

You can use dataset set_policy to indicate the name of a dataset storage policy you want to apply to a dataset.

Syntax

sdcli dataset set_policy-dn DatasetName-pn PolicyName

Description

You can use dataset set_policy to indicate the name of a dataset storage policy you want to apply to a dataset.

Parameters

-dn

Specifies the name of the dataset to which you want to apply the storage policy.

-pn

Specifies the name of the storage policy you want to apply to your specified dataset.

dataset transfer_now

You can use dataset transfer_now to transfer a dataset to a secondary storage system, while that dataset is waiting for a local backup operation.

sdcli dataset transfer_now-dn DatasetName-rt RetentionType

Description

You can use dataset transfer_now to transfer a dataset to a secondary storage system, while that dataset is waiting for a local backup operation.

Parameters

-dn

Specifies the name of the dataset you want to copy to a secondary storage system.

-rt

Specifies the retention type for the dataset you have copied to the secondary storage system. Your options are hourly, daily, weekly, monthly, and unlimited. The default is unlimited; if -rt is not specified, the default is applied.

dataset vss_backup_end

You can use dataset vss_backup_end to stop a VSS backup of your dataset.

Syntax

```
sdcli dataset vss_backup_end-dn DatasetName-bv BackupVersion-AP AccessPointList
```

Description

You can use dataset vss_backup_end to stop a VSS backup of your dataset.

Parameters

-dn

Specifies the name of the dataset for which you want to end the VSS backup.

-bv

Provides the version of the backup operation you want to stop.

-AP

Specifies the access point for the VSS backup operation you want to stop.

dataset vss_backup_prepare

You can use dataset vss_backup_prepare to prepare your dataset for VSS backup by providing Snapshot copy locators.

Syntax

```
sdcli dataset vss_backup_prepare-dn DatasetName-bv BackupVersion-AP AccessPointList-nn NodenameorId
```

Description

You can use dataset vss_backup_prepare to prepare your dataset for VSS backup by providing Snapshot copy locators.

Parameters

-dn

Specifies the dataset you want to prepare for VSS backup.

-bv

Specifies the version of the dataset backup you want to initiate.

-AP

Specifies the access point list for the dataset VSS backup you want to initiate.

-nn

Specifies the storage system node name or ID. If this parameter is omitted, the default is the primary storage system.

License commands

The sdcli utility provides command-line support for SnapDrive license operations.

The license set command

license set sets the license key for the specified module.

Syntax for this command is:

sdcli license set -module ModuleName -key LicenseKey

Example

sdcli license set -module LPSM -key ABCDEFGHIJKLMN

The license list command

license list displays all SnapDrive licenses installed.

Syntax for this command is:

sdcli license list

The license remove command

license remove removes an existing server license.

Syntax for this command is:

sdcli license remove -module ModuleName

Example

sdcli license remove -module LPSM

Initiator group management commands

The sdcli utility provides command-line support for initiator group management.

The igroup list command

The igroup list command displays all igroups on a storage system that have initiators on the local host or, if specified, a remote host.

Syntax for this command is:

sdcli igroup list [-m MachineName] -f StorageSystem

-f specifies the storage system name or IP address for which the igroups will be listed.

Examples

sdcli igroup list -f 172.17.167.45

The preceding example displays the igroup list for the storage system with the IP address 172.17.167.45.

```
sdcli igroup list -m server3 -f v34filer
```

The preceding example displays the igroup list for the storage system v34filer, which has initiators on the remote host server3.

The igroup create command

The igroup create command enables you to create a new igroup.

Syntax for this command is:

```
sdcli igroup create [-m MachineName] -f StorageSystem -I
NodeName InitiatorName -ig igroupName
```

-I lists the machine name and initiator name in pairs.

Note: A new igroup is created for only one machine, so you must specify the same machine name for each pair.

-ig specifies the name of the igroup you are creating.

Example

```
sdcli igroup create -f v34filer -I server3 10:00:00:00:c9:48:c9:5d
server3 10:00:00:00:c9:48:c9:5e -ig v3group1
```

The preceding example creates a new igroup called v3group1 on a storage system called v34filer for two initiators, 10:00:00:c9:48:c9:5d and 10:00:00:c9:48:c9:5e, on a host called server3.

The igroup rename command

The igroup rename command enables you to rename an existing igroup. Syntax for this command is:

```
sdcli igroup rename [-m MachineName] -f StorageSystem -ig igroupName -igNew
igroupNewName
```

-ig specifies the name of the existing igroup you are renaming.

-igNew specifies the new name of the igroup.

Example

sdcli igroup rename -f 172.17.167.45 -ig v3group1 -igNew v3group1fc

The preceding example renames an igroup from v3group1 to v3group1fc on a storage system with the IP address 172.17.167.45.

The igroup delete command

The igroup delete command enables you to delete an existing igroup if there is no LUN mapped to it.

Syntax for this command is:

```
sdcli igroup delete [-m MachineName] -f StorageSystem -ig igroupName
```

-ig specifies the name of the igroup you want to delete.

Example

```
sdcli igroup delete -f 172.17.167.45 -ig v3group1fc
```

The preceding example deletes the igroup named v3group1fc from the storage system with the IP address 172.17.167.45.

Fractional space reservation monitoring commands

The sdcli utility provides command-line support for fractional space reservation monitoring.

The spacemon list command

The spacemon get command displays the space reservation monitoring settings for the specified host.

Syntax for this command is:

sdcli spacemon list {-m MachineName}

MachineName is the machine name on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

The spacemon set command

The spacemon set command sets the space reservation monitoring settings for the specified host. Syntax for this command is:

sdcli spacemon set -mi Monitoring_interval -f StorageSystem -vn VolumeName
{-m MachineName} -rap Threshold_for_Reserved_Available_Percentage -roc
Threshold_for_Rate_of Change -ccs true|false

Monitoring_interval is the frequency, in minutes, at which you want to monitor fractional space available.

StorageSystem is the name of the storage system on which the LUNs reside.

VolumeName is the name of the volume you want to monitor.

Threshold _for_Reserved_Available_Percentage is the point at which you want to be warned of a low space reservation condition.

Threshold_for_Rate_of Change is the point at which you want to receive a notification. Use kb, mb, gb, or tb to specify the value as kilobytes, megabytes, gigabytes, or terabytes.

-ccs is used to monitor whether a Snapshot copy can be created. True indicates that you want to monitor whether a Snapshot copy can be created. False indicates that you do not want to monitor whether a Snapshot copy can be created.

MachineName is the machine name on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

Example

```
sdcli spacemon set -mi 30 -f controller1 -vn testvol -rap 90 -roc 500mb -ccs true
```

The preceding example shows that fractional space reservations will be monitored every 30 minutes on the volume named testvol on controller1. The threshold for testvol is 90 percent of

the reserved available percentage and the threshold for rate of change is 500 MB. SnapDrive will verify storage system and volume names and that space is available for Snapshot copies to be created.

The spacemon snap_delta command

The spacemon snap_delta command displays the rate of change between two Snapshot copies or between a Snapshot copy and the active file system of the storage system volume. Syntax for this command is:

```
sdcli spacemon snap_delta -f StorageSystem -vn VolumeName -s1 snapshot1 -s2 snapshot2 {-m MachineName}
```

StorageSystem is the name of the storage system on which the volume exists.

VolumeName is the name of the volume for which you want to display the snap delta.

snapshot1 is the name of the Snapshot copy you want to compare with either a second Snapshot copy or with the active file system.

snapshot2 is name of the second Snapshot copy.

MachineName is the machine name on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

The spacemon snap_reclaimable command

The spacemon snap_reclaimable command displays the space that can be reclaimed by deleting a Snapshot copy.

Syntax for this command is:

```
sdcli spacemon snap_reclaimable -f StorageSystem -vn VolumeName -s snapshot
```

StorageSystem is the name of the storage system on which the volume exists.

VolumeName is the name of the volume on which the Snapshot copy resides.

snapshot is the name of the Snapshot copy for which you want to view reclaimable space.

The spacemon vol_info command

The spacemon vol_info command displays information about fractional space reserved volumes. Syntax for this command is:

```
sdcli spacemon vol_info {-m MachineName}
```

MachineName is the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

Note: Output for the spacemon vol_info command is displayed in XML format.

The spacemon delete command

The spacemon delete command enables you to delete the fractional space reservation monitor settings for the specified storage system volume. Syntax for this command is:

sdcli spacemon delete -f StorageSystem -vn VolumeName {-m MachineName}

StorageSystem is the name of the storage system on which the volume exists.

VolumeName is the name of the volume from which you want to delete fractional space reservation settings.

Virtual Storage Console commands

The sdcli utility provides command-line support for managing communication between SnapDrive for Windows and Virtual Storage Console.

The vsc_config list command

The vsc_config list command displays Virtual Storage Console configuration settings. Syntax for this command is:

sdcli vsc_config list

The vsc_config set command

The vsc_config set command sets the where Virtual Storage Console server IP address and port. Syntax for this command is:

```
sdcli vsc_config set -host host [-vscport VSC_port_number] [-port SnapDrive
port] [-m MachineName]
```

-host specifies the host server name or IP address where Virtual Storage Console is registered to the vSphere Client.

-vscport specifies the where Virtual Storage Console service port number you use to communicate with where Virtual Storage Console. The default port is 8043.

-port specifies the Web service port number you use to communicate with SnapDrive. The default port is 808.

