




## SnapDrive® 7.1 for Windows®

### Administration Guide

#### For SAN Environments





# Contents

|                                                                                         |           |
|-----------------------------------------------------------------------------------------|-----------|
| <b>SnapDrive overview .....</b>                                                         | <b>11</b> |
| What's new in SnapDrive 7.1 for Windows .....                                           | 11        |
| Automated storage provisioning and data management using SnapDrive for<br>Windows ..... | 13        |
| List of current SnapDrive for Windows limitations .....                                 | 14        |
| Understanding SnapDrive for Windows components .....                                    | 14        |
| Understanding the Volume Shadow Copy Service .....                                      | 15        |
| Recommendations for using SnapDrive for Windows .....                                   | 16        |
| <b>Establishing a connection to the storage system .....</b>                            | <b>17</b> |
| Managing iSCSI sessions .....                                                           | 17        |
| iSCSI Software Initiator node naming requirements .....                                 | 17        |
| Establishing an iSCSI session to a target .....                                         | 17        |
| Disconnecting an iSCSI target from a Windows host .....                                 | 19        |
| Disconnecting a session to an iSCSI target .....                                        | 20        |
| Examining details of an iSCSI session .....                                             | 20        |
| SnapDrive for Windows support for ESX iSCSI initiators .....                            | 21        |
| ESX limitations .....                                                                   | 21        |
| Enabling storage system HTTP communication with SnapDrive .....                         | 22        |
| <b>Using SnapDrive in Microsoft environments .....</b>                                  | <b>23</b> |
| Understanding Data ONTAP DSM support for Windows MPIO .....                             | 23        |
| Support for Microsoft Cluster Shared Volumes .....                                      | 23        |
| Verifying the cluster group owner .....                                                 | 24        |
| Changing the cluster group owner .....                                                  | 24        |
| Creating LUNs .....                                                                     | 25        |
| Rules for creating LUNs .....                                                           | 25        |
| What volume mount points are .....                                                      | 26        |
| Creating a dedicated LUN .....                                                          | 26        |
| Creating a shared LUN .....                                                             | 28        |
| Windows Server failover cluster support .....                                           | 29        |
| Configuring a Windows Server failover cluster witness disk .....                        | 30        |
| Support for creating disks on a SVM operating in clustered Data<br>ONTAP .....          | 31        |

|                                                                                                       |           |
|-------------------------------------------------------------------------------------------------------|-----------|
| GPT partition support .....                                                                           | 31        |
| Configuring SnapDrive for Windows to create LUNs with virtual Fibre Channel .....                     | 32        |
| Managing LUNs .....                                                                                   | 32        |
| How LUNs work .....                                                                                   | 32        |
| List of guidelines for connecting LUNs .....                                                          | 33        |
| Connecting to a LUN .....                                                                             | 33        |
| Making drive letter or path modifications to a LUN .....                                              | 34        |
| Guidelines for disconnecting or deleting LUNs .....                                                   | 35        |
| Disconnecting a LUN .....                                                                             | 36        |
| Deleting a LUN .....                                                                                  | 37        |
| Deleting folders within volume mount points .....                                                     | 38        |
| Guidelines for resizing disks .....                                                                   | 38        |
| Resizing a disk .....                                                                                 | 39        |
| Requirements for dynamically adding and removing pass-through disks on Hyper-V virtual machines ..... | 40        |
| Managing Snapshot copies .....                                                                        | 42        |
| Reasons for creating Snapshot copies .....                                                            | 42        |
| Restrictions on Snapshot copy creation .....                                                          | 42        |
| Creating a Snapshot copy .....                                                                        | 43        |
| Scheduling Snapshot copies .....                                                                      | 44        |
| Support for FlexClone volumes .....                                                                   | 45        |
| Restrictions when working with Snapshot copies .....                                                  | 46        |
| Connecting to a LUN in a Snapshot copy .....                                                          | 46        |
| Data protection through archiving and restoring Snapshot backup copies ...                            | 47        |
| Deleting a Snapshot copy .....                                                                        | 50        |
| Managing space on storage system volumes .....                                                        | 52        |
| What SnapDrive fractional space reservation monitoring does .....                                     | 52        |
| Configuring space reservation monitoring .....                                                        | 52        |
| Using the storage access control tool to enable thinly provisioned LUNs ....                          | 53        |
| What Space Reclaimer does .....                                                                       | 54        |
| <b>Using SnapDrive in VMware environments .....</b>                                                   | <b>58</b> |
| Supported VMware guest OS platforms .....                                                             | 58        |
| VMware ESX server-related limitations .....                                                           | 58        |
| Enabling and disabling vCenter or ESX login from SnapDrive for Windows MMC .....                      | 59        |