The vsc_config delete command

The vsc_config delete command deletes the specified Virtual Storage Console server entry. Syntax for this command is:

```
sdcli vsc_config delete [-m MachineName] [-port SnapDrive port]
```

-port specifies the Web service port number you use to communicate with SnapDrive. The default port is 808.

Space reclamation commands

The sdcli utility provides command-line support for space reclamation operations.

The spacereclaimer start command

The spacereclaimer start command starts the SnapDrive space reclamation process. The space reclamation process initiates SnapDrive space optimization and significantly improves performance.

Syntax

sdcli spacereclaimer start[-m MachineName][-d MountPoint][-t TimetoRun]

Description

After you start space reclamation using spacereclaimer start, Space Reclaimer continues to run, even if the analyzer determines that there is no space to reclaim.

Parameters

m

Specifies the *MachineName* on which you want to start the space reclamation process.

-d

Specifies the LUN *MountPoint* volume name or CSV reparse point on which you want to start Space Reclaimer.

-t

Specifies the amount of time (*TimeToRun*) that you want Space Reclaimer to run on the specified LUN. Specify a time from 1 to 10080 minutes (7 days).

Example

sdcli spacereclaimer start -d C:\ClusterStorage\Volume8

This example starts the space reclamation process by specifying the CSV reparse point of a CSV disk.

The spacereclaimer stop command

The spacereclaimer stop command stops the space reclamation process. Syntax for this command is:

sdcli spacereclaimer stop [-m MachineName]-d MountPoint
-d *MountPoint* specifies the LUN mount point, volume name, or CSV reparse point on which you want to stop Space Reclaimer.

Example

```
sdcli spacereclaimer stop -d \\?\Volume{0944ca87-be05-45ad-8606-ba13ee7388a0}
```

The preceding example stops the space reclamation process on a CVS volume.

The spacereclaimer analyze command

The spacereclaimer analyze command checks whether space reclamation is needed for the LUN specified.

Syntax for this command is:

```
sdcli spacereclaimer analyze [-m MachineName] -d MountPoint
```

-d *MountPoint* specifies the LUN mount point, volume name, or CSV reparse point on which you want to analyze.

The spacereclaimer status command

The spacereclaimer status command displays the space reclamation status for the LUN specified.

If you specified a time to run when you started Space Reclaimer, the status displays the number of minutes remaining. If no time was specified, the status displays the percentage of space remaining for space reclamation.

Syntax for this command is:

```
sdcli spacereclaimer status [-m MachineName] [-D MountPointList]
```

-DMountPointList specifies a list of LUN mount points. This list is optional. If no mount points are specified, SnapDrive displays the status for all Space Reclaimer operations.

Preferred IP address commands

The sdcli utility provides command-line support for managing preferred storage system IP addresses.

The preferredIP set command

The preferredIP set command sets the SnapDrive preferred IP address for the specified storage system.

Syntax for this command is:

sdcli preferredIP set -f StorageSystem -IP PreferredIPAddress

Example

```
sdcli preferredIP set -f Storage1 -IP 172.18.53.94
```

The preceding example sets the SnapDrive preferred IP address for the storage system named Storage1 to 172.28.53.94.

The preferredIP list command

The preferredIP list command displays the storage system names and IP addresses that you set as the preferred IP addresses for SnapDrive to use for management traffic. Syntax for this command is:

sdcli preferredIP list

The preferredIP delete command

The preferredIP delete command deletes the preferred IP address for the specified storage system.

Syntax for this command is:

```
sdcli preferredIP delete -f StorageSystem
```

iSCSI connection commands

The sdcli utility provides command-line support for managing connections to iSCSI targets.

The iscsi_target disconnect command

The iscsi_target disconnect command disconnects the specified iSCSI initiator from the specified iSCSI target on all portals.

Syntax for this command is:

```
sdcli iscsi_target disconnect -t TargetName
```

Example

sdcli iscsi_target disconnect -t iqn.1992.08.com.ibm:sn.33604307

The preceding example disconnects the specified iSCSI target.

The iscsi_target list command

The iscsi_target list command displays a list of all iSCSI targets. For each target, the command displays all portals through which the target is available or to which the target is connected.

Syntax for this command is:

sdcli iscsi_target list {-f Storage_System | -i InitiatorPortName}

-f displays all targets on the specified storage system.

Example

sdcli iscsi_target list -f Storage2

The preceding example lists all the iSCSI targets on the Storage2 storage system, as well as all portals those targets are available through or connected to.

iSCSI initiator commands

The sdcli utility provides command-line support for managing iSCSI initiators.

The iscsi_initiator list command

The iscsi_initiator list command displays a list of all iSCSI sessions on the specified machine.

Syntax for this command is:

sdcli iscsi_initiator list {-m MachineName} -s

MachineName is the machine name on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

-s enumerates the iSCSI sessions.

The iscsi_initiator establish_session command

The iscsi_initiator establish_session command establishes a session with a target using the specified HBA.

Syntax for this command is:

```
sdcli iscsi_initiator establish_session {-m MachineName} {-h HBA_ID} {-hp
HBA_Portal_ID} -t TargetName -np IPAddress IPPort {-c CHAPName
CHAPPassword}
```

-h HBA_ID is used to establish the iSCSI session. The HBA ID can be obtained by using the sdcli sysconfig list command.

-hp HBA Portal ID is used to specify the portal on the iSCSI HBA to be used to establish the iSCSI session. The HBA Portal ID can be obtained by using the sdcli sysconfig list command.

-t TargetName is the name of the iSCSI target.

-np IP Address IPPort specify the IP address and IP port of the network portal on the target. The IP Port can be obtained by using the sdcli iscsi_initiator list command.

Example

sdcli iscsi_initiator establish_session -h 0 -t iqn. 1992-8.com.ibm:maya -np 172.18.53.94 3260

The preceding example establishes an iSCSI session with the specified target using the specified HBA ID.

The iscsi_initiator terminate_session command

The iscsi_initiator terminate_session command terminates the session. Syntax for this command is:

```
sdcli iscsi_initiator terminate_session {-m MachineName} -s Session_ID
```

MachineName is the machine name on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

-s Session_ID is the session ID of the session you want to terminate.

Example

```
sdcli iscsi_initiator terminate_session -s
0xffffffff868589cc-0x4000013700000006
```

The preceding example terminates the specified iSCSI session on the local machine.

LUN commands

The sdcli utility provides command-line support for managing LUNs in SnapDrive.

The disk create command

The disk create command creates a new LUN.

Syntax

```
sdcli disk create
[-m MachineName]
-d MountPoint
-p LUNpath
-z DriveSize
[-rs Reserve Snapshot Space y | n]
[-I NodeMachineName InitiatorName +] | [-IG NodeMachineName GroupName +]
-dtype {shared | dedicated}
[-ds "datastore name"
{[-e "ResourceGroupName"] | [-n "ResourceGroupName" "ResourceGroupDesc"]
    [-csv ]}
-port "PortNumber"
```

-passthrough

Parameters

-rs enables you to limit the maximum disk space of the LUN you are creating to allow for at least one Snapshot copy on the volume.

-ds enables you to specify the name of a VMFS datastore on which to store the VMDK file if you are creating a LUN on an ESX server guest OS. When the -ds option is not specified, the VMDK file is stored by default on the same VMFS datastore as the virtual machine. If the virtual machine resides on an NFS datastore, you must specify a VMFS datastore for the VMDK file.

-e enables you to specify the name of an existing resource group.

-n enables you to specify the name of a new resource group.

-csv enables you to add a shared disk to a cluster shared volume.

-passthrough enables you to dynamically add a pass-through disk. This option is applicable only when a Hyper-V virtual machine is present.

Examples

The following example creates a dedicated, 1-GB LUN named mktng.lun in the storage2 volume named sd_vds_only. Next, it connects this LUN to the host as drive R:.

```
sdcli disk create -dtype dedicated -z 1024 -p storage2:/vol/
sd_vds_only/mktng.lun -d R -I host3 10:00:00:00:C9:2B:FD:12
```

The following example creates a shared, 4-GB LUN on host4 (the local machine running the SDCL command) and maps it to drive R:, using a pair of initiators.

```
sdcli disk create -p \\133.25.61.62\sd_vds_only\mktng.lun -d r -z 4096
-dtype shared -e "mktng" -I host4 10:00:00:00:C9:2B:FD:12 host4
10:00:00:C9:2B:FD:11 host5 10:00:00:C9:2B:FC:12 host5
10:00:00:C9:2B:FC:11
```

The following example creates a 65-MB RDM LUN in the NFS volume named vol1_nfs, and maps the LUN to drive T. The VMFS datastore named "big guy" is specified to store the VMDK file.

```
sdcli disk create -d T -p rlabf6:/vol/vol1_nfs/a -z 65MB -IG vmware229
igroup_fcp -dtype dedicated -ds "big guy"
```

The disk connect command

The disk connect command connects a LUN to a host by mapping the LUN to a Windows drive letter.

Syntax

```
sdcli disk connect
[-m MachineName]
-p LUNpath
-d MountPoint
[-I NodeMachineName InitiatorName ...] | [-IG
NodeMachineName GroupName ...]
-dtype {shared | dedicated}
-ds "datastore name"
{[-e "ResourceGroupName"] | [-n "ResourceGroupName" "ResourceGroupDesc"]
| [-csv] } [-c "ClusterName"]
-port PortName -passthrough
```

Parameters

-ds enables you to specify the name of a VMFS datastore on which to store the VMDK file if you are creating a LUN on an ESX server guest OS. When the -ds option is not specified, the VMDK file is stored by default on the same VMFS datastore as the virtual machine. If the virtual machine resides on an NFS datastore, you must specify a VMFS datastore for the VMDK file.