|                                                                                  |           |
|----------------------------------------------------------------------------------|-----------|
| Minimum vCenter privileges required for SnapDrive operations .....               | 60        |
| Creating LUNs in VMware environments .....                                       | 60        |
| Creating an RDM LUN on a guest OS .....                                          | 60        |
| Using FC RDM LUNs in a Microsoft cluster .....                                   | 63        |
| Managing LUNs in VMware environments .....                                       | 64        |
| Connecting to an RDM LUN on a guest OS .....                                     | 65        |
| Guideline for managing RDM LUNs not created in SnapDrive .....                   | 66        |
| Managing Snapshot backups in VMware environments .....                           | 66        |
| Support requirements for performing Snapshot copy operations in                  |           |
| VMDKs on NFS and VMFS datastores .....                                           | 66        |
| Snapshot copy support limitations on VMDKs .....                                 | 67        |
| Troubleshooting VMDKs .....                                                      | 67        |
| Support requirements for space reclamation in VMDK files in NFS datastores ..... | 69        |
| <b>Performing SnapVault and SnapMirror operations .....</b>                      | <b>70</b> |
| Using SnapVault with SnapDrive .....                                             | 70        |
| Requirements for using SnapVault .....                                           | 70        |
| Initiating SnapVault backup jobs from SnapDrive for Windows .....                | 71        |
| SnapVault operations supported in a clustered Data ONTAP .....                   | 71        |
| Using SnapMirror with SnapDrive for Windows .....                                | 72        |
| SnapMirror overview .....                                                        | 72        |
| Types of SnapMirror replication .....                                            | 72        |
| Requirements for using SnapMirror with SnapDrive for Windows .....               | 74        |
| Initiating replication manually .....                                            | 75        |
| Connecting to a LUN in a mirrored destination volume .....                       | 76        |
| Restoring a volume on a SnapMirror destination .....                             | 77        |
| SnapDrive integration with OnCommand Unified Manager Core Package data           |           |
| protection capabilities .....                                                    | 78        |
| How SnapDrive for Windows integrates with OnCommand Unified                      |           |
| Manager Core Package data protection capabilities .....                          | 78        |
| <b>Configuring and managing access control .....</b>                             | <b>80</b> |
| Support for storage system access control .....                                  | 80        |
| Using storage system access control .....                                        | 80        |
| Storage system access control reference .....                                    | 81        |
| Enabling RBAC for use with SnapDrive .....                                       | 87        |
| Using RBAC with the OnCommand Unified Manager Core Package                       |           |
| server .....                                                                     | 87        |

|                                                                       |           |
|-----------------------------------------------------------------------|-----------|
| Enabling RBAC on the storage system .....                             | 89        |
| Configuring SnapDrive for Windows to use RBAC .....                   | 89        |
| Creating SnapDrive user roles on DataFabric Manager server .....      | 90        |
| Assigning roles to SnapDrive users on DataFabric Manager server ..... | 91        |
| SnapDrive for Windows to DataFabric Manager role mappings .....       | 91        |
| <b>SnapDrive command-line reference .....</b>                         | <b>95</b> |
| About sdcli commands .....                                            | 95        |
| Executing sdcli commands .....                                        | 95        |
| Common command switches .....                                         | 96        |
| Configuration commands .....                                          | 98        |
| sysconfig list .....                                                  | 98        |
| Dataset management commands .....                                     | 98        |
| dataset add_members .....                                             | 98        |
| dataset backup_add .....                                              | 99        |
| dataset backup_change_retention_type .....                            | 100       |
| dataset backup_delete .....                                           | 100       |
| dataset backup_end .....                                              | 101       |
| dataset backup_get_metadata .....                                     | 102       |
| dataset backup_list .....                                             | 102       |
| dataset backup_set_metadata .....                                     | 103       |
| dataset backup_start .....                                            | 104       |
| dataset backup_status .....                                           | 104       |
| dataset backup_version_convert .....                                  | 105       |
| dataset create .....                                                  | 105       |
| dataset create_local_backup .....                                     | 106       |
| dataset delete .....                                                  | 106       |
| dataset dfm_request .....                                             | 107       |
| dataset get_available_policies .....                                  | 107       |
| dataset get_backup_location .....                                     | 108       |
| dataset get_backup_version_info .....                                 | 108       |
| dataset get_metadata .....                                            | 109       |
| dataset get_policy .....                                              | 109       |
| dataset get_retention_info .....                                      | 110       |
| dataset info .....                                                    | 110       |
| dataset initiate_conformance .....                                    | 111       |
| dataset list_members .....                                            | 111       |