-csv enables you to add a shared disk to a cluster shared volume.

-passthrough enables you to dynamically connect a pass-through disk. This option is only applicable when a Hyper-V virtual machine is present.

Note:

- The SnapDrive CLI fails to notify you about an active SnapMirror relationship between the source and the destination. The disk connect command connects a LUN to a host and the SnapMirror relationship is broken.
- When you are trying to connect to a LUN using this command, SnapDrive might not store the RDM LUN in the datastore where a virtual machine is hosted.

Examples

The following example connects a LUN in the storage2 volume sd_vds_only and named mktng.lun, which belongs to the MSCS cluster resource group tech_mktng on the mktng cluster.

```
sdcli disk connect -d s -dtype shared -p storage2:/vol/sd_vds_only/
mktng.lun -I host3 10:00:00:C9:2B:FD:1B host3 10:00:00:C9:2B:FD:
1C host4 10:00:00:C9:2B:FD:12 host4 10:00:00:C9:2B:FD:11 -e
"tech_mktng" -c "mktng"
```

The following example connects an RDM LUN on the NFS volume named vol1_nfs. The VMFS datastore named "big guy" is specified to store the VMDK file.

```
sdcli disk connect -d T -p rlabf6:/vol/vol1_nfs/a -z 65MB -IG
vmware229 igroup_fcp -dtype dedicated -ds "big guy"
```

The following example connects a CSV LUN called csvLun12 as a shared disk to the cluster.

```
sdcli disk connect -p ab-270-2:/vol/sdw_vol1/csvLun12 -I CLAB-A9-8
iqn.1991-05.com.microsoft:clab-a9-8.sddev.ibm.com CLAB-A9-7 iqn.
1991-05.com.microsoft:clab-a9-7.sddev.ibm.com -dtype shared -csv
```

The disk delete command

The disk delete command deletes a LUN. The LUN must be connected (mapped to a Windows drive letter or mount point) for the command to succeed. Syntax for this command is:

```
sdcli disk delete [-m MachineName] {-p LUNpath | -d MountPoint}
```

Example

```
sdcli disk delete -p \\133.25.61.62\sd_vds_only\mktng.lun
```

The preceding example deletes the LUN mktng.lun from the sd_vds_only volume on the storage system identified by the IP address 133.25.61.62.

```
sdcli disk delete -d \\?\Volume{f1816466-f1d8-4b96-b547-2ce12415aee4}\
```

The preceding example specifies the CSV volume name to delete a CSV disk.

```
sdcli disk delete -d C:\ClusterStorage\Volume8
```

The preceding example specifies the CSV reparse point to delete a CSV disk.

The disk disconnect command

The disk disconnect command disconnects a LUN from the host. The LUN must be connected (mapped to a Windows drive letter) for the command to succeed.

Note: You must make sure that the LUN you are disconnecting is not monitored with the Windows Performance Monitor (perfmon).

Syntax for this command is:

```
sdcli disk disconnect [-m MachineName] {-p LUNpath | -d MountPoint} [-f]
```

Attention: The -f switch causes the LUN to be forcibly unmounted, even if an application or the Windows operating system is using it. Therefore, use this feature with extreme care.

Example

sdcli disk disconnect -d z

The preceding example disconnects the LUN mapped to the drive letter "Z:" on the SnapDrive host running the sdcli command.

```
sdcli disk disconnect -p \\storage2\sd_vds_only\mktng.lun -f
```

The preceding example forces disconnection of the LUN mktng.lun, which is in the sd_vds_only volume on storage2. Because the -f switch is being used, all open files in the LUN might be lost or corrupted.

The disk resize command

The disk resize command increases or decreases the disk by a user-specified size, as long as that figure falls within the SnapDrive-specified minimum and maximum values. Syntax for this command is:

```
sdcli disk resize
[ -m MachineName ]
```

{ -p LUNpath | UNCPath

```
-d MountPoint}
```

```
-z DriveSizeIncrement
```

```
[ -s SnapshotName]
```

```
[ -x LunSnapshot ]
```

-m MachineName

indicates the machine on which you want to execute the disk resize operation. If no machine name is specified, the command executes on your local machine.

-p LUNPath | UNCPath

specifies the LUN path include the storage system name or the UNC path to the location of the LUN file on the storage system.

-d MountPoint

specifies the mount point volume name or CSV reparse point of the LUN.

DriveSizeIncrement is the amount by which you want to increase or decrease the size of the disk. When decreasing the size of the disk, use a dash (-) before the amount to indicate a negative value. You can specify a postfix of MB, GB, or TB. If no postfix is specified, the default is MB.

-s SnapshotName

specifies the Snapshot copy name to use when making a Snapshot copy for restore.

-x LunSnapshot

specifies that a Snapshot copy is taken only for the LUNs explicitly specified in the *MountPointList*. When no value is specified, a Snapshot copy is made of all the LUN drives on the storage system volume.

Examples

```
sdcli disk resize -d G: -z -500
```

The preceding example shrinks the size of the disk mapped to "G:" by 500 megabytes.

```
sdcli disk resize -d P: -z 1GB
```

The preceding example increases the LUN mapped to "P:" by 1 GB. (In practice, SnapDrive expands the disk by the amount specified by -z, plus a certain increment required for system overhead.)

```
sdcli disk resize -d \\?\Volume{f1816466-f1d8-4b96-b547-2ce12415aee4}\
-z +100
```

The preceding example increases a CSV disk by 100 MB using the CSV volume name.

The disk expand command

The disk expand command expands the disk by a user-specified size, as long as that figure falls within the minimum and maximum values required by SnapDrive. Syntax for this command is as follows:

```
sdcli disk expand [-m MachineName] {-p LUNpath | -d MountPoint} -z DriveSizeIncrement [-rs Reserve Snapshot copy space \{y/n\}]
```

Section

-z DriveSizeIncrement is the amount by which you want to expand the disk. You can specify a postfix of MB, GB, or TB. If no postfix is specified, the default is MB.

-rs enables you to limit the maximum disk space for the disk you are expanding, to allow for at least one Snapshot copy on the volume. This option is recommended.

Example

The following example increases the LUN mapped to "P:" by 1 GB. (In practice, SnapDrive expands the disk by the amount specified by -z, plus a certain increment required for system overhead.)

```
sdcli disk expand -d p -z 1GB
```

The following example expands a CSV disk by 100 MB on the CSV reparse point C: \ClusterStorage\Volume8.

sdcli disk expand -d C:\ClusterStorage\Volume8 -z 100

disk add_initiator

You can use the disk add_initiator command to add a new initiator to a LUN. Syntax for this command is:

```
sdcli disk add_initiator [-m MachineName] {-p LUNpath | -d MountPoint} -i
InitiatorPortName
```

Examples

The following example adds an initiator to a LUN mapped to drive E: on the SnapDrive host from which the sdcli command was executed.

```
sdcli disk add_initiator -d E -i 21:00:00:e0:8b:85:19:ba
```

The following example adds an initiator to a LUN located at the path $\$ $\sdwatf2\sd$

sdcli disk add_initiator -p \\sdwatf2\sdwatf2_vol1\sdwath2_EEE.lun -i 21:00:00:e0:8b:85:19:ba

The following example specifies the CSV reparse point C:\ClusterStorage\Volume4 to add an initiator to a CSV disk.

```
sdcli disk add_initiator -d C:\ClusterStorage\Volume4 -i 50:0a:
09:80:85:f4:69:37
```

disk remove_initiator

The disk remove_initiator command removes an initiator from the specified LUN. Syntax for this command is:

```
sdcli disk remove_initiator [-m MachineName] {-p LUNpath | -d MountPoint} -
i InitiatorPortName
```

Examples

sdcli disk remove_initiator -d E -i 21:00:00:e0:8b:85:19:ba

The preceding example removes an initiator from a LUN mounted on drive letter "E:" on the SnapDrive host running the SDCLI command.

```
sdcli disk remove_initiator -p \\sdwatf2\sdwatf2_vol1\sdwath2_EEE.lun
-i 21:00:00:e0:8b:85:19:ba
```

The preceding example removes an initiator from a LUN located at the path \\sdwatf2\sdwatf2_voll\sdwath2_EEE.lun on the SnapDrive host from which the SDCLI command was executed.

```
sdcli disk remove_initiator -d \\?\Volume{24cd94c5-
d201-4474-9246-2ad491b155ea}\ -i 50:0a:09:80:85:f4:69:37
```

The disk list command

The disk list command displays a list of all the LUNs connected to the host. Syntax for this command is:

sdcli disk list [-m MachineName]

Example

sdcli disk list

The preceding example lists all the SnapDrive LUNs mapped to drive letters on the local host. The disk list command also provides the following information for each LUN:

- LUN path (storage system name, share name, virtual disk file name, and might also include qtree name)
- Storage System
- Storage System Path (storage system-side path, which includes volume name and LUN name)
- Hyper-V VHD present (a Hyper-V VHD exists)
- Hyper-V VM name
- Type
- Disk serial number
- Backed by Snapshot copy (if this is a LUN in a Snapshot copy, this displays the storage system-side path to the Snapshot copy)
- Shared (whether the disk is dedicated or shared)
- CSV disk
- Boot or System Disk
- SCSI port
- Bus
- Target
- LUN
- Read only
- Disk size (in megabytes)
- SnapMirror source
- SnapVault primary
- Disk Partition Style (either MBR or GPT)
- Clone Split Restore status
- Disk ID
- Volume name
- Mount points (the drive letter and path to which the LUN is mapped on the host)

156 | SnapDrive 7.0 for Windows Administration Guide for SAN Environments

- · CSV reparse point
- IP Addresses (IP addresses on the target storage system)
- Initiator name

Note: If you are using ESX 3.5, the initiator name field might be blank for RDM LUNs.

The disk add_mount command

The disk add_mount command adds a volume mount point. Syntax for this command is:

```
sdcli disk add_mount {-m MachineName} -vn Volume_Name -mp
Volume_Mount_Point {-create_folder}
```

Volume_Name is the name of the volume that you are trying to add or move. The volume name can be located in the output from the disk list command.

Volume_Mount_Point is the location you want to mount the LUN. This can also be a drive letter.

-create_folder indicates that a folder should be created for the new mount point if one does not already exist.

Example

```
sdcli disk add_mount -vn \\?
\Volume{db6160d8-1f14-11da-8ef3-000d5671229b} -mp G:\mount_vol1 -
create_folder
```

The disk remove_mount command

The disk remove_mount command removes a volume mount point or drive letter.

Note: This command will not delete the folder that was created at the time the volume mount point was added. After you remove a mount point, an empty folder will remain with the same name as the mount point you removed.

Syntax for this command is:

```
sdcli disk remove_mount {-m MachineName} -vn VolumeName -mp
Volume_Mount_Point
```

The disk rename_flexclone command

The disk rename_flexclone command enables you to rename the FlexClone volume from the default.

Syntax for this command is:

sdcli disk rename_flexclone

-d MountPoint

-n NewName

Mount Point is the name of the LUN in the FlexClone volume.

NewName is the new name that you can assign for the FlexClone volume. The volume name can be located in the output from the disk list command.

Note: You cannot rename VMDK volumes in SnapDrive for Windows because Virtual Storage Console is not supported.

Snapshot copy commands

The sdcli utility provides command-line support for managing Snapshot copies of SnapDrive LUNs.

The snap create command

The snap create command creates a new Snapshot copy of the specified LUNs on the SnapDrive system.

Syntax for this command is:

sdcli snap create [-m MachineName] -s SnapshotName -D MountPointList [...] [-x]

-x causes data to be flushed and consistent Snapshot copies to be created only for the drives and mount points specified by the -D switch. Otherwise, SnapDrive flushes data and creates consistent Snapshot copies for all LUNs connected to the host and residing on storage system volumes.

Note: Snapshot copies are created at the volume level. When a Snapshot copy is created using -x with the -D switch, Snapshot copies are also created for any additional disks mapped to the host that reside on the same volumes as the disks specified. Snapshot copies for the unspecified disks are dimmed in the SnapDrive MMC because they are inconsistent.

Example

```
sdcli snap create -s Jun_13_03 -D j k l
```

The preceding example creates a Snapshot copy named Jun_13_03 for each volume containing one or more of the LUNs mapped to the specified drives (that is, J:, K:, and L:). The Snapshot copies created are consistent for all LUNs contained by those volumes.

```
sdcli snap create -s csvLun11_051109 -D C:\ClusterStorage\Volume8
```

The preceding example creates a Snapshot copy named csvLUN11_051109 on the CSV reparse point C:\ClusterStorage\Volume8.

The snap delete command

The snap delete command deletes an existing Snapshot copy.

Note: You must make sure that the LUN whose Snapshot copy you are deleting is not being monitored with the Windows Performance Monitor (perfmon).

Syntax for this command is:

```
sdcli snap delete [-m MachineName] -D MountPointList [. . .] -s
SnapshotName
```

Example

```
sdcli snap delete -D k -s Jun_14_09
```

The preceding example deletes the Snapshot copy named Jun_14_09 that is associated with the LUN mapped to K: on the local host.

```
sdcli snap delete -d \\?\Volume{239889f5-3a36-4993-b957-0a85f56cab45}\
-s csvLun11_051109_new
```

The preceding example deletes the Snapshot copy named csvLun11_051109_new from the CSV volume $\?\$ volume $\?\$ volume $\$

The snap list command

The snap list command lists all the Snapshot copies that exist for the specified LUN. Syntax for this command is:

sdcli snap list [-m MachineName] -d MountPoint

MountPoint is a drive letter, mount point path, volume name, or CSV reparse point for the LUN.

Example

```
sdcli snap list -d j
```

The preceding example displays all the Snapshot copies that exist for the volume containing the LUN mapped to "J." on the local host.

The snap mirror_list command

The snap mirror_list command displays the SnapMirror relationships associated with the SnapMirror source volume, including the SnapMirror source storage system, volume, and Snapshot copy; the destination storage system and volume; the current state of the SnapMirror relationship; and whether FlexClone volumes can be created on the destination to allow mirror verification without breaking the mirror.

Syntax for this command is:

sdcli snap mirror_list -d MountPoint

MountPoint is a drive letter, mount point path, volume name, or CSV reparse point of the SnapMirror source drive.

Example

C:\Program Files\IBM\SnapDrive>sdcli snap mirror_list -d e

1 SnapMirror destination(s) Source: andes-1:s Snapshot: andes-2(0084186538)_d.27 Destination: andes-2:d Snapmirrored [FlexClone Success]

The operation completed successfully.

The preceding example displays the SnapMirror relationship between source drive E: on the storage system volume andes-1 and the destination storage system volume andes-2. The source has one Snapshot copy named andes-2(0084186538)_d.27. FlexClone volumes are enabled on the destination.

The snap mount command

The snap mount command mounts a Snapshot copy of a LUN. Snapshot copies are always mounted in read/write mode.

Syntax for this command is:

```
sdcli snap mount [-m MachineName] [-r LiveMachineName] -k LiveMountPoint -s
SnapshotName -d MountPoint
```

LiveMachineName refers to the name of the host connected to the LUN in the active file system. When left unspecified, -r defaults to the local host.

Note: When using this option to mount a Snapshot copy on a remote host, both the local and remote hosts must be running the same version of SnapDrive.

LiveMountPoint refers to the drive letter, mount point, volume name, or CSV reparse point assigned to the LUN in the active file system.

Example

```
sdcli snap mount -r host3 -k j -s Jun_14_09 -d t
```

The preceding example maps the Snapshot copy named Jun_14_09 to drive T: on the local host. This Snapshot copy represents a point-in-time image of the LUN mapped to J: on host3.

The snap rename command

The snap rename command enables you to change the name of an existing Snapshot copy. Syntax for this command is:

sdcli snap rename [-m MachineName] -d MountPoint -o OldSnapshotName -n NewSnapshotName

MountPoint is a drive letter, mount point path, volume name, or CSV reparse point.

Example

sdcli snap rename -d j -o Jun_14_09 -n last_known_good

The preceding example changes the name of the June_14_09 Snapshot copy associated with the J: drive to last known good.

```
sdcli snap rename -d C:\ClusterStorage\Volume8 -o csvLun11_051109 -n
csvLun11_051109_new
```

The preceding example changes the name of the csvLun11_051109 Snapshot copy associated with CSV reparse point C:\ClusterStorage\Volume8 to csvLun11_051109_new.

The snap restore command

The snap restore command replaces the current LUN image in the active file system with the point-in-time image captured by the specified Snapshot copy.

Note: You must make sure that the LUN you are disconnecting is not being monitored with the Windows Performance Monitor (perfmon).

Syntax for this command is:

```
sdcli snap restore [-m MachineName] -d MountPoint -flr [-copy] -s
SnapshotName -files filepath[...]
```

MountPoint is a drive letter, mount point path, volume name, or CSV reparse point.

-flr specifies that you are performing a file level restore.

-copy is an optional setting that forces a copy-restore operation if other restore options fail.

-files enables you to specify the path to the files you want to restore.

filepath is the list of files you want to restore, including the path and drive letter or volume mount point to those files.

Use a hyphen (-) at the end of the file path when you want to restore an entire directory, as well as all subdirectories within that directory, recursively.

Use an asterisk (*) at the end of the specified file when you want to restore all files in that directory only. Subdirectories in that directory will not be recursively restored.

Note: File-level wildcard entries are not supported; therefore, the hyphen and asterisk cannot be combined with a filename at the end of the path you want to restore.

Examples

sdcli snap restore -d l -s Jun_14_09

The preceding example restores the LUN mapped to L: on the local host to its state when the Snapshot copy named Jun_14_09 was taken.

sdcli snap restore -flr -s st_10_18_2009 -files c:\mnt1\v1\v1_d.vhd f: \vm1\vm1.vhd "f:\vm1\Virtual Machines\-"

The preceding example restores a single file called v1_d.vhd in the volume mount point c: mnt1, a single file called vm1.vhd from the volume mount point f:vm1, and restores recursively the folder "Virtual Machines" from the Snapshot copy st_10_18_2009.

The snap unmount command

The snap unmount command disconnects a Snapshot copy of a LUN that is mounted as a LUN.

Note: You must make sure that the LUN whose Snapshot copy you are disconnecting is not being monitored with the Windows Performance Monitor (perfmon).

Attention: If you unmount a LUN on a FlexClone volume that SnapDrive for Windows created and it is the last LUN connected on the volume, SnapDrive deletes that volume resulting in the deletion of all LUNs in the FlexClone volume.

To avoid automatic deletion of the FlexClone volume, rename the volume before unmounting the last LUN. When you rename the volume, be sure to change more than just the last integers in the name. For instance, if the FlexClone volume is named sdw_cl_myvol_0, rename it to new_sdwvol_0, and not to sdw_cl_myvol_20. If you rename only the last integers in the volume name, SnapDrive still recognizes that it created that volume and it will delete the volume when you disconnect the last LUN.

Syntax for this command is:

sdcli snap unmount [-m MachineName] -d MountPoint [-f]

Attention: The -f argument forcibly unmounts the LUN, even if it is in use by an application or Windows. Such a forced operation could cause data loss, so use it with extreme caution.

Example

sdcli snap unmount -d k

The preceding example disconnects the Snapshot copy mapped to K: on the local host.

sdcli snap unmount -d k -f

The preceding example forces disconnection of the Snapshot copy mapped to the K: drive on the local host.

The snap update_mirror command

The snap update_mirror command updates the LUN to a SnapMirror destination volume residing on the same or a different storage system. Syntax for this command is:

sdcli snap update_mirror [-m MachineName] -d MountPoint

Example

sdcli snap update_mirror -d l

The preceding example updates the SnapMirror destination for the LUN mapped to the L: drive on the local host. You do not need to specify the location of the SnapMirror destination because that information was entered when mirroring was set up for the LUN.

The snap restore_volume_check command

The snap restore_volume_check command verifies whether a restore operation can be performed on a volume.

Syntax for this command is:

sdcli snap restore_volume_check [-f StorageSystemName] -volume StorageSystemVolumeName -s SnapshotCopyName [-m MachineName]

-f StorageSystemName is the name of the storage system on which the volume resides.

-volume *StorageSystemVolumeName* indicates the name of the volume on which the restore operation will be performed.

-s *SnapshotCopyName* indicates the name of the Snapshot copy from which the volume will be restored.

Example

sdcli snap restore_volume_check -f clpubs-storage1 -volume vol3 -s
my_snap

The preceding example checks whether a volume restoration from the Snapshot copy named my_snap can be performed on a volume called vol3 that resides on a storage system called clpubs-storage1.

The snap restore_volume command

The snap restore_volume command restores a storage system volume from the specified Snapshot copy.

Syntax for this command is:

```
sdcli snap restore_volume [-f StorageSystemName] -volume
StorageSystemVolumeName -s SnapshotCopyName [-force] [-m MachineName]
```

-f StorageSystemName is the name of the storage system on which the volume resides.

-volume *StorageSystemVolumeName* indicates name of the volume on which the restore operation will be performed.

-s *SnapshotCopyName* indicates the name of the Snapshot copy from which the volume will be restored.

-force is an optional switch that you use to ensure the volume restoration is performed even when non-LUN files or newer Snapshot copies are found on the volume.

Example

```
sdcli snap restore_volume -f clpubs-storage1 -volume vol3 -s my_snap
```

The preceding example restores a volume from the Snapshot copy named my_snap on a volume called vol3 that resides on a storage system called clpubs-storage1

SnapVault commands

The sdcli utility provides command-line support for SnapVault management using SnapDrive.

The snapvault verify_configuration command

The snapvault verify_configuration command enables you to check the SnapVault configuration to ensure that it configured correctly. Syntax for this command is:

sdcli snapvault verify_configuration [-m MachineName] {-D MountPoint | -G guidlist}

-m specifies the name of the remote system on which you want to execute the command. If no machine name is specified, the command is executed on the local system.

-D specifies a list of mount points of disks on the primary system.

-G specifies a list of GUIDs of disks on the primary system.

The snapvault snapshot_rename command

The snapvault snapshot_rename command enables you to rename an existing Snapshot copy on a secondary system

Syntax for this command is:

```
sdcli snapvault snapshot_rename [-m MachineName] -o OldName -n NewName {-d
MountPoint | -G guidlist}
```

 $-\circ OldName$ specifies the name of the existing Snapshot copy that you want to change on the secondary system.

-n NewName specifies the new name of the Snapshot copy on the secondary system. The new name must not yet exist.

-d *MountPoint* specifies the mount point that identifies the disk on the primary system.

-G guidlist specifies a list of GUIDs of disks on the primary system.

The snapvault snapshot_delete command

The snapvault snapshot_delete command deletes an existing Snapshot copy on a SnapVault secondary system.

Syntax for this command is:

```
sdcli snapvault snapshot_delete [-m MachineName] {-D MountPoint | -G
guidlist} -a ArchivalSnapshotName
```

-D MountPoint specifies a list of mount points of disks on the primary system.

-G guidlist specifies a list of GUIDs of disks on the primary system.

-a ArchivalSnapshotName specifies the name of the Snapshot copy that you want to delete.

The snapvault archive command

The snapvault archive command archives a backup set to a secondary system. Syntax for this command is:

```
sdcli snapvault archive [-m MachineName] [-force] -a ArchivalSnapshotName -
DS MountPointandSnapshotList [...]
```

-force forces the secondary Snapshot copy to be made, regardless of the possible failure of some qtree updates.

-a ArchivalSnapshotName specifies the name of the Snapshot copy on the secondary system. This Snapshot copy name must not already exist.

-DS *MountPointandSnapshotList* specifies a list of mount points and Snapshot copies to be archived.

The snapvault relationship_status command

The snapvault relationship_status command displays the relationship status of the primary system for the disk specified.

Syntax for this command is:

sdcli snapvault relationship_status [-m MachineName] {-D MountPoint | -G guidlist}

-D MountPoint specifies a list of mount points of disks on the primary system.

-G guidlist specifies a list of GUIDs of disks on the primary system.

The snapvault snap_list command

The snapvault snap_list command displays the Snapshot copies on the volume specified by the mount point or GUID on the SnapVault secondary system. Syntax for this command is:

sdcli snapvault snap_list [-m MachineName] {-D MountPoint | -G guidlist}

-D MountPoint specifies a list of mount points of disks on the primary system.

-G guidlist specifies a list of GUIDs of disks on the primary system.

OnCommand commands

The sdcli utility provides command-line support for managing OnCommand Manager credentials after SnapDrive has been installed.

The oncommand_config list command

The oncommand_config list command displays a list of already configured OnCommand servers.

Syntax for this command is:

sdcli oncommand_config list

The oncommand_config set command

The oncommand_config set command enables you to set OnCommand server credentials. Syntax for this command is:

sdcli oncommand_config set

-host Host

-user Username

-pwd Password

[-port Port]

-host specifies the hostname or IP address of the host running a DataFabric Manager server.

-user specifies the username for the DataFabric Manager server.

-pwd specifies the password to be used for the DataFabric Manager server.

-port specifies a new TCP port. The default port is 8088 if a new port is not specified.

The oncommand_config delete command

The oncommand_config delete command enables you to remove a Unified Manager server from the SnapDrive Unified Manager server list. Syntax for this command is:

sdcli oncommand config delete

```
-host Host
```

-host specifies the hostname or IP address of the host running the Unified Manager server you want to remove from the list.

The oncommand_config rbaccache command

The oncommand_config rbaccache command enables SnapDrive to cache RBAC operations for use when the OnCommand server is down and unavailable for less than 24 hours. Syntax for this command is:

```
sdcli oncommand_config rbaccache
```

[-m MachineName]

```
-rc Enable | Disable
```

Use -rc Enable or -rc Disable to either enable or disable OnCommand RBAC caching.

Note: When dfm_config rbaccache is enabled, updates to the RBAC cache occur automatically whenever RBAC information changes on the storage system. If the OnCommand is unavailable for more than 24 hours, SnapDrive no longer operates.

Transport protocol commands

The SDCL utility provides command-line support for managing transport protocols used by SnapDrive.

The transport_protocol list command

The transport protocol list command displays the transport protocol configuration settings SnapDrive uses on the storage system. Syntax for this command is:

sdcli transport_protocol list [-m MachineName]

Example

```
sdcli transport_protocol list
Default protocol: HTTP
User Name: root
Port: 80
```

```
Storage System: Storage1
Other IP address(es)/Name: 172.17.176.44
Protocol: HTTP
Username: root
Port: 80
The operation completed successfully.
```

The preceding example indicates that HTTP is the default transport protocol setting on the local SnapDrive system.

The transport_protocol set command

The transport_protocol set command sets or modifies the transport protocol on the storage system.

Syntax

```
sdcli transport_protocol set
[-m MachineName]
-f StorageSystem | -default
-type HTTP| HTTPS |RPC [ -port port]
[ -user UserName]
[-pwd password]
```

Parameters

-f *StorageSystem* specifies the storage system name or IP address.

-default default protocol.

-type specifies the protocol type that will be used. Protocol type is either HTTP, HTTPS, or RPC.

-port *port* specifies the port number the protocol will use. The default port for HTTPS is 443. The default port for HTTP is 80.

-user *UserName* specifies the user with permission on the storage system. A username is required if protocol the type is HTTP/HTTPS.

-pwd password is the password for the user. A password is required if the protocol type is HTTP or HTTPS. You are prompted for a password if it is not specified.

Note:

- When you use clustered Data ONTAP 8.1 or later with SnapDrive, specify the virtual storage server data LIF when you set the transport protocol. The management firewall policy, data role, and protocols must be set to none in the virtual storage server data LIF.
- You must add cluster credentials in the transport protocol setting to enable a cluster-wide user to query any license.

Example

The following example sets the transport protocol to HTTPS on the storage system called atlas-1 using the user name "admin". A password is required but was not specified, so the command prompts the user to enter a password.

```
C:\sdcli transport_protocol set -f atlas-1 -type HTTPS -user admin
Type password for the user:
New transport protocol has been set.
```

The transport_protocol delete command

The transport_protocol delete command deletes a transport protocol from a storage system. Syntax for this command is:

```
sdcli transport_protocol delete [-m MachineName] -f StorageSystem | -
default
```

-f StorageSystem specifies the storage system name or IP address.

-default indicates that the protocol will be the default on the specified storage system.

Virtual server commands

The SDCLI utility provides command-line support for managing virtual server configurations on a VMware Guest OS with SnapDrive either on an ESX server or the vCenter Server.

The vsconfig list command

The vsconfig list command displays the virtual server configuration settings. Syntax for this command is:

sdcli vsconfig list

The vsconfig set command

The vsconfig set command enables you to set the virtual server configuration. Syntax for this command is:

sdcli vsconfig set -ip IP Address -user User Name -pwd Password

-ip specifies the IP address of the virtual server.

-user specifies the virtual machine user name.

-pwd specifies the password for the virtual machine.

Note: You can change the IP address of a virtual server that you have previously configured during SnapDrive installation without re-entering the username and password. A valid IP address,

username, and password must already exist in the Windows registry; otherwise, an error message is displayed indicating that the user credentials are not set.

The vsconfig dslist command

The vsconfig dslist command displays datastores available on ESX servers. Syntax for this command is:

sdcli vsconfig dslist -m MachineName -port "Port Number"

-port specifies the Web service port number you use to communicate with SnapDrive. The default port is 808.

The vsconfig delete command

The vsconfig delete command disables ESX server and vCenter Server settings. Syntax for this command is:

sdcli vsconfig delete

Note: When you disable vCenter Server or ESX settings, SnapDrive cannot display WWPNs for FC HBAs on the ESX server.

Note: You cannot disable vCenter Server or ESX settings when FC RDM LUNs are present.

Hyper-V configuration commands

The SnapDrive for Windows GUI enables you to add only one Hyper-V node when you configure pass-through disk provisioning. You can use the hyperv_config command group in SnapDrive for Windows to specify additional Hyper-V nodes, as well as delete and list Hyper-V nodes; for instance, when you have a Cluster Shared Volume on a Microsoft cluster.

The hyperv_config list command

The hyperv_config list command enables you to list all the configured Hyper-V parent nodes. You can use this command to verify that SnapDrive recognizes all of the existing Hyper-V nodes, for example, if they belong to Cluster Shared Volumes. Syntax for this command is as follows:

Syntax for this command is as follow

sdcli hyperv_config list

The hyperv_config set command

You can specify additional Hyper-V nodes in Cluster Shared Volumes on a Microsoft cluster using the hyperv_config set command. This is useful, for instance, when SnapDrive must locate a LUN on a virtual machine that has been moved to another node in the cluster. Syntax for this command is as follows:

sdcli hyperv_config set -host hostname -IP IP_address [-port port_number]

-host hostname specifies the Hyper-V parent node host name.

-IP IP_address specifies the parent node IP address.

-port *port_number* specifies the parent node SnapDrive Web service TCP port number. The default is 808.

Example

sdcli hyperv_config set -host NN-HYP-001-P -IP 10.20.1.150 -port 808

The hyperv_config delete command

The hyperv_config delete command enables you to delete a Hyper-V parent node configuration. This is helpful, for instance, when the Hyper-V configuration is no longer required. Syntax for this command is as follows:

sdcli hyperv_config delete -host hostname

Typical SnapDrive configurations

SnapDrive for Windows supports a variety of configurations for your iSCSI, FC, or MPIO environment.

SnapDrive iSCSI configurations

SnapDrive for Windows supports several different iSCSI configurations.

Single host direct-attached to a single storage system using iSCSI

You can configure SnapDrive for Windows to use a GbE crossover cable to attach the host directly to the storage system, an arrangement that minimizes latency and eliminates unwanted network broadcasts.

The host and storage system in this configuration each use the following connection hardware:

- 1 GbE NIC dedicated to host-storage system data transfer
- 1 Fast Ethernet (or GbE) NIC to connect to the data-center fabric

Note: Both the storage system and the host must be within the same broadcast domain.

Note: LUN traffic and management traffic in an iSCSI configuration can be performed over a single GbE connection; however, for best results, you should separate the traffic as shown in the following illustration.



Single host attached to a single storage system through a GbE switch

You can configure SnapDrive for Windows to use a single-homed configuration that places a network switch between the storage system and the host, an arrangement that provides good

172 | SnapDrive 7.0 for Windows Administration Guide for SAN Environments

performance and also segregates host-storage system traffic by directing it through a single pair of switch ports.

Because the switch connects to the data-center fabric, the host and storage system in this configuration each use a single GbE NIC both for host-storage system data transfers and for connecting to the data-center fabric.

Note: LUN traffic and management traffic in an iSCSI configuration can be performed over a single GbE connection; however, for best results, you should separate the traffic as shown in the following illustration.



Single host attached to a single storage system through a dedicated switch

You can configure SnapDrive for Windows to use a multihomed configuration with a GbE switch between the storage system and the host, an arrangement that, in addition to providing good performance and segregating host-storage system traffic to the dedicated switch, also minimizes disruptions in situations where network routing configuration changes frequently.

The host and storage system in this configuration each use the following hardware for the connection:

- 1 GbE NIC dedicated to data transfer between host and storage system
- 1 Fast Ethernet (or GbE) NIC to connect to the data-center fabric

Note: LUN traffic and management traffic in an iSCSI configuration can be performed over a single GbE connection; however, for best results, you should separate the traffic as shown in the following illustration.



Windows cluster connected to a storage system cluster through a dedicated GbE switch

You can configure SnapDrive for Windows to use both a Windows cluster and a storage system cluster.

The following illustration shows a Windows cluster and storage system cluster with an optional but recommended "private" network that manages internal cluster traffic (rather than data traffic between host and storage system).

You can also create configurations that connect the host cluster to multiple storage systems or storage system active/active configurations, and you can connect a storage system or storage system active/ active configuration to multiple hosts.

Note: LUN traffic and management traffic in an iSCSI configuration can be performed over a single GbE connection; however, for best results, you should separate the traffic as shown in the following illustration.



SnapDrive FC configurations

SnapDrive for Windows supports several different FC configurations.

Single host direct-attached to a single storage system using FC

You can configure SnapDrive for Windows to use a crossover FC cable to attach the host directly to the storage system.

The host and storage system in this configuration each use the following connection hardware:

- 1 HBA to transfer LUN data between storage system and host
- 1 Fast Ethernet (or GbE) NIC to connect to the data-center fabric

Note: Both the storage system and the host must be within the same broadcast domain.



Figure 1: Single host direct-attached to single storage system using FC

Single host attached to a single storage system through an FC switch

You can configure SnapDrive for Windows to use a dedicated FC switch between the storage system and the host.

The host and storage system in this configuration each use the following connection hardware:

- 1 HBA to transfer LUN data between storage system and host
- 1 Fast Ethernet (or GbE) NIC to connect to the data-center fabric

Note: LUN traffic and management traffic in an FC configuration can be performed over a single GbE connection, however, for best results, you should separate the traffic as shown in the following illustration.



Windows cluster attached to a storage system active/active configuration through an FC switch

You can configure SnapDrive for Windows to use both a Windows cluster and a storage system active/active configuration connected through an FC switch.

The following illustration shows a Windows cluster and a storage system active/active configuration with an optional but recommended dedicated network for internal cluster traffic.

You can also create configurations that connect the Windows cluster to multiple storage systems or storage system active/active configurations.



SnapDrive MPIO configurations

SnapDrive for Windows supports several different MPIO configurations.

If you plan to use MPIO configurations with SnapDrive, you should download Data ONTAP DSM for Windows MPIO from the N series support website (accessed and navigated as described in *Websites* on page 12). MPIO is not included with the SnapDrive installation. For more information, see the *Data ONTAP DSM for Windows MPIO Installation and Administration Guide*.

For more information about the latest supported MPIO configurations, see the N series support website (accessed and navigated as described in *Websites* on page 12).

Single host direct-attached to a single storage system using MPIO

You can configure SnapDrive for Windows to employ FC or iSCSI HBAs to support MPIO between a host and a single direct-attached storage system.

Using FC HBAs, the host and storage system in this configuration each use the following connection hardware:

- 2 FC HBAs to transfer multipathed LUN data between storage system and host
- 1 Fast Ethernet (or GbE) NIC to connect to the data-center fabric

Using iSCSI HBAs or the Microsoft iSCSI Software Initiator, the storage system in this configuration has two GbE adapters, and the host has one or both of the following:

- 2 or more iSCSI HBAs
- The Microsoft iSCSI Software Initiator and 2 GbE NICs

176 | SnapDrive 7.0 for Windows Administration Guide for SAN Environments



Windows cluster attached to a storage system active/active configuration through a GbE switch using MPIO

You can configure SnapDrive for Windows to employ both a Windows cluster and a storage system active/active configuration connected through a GbE switch using MPIO.

The following illustration shows a Windows cluster and a storage system active/active configuration with an optional but recommended dedicated network for internal cluster traffic.

Each host in this configuration uses the following connection hardware:

- 2 GbE (or iSCSI HBAs) to transfer multipathed LUN data between storage system and host
- 1 Fast Ethernet (or GbE) NIC to connect to the data-center fabric
- 1 optional Fast Ethernet, GbE, or 10/100 NIC for internal cluster traffic

Each storage system in this configuration requires at least two Fast Ethernet (or GbE) NICs to connect to the data-center fabric. (See your *Data ONTAP SAN Administration Guide for 7-Mode* for details.)



Windows cluster attached to a storage system active/active configuration through an FC switch using MPIO

You can configure SnapDrive for Windows to use both a Windows cluster and a storage system active/active configuration connected through an FC switch using MPIO.

The following illustration shows a Windows cluster and a storage system active/active configuration with an optional but recommended dedicated network for internal cluster traffic.

Each host in this configuration uses the following connection hardware:

- 2 HBAs, to transfer multipathed LUN data between storage system and host
- 1 Fast Ethernet (or GbE) NIC, to connect to the data-center fabric
- 1 optional Fast Ethernet, GbE, or 10/100 NIC, for internal cluster traffic

Each storage system configuration requires two dual-port FC adapters and a Fast Ethernet (or GbE) NIC to connect to the data-center fabric. (See your *Data ONTAP SAN Administration Guide for 7-Mode* for details.)



SAN booting with SnapDrive

SnapDrive for Windows supports SAN booting, but some tasks are restricted.

What SAN booting is

SAN booting is the general term for booting a host from a storage system LUN instead of an internal hard disk.

Fibre Channel SAN booting does not require support for special SCSI operations; it is not different from any other SCSI disk operation. The HBA uses special code in the BIOS that enables the host to boot from a LUN on the storage system.

iSCSI SAN booting also uses special code in the BIOS that enables the host to boot from a LUN on the storage system. You need to set specific parameters in the BIOS to enable iSCSI SAN booting.

The general process is as follows.

- 1. After the HBA has accessed the BIOS, use the Emulex or QLogic BootBIOS utility to configure the LUN as a boot device.
- 2. Configure the PC BIOS to make the LUN the first disk device in the boot order.
- **3.** Install the following on the LUN.
 - · Windows operating system
 - HBA driver

Note: Following a system failure, the bootable virtual disk is no longer the default boot device. You need to reconfigure the hard disk sequence in the system BIOS to set the bootable virtual disk as the default boot device.

How SnapDrive supports SAN booting

SnapDrive for Windows identifies bootable LUNs and prevents you from performing some of the operations you would normally perform on a nonbootable LUN.

SnapDrive detects both bootable LUNs (SAN booting) and nonbootable LUNs and differentiates between the two in MMC by representing each LUN type with a unique icon. SAN bootable LUNs are represented by an icon containing a disk with a red letter "s" in the upper left corner.

When a LUN is a boot disk, the following actions are disabled or unavailable in SnapDrive:

- Disconnect
- Delete
- Expand
- Restore

SnapDrive does support the following Snapshot copy-related actions on bootable LUNs:

- Create
- Rename
- Delete

Note: Restoring Snapshot copies of bootable LUNs is not allowed by SnapDrive. For important information about Snapshot copies of bootable LUNs, see the technical white papers on the N series support website (accessed and navigated as described in *Websites* on page 12).

Copyright and trademark information

Copyright ©1994 - 2013 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2013 IBM Corporation. All rights reserved.

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

No part of this document covered by copyright may be reproduced in any form or by any means— graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

No part of this document covered by copyright may be reproduced in any form or by any means— graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service
Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing IBM Corporation North Castle Drive Armonk, N.Y. 10504-1785 U.S.A.

For additional information, visit the web at: http://www.ibm.com/ibm/licensing/contact/

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Index

7-Mode limitations SnapVault *90*7-Mode SAN environment initiating SnapVault backup jobs in *91*

A

access control operations list and descriptions *104* AccessControl.xml file operations *104* using to manage storage system access control *103* adding a backup using dataset backup_add *121* adding metadata to a dataset using the dataset set_metadata command *135* adding objects to a dataset using dataset add_members *120* application transaction log changes introduced in Windows Server 2012 *30* assigning roles using RBAC *113*

B

backing up a dataset using the dataset backup start command 125 backup applications using VSS with 19 backup copy location of datasets 129 backup jobs using SnapVault with SnapDrive in 7-Mode SAN 91 backup retention policy information obtaining using dataset get retention info 130 backup sets SnapVault 90 backup version numbers converting to timestamps 126 BackupComplete call changes introduced in Windows Server 2012 30 backups excluding independent VMDK disks 87 typical process using VSS 20

С

changing dataset backup metadata 124 CHAP authentication about using with SnapDrive 26 cluster configurations support of SnapDrive for Windows 32 cluster service 99 Cluster Shared Volume See CSV Cluster Shared Volume File System (CSVFS) 30 Cluster Shared Volumes about 32cluster support of SnapDrive for Windows 32 clustered Data ONTAP creating disks in 45 SnapVault operations supported in 91 clusters 83 command-line interface spacereclaimer 144 commands dataset 120 dataset add members 120 dataset backup add 121 dataset backup change retention type 121 dataset backup delete 122 dataset backup end 123 dataset backup get metadata 123 dataset backup list 124 dataset backup set metadata 124 dataset backup start 125 dataset backup status 125 dataset backup version convert 126 dataset create 126 dataset create local backup 127 dataset delete 127 dataset dfm request 128 dataset get available policies 128 dataset get backup location 129 dataset get backup version info 129 dataset get metadata 130 dataset get policy 130 dataset get retention info 130 dataset info 131 dataset initiate conformance 131

dataset list members 132 dataset mount backup 132 dataset protect 133 dataset remove members 133 dataset restore 134 dataset restore status 135 dataset set metadata 135 dataset set policy 136 dataset transfer now 136 dataset vss backup end 137 dataset vss backup prepare 137 sending a request to Protection Manager 128 components SnapDrive, described 18 configuration of a failover cluster witness disk 42 verifying for VSS 22 configuration requirements for SnapVault support 91 SnapVault 90 configurations FC, single host, dedicated switch 174 FC, single host, direct-attached 173 FC, Windows cluster 174 iSCSI, single host, dedicated switch 172 iSCSI, single host, direct-attached 171 iSCSI, single host, GbE switch 171 iSCSI, Windows cluster 173 MPIO, single host, direct-attached 175 MPIO, Windows cluster, FC switch 177 MPIO, Windows cluster, GbE switch 176 configuring RBAC on SnapDrive 112 space reservation monitoring 70 conformance check initiating using dataset initiate conformance 131 Connect Disk wizard using 48 connecting LUNs, list of guidelines 47 connections to a disk 48 to a LUN 48 copy management provided to a guest OS by ESX iSCSI initiators 28 creating LUNs 35 RDM LUNs 78 creating a local dataset backup using dataset create local backup 127

creating a new dataset using the dataset create command 126 creating disks on a virtual storage server 45 creating roles using RBAC 112 CSV 2.0 changes to architecture 30 CSVs recommendation when using SnapDrive 33 requirements for running Space Reclaimer on 75

D

Data ONTAP DSM using with SnapDrive 31 data protection capabilities of OnCommand Unified Manager Core Package 101 Data protection capabilities integration with SnapDrive 101 data servers using VSS with 19 dataset creating a local backup 127 transferring a 136 viewing members of 132 dataset commands usage 120 dataset operations using dataset commands 120 datasets adding objects to 120 changing the retention type 121 creating a new 126 deleting 127 general concepts 101 getting a backup copy location 129 getting available retention policies 128 getting backup metadata 123 getting backup retention policy information 130 getting information about 131 getting information about the storage policy 130 getting metadata for 130 getting your backup version number 129 initiating a conformance check 131 listing backup copies 124 mounting a backup of 132 preparing for a VSS backup 137 removing objects 133 restoring from a backup 134 restoring the status of 135

185 | SnapDrive 7.0 for Windows Administration Guide for SAN Environments

scheduling a backup 121 securing 133 setting a storage policy 136 setting backup metadata 124 starting a backup 125 stopping a VSS backup 137 viewing a backup status 125 dedicated LUN creating 35 Delete Disk 53 deleting a LUN 53 folder within volume mount point 54 volume mount point 50 deleting a dataset using the dataset delete command 127 deleting a dataset backup using dataset backup delete 122 deleting a remote backup using dataset backup delete 122 disconnecting a LUN 52 forced (of LUN) 52 iSCSI session 27 iSCSI target 26 disks creating in clustered Data ONTAP 45 creating on a virtual storage server 45 expanding 54 resizing guidelines 54 shrinking 54 documentation Microsoft cluster with ESX 89 reference material Microsoft cluster with ESX 89

E

enabling ESX or vCenter logon from SnapDrive MMC 76
ending dataset backup operations

using dataset backup_end 123

ESX iSCSI initiator support

limitations 28

ESX iSCSI initiators

supported by SnapDrive for Windows 28

ESX logon

enabling and disabling from SnapDrive MMC 76

ESX server

limitations to SnapDrive support 76

examples storage system access control commands 107 expanding a LUN 55 a quorum disk 56 expanding LUNs about 54

F

failover cluster witness disks configuring 42 FC configurations 173 FC RDM LUN support with Microsoft clusters 82 features list Windows Server 2012 30 Fibre Channel sessions managing 24 file-level restore operation described 68 support for 68 FlexClone volumes about using in SnapDrive 63 prerequisites for using with SnapDrive 63 forced disconnect (of LUN) 52 fractional space reservation monitoring about 70

G

getting backup 123 getting dataset information using the dataset info command 131 getting retention policy information using dataset get_retention_info 130 GPT partition support 45 guest OS providing LUN provisioning and copy management to 28

Η

high availability providing 31 virtual machine 44 hot add and remove feature requirements for using 58 Hyper-V enhancements with Windows Server 2012 *30*Hyper-V pass-through disk support list of limitations *59*Hyper-V pass-through disks requirements for dynamically adding and removing *58*Hyper-V replica *30*Hyper-V Server role CSV creation *33*Hyper-V VSS backup *30*

I

initiating a conformance check using dataset initiate_conformance 131
initiators adding to LUNs using SDCLI 154
iSCSI configurations 171 disconnecting target from Windows host 26 establishing a session to a target 24 examining session details 27
iSCSI session disconnecting from target 27
iSCSI sessions managing 24
iSCSI Software Initiator node naming standards 24

L

limitations SnapDrive 16 SnapVault 90 to SnapDrive support of ESX iSCSI initiators 28 to SnapDrive support of Hyper-V pass-through disks 59 VMDK and Snapshot functionality interoperability in SnapDrive 86 listing backups using dataset backup list 124 listing members of a dataset using dataset list members 132 LUN protocols 47 LUN clone split feature about 66 LUN provisioning provided to a guest OS by ESX iSCSI initiators 28

LUN restore checking status 68 LUNs about disconnecting in FlexClone volume 51 about disconnecting or deleting 51 adding an initiator to using SDCLI 154 adding, removing, or changing a drive letter or path 50 connecting to 48creating 35 creating and using in VMware environments 76 creating and using SnapDrive in Microsoft environments 30 creating in VMware environments 78 creating shared 38 deleting 53 disconnecting 52 expanding 54, 55 expanding a quorum disk 56 forced disconnect 52 how the storage system interacts with 46list of guidelines for connecting 47managing 30, 46, 76 managing in VMware environments 83 managing non-SnapDrive LUNs 56 moving a mount point 51 rules for creating 34LUNS on a SnapMirror destination volume 96

M

managing LUNs 46 managing space 70metadata adding to a dataset 135 obtaining for specified datasets 130 Microsoft cluster with ESX related documentation 89 Microsoft clusters 83 Microsoft MPIO integrating with Data ONTAP DSM for Windows MPIO 31 Microsoft Windows Server 2012 features 30 mirroring operations performing by integrating SnapDrive with SnapMirror 90 monitoring space reservation 70

mounting a dataset backup using dataset mount_backup *132* MPIO configurations *175*

N

naming requirements iSCSI Software Initiator 24 naming settings definition of 101 new features Microsoft Windows Server 2012 30 new quorum disk 99 nodes naming standards for the iSCSI Software Initiator 24 non-SnapDrive LUNs managing 56 preparing for SnapDrive 56

0

operations on datasets 120

P

pass-through disks requirements for dynamically adding and removing 58 pass-through LUNs running Space Reclaimer on 74 path-failover methods using 31 performance optimization 144 planned VHD 30 preparing for a VSS backup using dataset vss backup prepare 137 Protection Manager sending requests to for backup version information 128 protection or provisioning related objects definition of 101 protection policies overview 101 protocols for accessing LUNs 47 provisioning policies overview 101

Q

quorum disk expanding 56

R

rapid LUN restore about 66 RBAC configuring SnapDrive to use 112 enabling on a storage system 111 enabling to configure access control 103 support for 110 See also RBAC RBAC (role-based access control assigning roles 113 creating roles 112 RDM 51 RDM LUN creating 78 FC support 82 shared 83 RDM LUNs troubleshooting creation of 81 **RDM** operations vCenter privileges for 77 removing a dataset backup using dataset backup delete 122 removing objects from a dataset using dataset remove members 133 removing stale RDMs 51 requirements for space reclamation in VMDK files in NFS datastores 88 hot add and remove 58 restoring a dataset status using dataset restore status 135 restoring from a dataset backup using dataset restore 134 retention policy information obtaining using dataset get retention info 130 role mappings SnapDrive to DataFabric Manager server 113 role-based access control See RBAC roles storage system access control 105 rolling Snapshot copies 93

S

SAN booting SnapDrive support for 178 SAN environment initiating SnapVault backup jobs in 91 schedules Snapshot copy 62 scheduling a backup using dataset backup add 121 SDCLI dataset commands 120 sdcli spacereclaimer start 144 setting space reservation 71 setting a dataset storage policy using the dataset set policy command 136 shadow copies 19 shared FC RDM LUN creating 83 shared LUNs creating 38 shared RDM LUN using in a Microsoft cluster 83 SnapDrive capabilities 15 components, described 18 configurations 171 configuring to enable virtual Fibre Channel port connection 46 in SMB 3.0 environments 15 integrating with OnCommand Unified Manager Core Package 101 integration SnapDrive with OnCommand Unified Manager Core Package data protection 101 integration with data protection capabilities 101 integration with Windows Volume Manager 15 not listing all VMDK disks 88 recommendations for using 16 using in Microsoft environments 30 using in VMware environments 76 SnapDrive to DataFabric Manager server role mappings 113 SnapManager for SQL and SnapVault support in cluster environments 91 SnapMirror integrating with SnapDrive to perform mirroring operations 90 requirements for using with SnapDrive 94

Snapshot backup copies managing using SnapDrive in Microsoft environments 30 managing using SnapDrive in VMware environments 76 Snapshot backups managing in VMware environments 86 Snapshot copies backing up to a secondary storage system using SnapVault 90 creating 61 creating in VMDKs 86 deleting 69 deleting in VMDKs 86 reasons for creating 60scheduling 62support with VMDKs on NFS and VMFS datastores 86 Snapshot copy management provided to a guest OS by ESX iSCSI initiators 28 Snapshot copy support interoperability with VMDKs in SnapDrive 86 SnapVault 7-Mode limitations 90 backup sets 90 configuration requirements 90 initiating backup jobs from SnapDrive 7-Mode SAN 91 integrating with SnapDrive to perform vaulting operations 90 limitations 90 system requirements 90 usage described 90 SnapVault support by SnapDrive for Windows in cluster environments 91 space managing 70 space optimization space reclaimer 144 Space Reclaimer about 71 enabling 74 guidelines for using 72reasons for SnapDrive to automatically stop 73 requirements for running on CSVs 75 running on pass-through LUNs 74 starting 72 stopping 73 support with VMDK 88

space reclamation requirements for VMDK files in NFS datastores 88 space reservation monitoring 70 space reservation settings managing with the storage access control tool 71 starting Space Reclaimer 72 starting a backup using the dataset backup start command 125 starting a conformance check using dataset initiate conformance 131 stopping Space Reclaimer 73 stopping a VSS backup using dataset vss backup end 137 storage management using VSS with 19 storage policy information getting using the dataset get policy command 130 storage services overview 101 storage system configuring access control 103 managing access control 103 storage system access allowing using AccessControl.xml operations 104 restricting using AccessControl.xml operations 104 storage system access control command examples 107 described 103 roles, described 105 support for 103 storage system access control tool (storacl.exe) using 103 storage system capabilities adding for HTTP users 29 storage systems access control operations described 104 access control tool 107 controlling user activity on 107 enabling RBAC on 111 establishing a connection to 24 interacting with LUNs 46 managing space on using SnapDrive in Microsoft environments 30 managing space on using SnapDrive in VMware environments 76 surprise removal (of LUN) 52 system requirements SnapVault 90

Т

thinly provisioned LUNs enabling 71 timestamps converting backup version numbers to 126 transferring a dataset using dataset transfer_now 136 transport protocol settings managing 24 troubleshooting RDM LUN creation 81 VMDKs 87, 88 VSS Hardware Provider 21 truncated transaction logs 30

U

Update Mirror operations using the storage system access control tool (storacl.exe) to determine operations enabled for users to determine user roles

V

vaulting operations performing by integrating SnapDrive with SnapVault 90 vCenter logon enabling and disabling from SnapDrive MMC 76 vCenter privileges for performing RDM operations 77 minimum 77 viewing backup copies using dataset backup list 124 Virtual Fibre Channel configuring SnapDrive for 46 using to connect a guest OS to the storage system 46virtual machine creating a highly available 44 Virtual Storage Console 143 virtual storage servers configuration requirements in clustered Data ONTAP 45 VMDK disk partition style 87 inconsistent disk enumeration 87 interoperability with Snapshot functionality and SnapDrive 86 Snapshot copy support 86

Space Reclaimer support 88 troubleshooting 87 VMDKs not listed in SnapDrive enumeration 88 troubleshooting 87, 88 vMotion requirements 77 support 77 VMware support 76 VMware ESX server limitations to SnapDrive support 76 volume mount points about 35 adding 50 changing 50 deleting folder within 54 limitations 35 moving 51 removing 50 volume-based Snapshot copy restoration about <u>68</u> volume-level SnapVault support

by SnapDrive for Windows in cluster environments 91 VSS

about 19 troubleshooting 21 typical backup process 20 verifying configuration 22 verifying provider used 22 viewing installed providers 21 VSS backup 30 VSS backups preparing for 137

W

Windows Server 2008 failover cluster support 42
Windows Server 2012 changes introduced to Hyper-V VSS backups 30 CSV changes 30
Windows Server 2012 features 30
Windows Volume Manager integration with SnapDrive 15

IBM.®

NA 210-06113_A0, Printed in USA

SC27-5983-00