|                                                        |     |
|--------------------------------------------------------|-----|
| dataset mount_backup .....                             | 112 |
| dataset protect .....                                  | 113 |
| dataset remove_members .....                           | 113 |
| dataset restore .....                                  | 114 |
| dataset restore_status .....                           | 114 |
| dataset set_metadata .....                             | 115 |
| dataset set_policy .....                               | 116 |
| dataset transfer_now .....                             | 116 |
| dataset vss_backup_end .....                           | 117 |
| dataset vss_backup_prepare .....                       | 117 |
| License commands .....                                 | 118 |
| license set .....                                      | 118 |
| license list .....                                     | 119 |
| license remove .....                                   | 119 |
| Initiator group management commands .....              | 120 |
| igroup list .....                                      | 120 |
| igroup create .....                                    | 121 |
| igroup rename .....                                    | 122 |
| igroup modify .....                                    | 123 |
| igroup delete .....                                    | 123 |
| Fractional space reservation monitoring commands ..... | 124 |
| spacemon list .....                                    | 124 |
| spacemon set .....                                     | 125 |
| spacemon snap_delta .....                              | 126 |
| spacemon snap_reclaimable .....                        | 127 |
| spacemon vol_info .....                                | 127 |
| spacemon delete .....                                  | 128 |
| Space reclamation commands .....                       | 129 |
| spacereclaimer start .....                             | 129 |
| spacereclaimer stop .....                              | 130 |
| spacereclaimer analyze .....                           | 130 |
| spacereclaimer status .....                            | 131 |
| Virtual Storage Console commands .....                 | 132 |
| vsc_config list .....                                  | 132 |
| vsc_config set .....                                   | 132 |
| vsc_config delete .....                                | 133 |

|                                         |     |
|-----------------------------------------|-----|
| Preferred IP address commands .....     | 134 |
| preferredIP set .....                   | 134 |
| preferredIP list .....                  | 135 |
| preferredIP delete .....                | 135 |
| iSCSI connection commands .....         | 136 |
| iscsi_target disconnect .....           | 136 |
| iscsi_target list .....                 | 136 |
| iSCSI initiator commands .....          | 137 |
| iscsi_initiator list .....              | 137 |
| iscsi_initiator establish_session ..... | 138 |
| iscsi_initiator terminate_session ..... | 139 |
| LUN commands .....                      | 140 |
| disk create .....                       | 140 |
| disk connect .....                      | 143 |
| disk delete .....                       | 146 |
| disk disconnect .....                   | 147 |
| disk resize .....                       | 148 |
| disk expand .....                       | 150 |
| disk add_initiator .....                | 151 |
| disk remove_initiator .....             | 153 |
| disk list .....                         | 154 |
| disk add_mount .....                    | 156 |
| disk remove_mount .....                 | 157 |
| disk rename_flexclone .....             | 158 |
| Snapshot copy commands .....            | 159 |
| snap create .....                       | 159 |
| snap delete .....                       | 160 |
| snap list .....                         | 161 |
| snap mirror_list .....                  | 162 |
| snap mount .....                        | 163 |
| snap rename .....                       | 164 |
| snap restore .....                      | 165 |
| snap unmount .....                      | 167 |
| snap update_mirror .....                | 168 |
| snap restore_volume_check .....         | 169 |
| snap restore_volume .....               | 170 |

|                                                                                     |            |
|-------------------------------------------------------------------------------------|------------|
| SnapVault commands .....                                                            | 171        |
| snapvault verify_configuration .....                                                | 171        |
| snapvault snapshot_rename .....                                                     | 172        |
| snapvault snapshot_delete .....                                                     | 173        |
| snapvault archive .....                                                             | 174        |
| snapvault relationship_status .....                                                 | 174        |
| snapvault snap_list .....                                                           | 175        |
| OnCommand commands .....                                                            | 176        |
| oncommand_config list .....                                                         | 176        |
| oncommand_config set .....                                                          | 176        |
| oncommand_config delete .....                                                       | 177        |
| oncommand_config rbaccache .....                                                    | 178        |
| Transport protocol commands .....                                                   | 179        |
| transport_protocol list .....                                                       | 179        |
| transport_protocol set .....                                                        | 180        |
| transport_protocol delete .....                                                     | 181        |
| Virtual server commands .....                                                       | 182        |
| vsconfig list .....                                                                 | 182        |
| vsconfig set .....                                                                  | 182        |
| vsconfig dslist .....                                                               | 183        |
| vsconfig delete .....                                                               | 184        |
| Hyper-V configuration commands .....                                                | 184        |
| hyperv_config list .....                                                            | 184        |
| hyperv_config set .....                                                             | 185        |
| hyperv_config delete .....                                                          | 186        |
| Clone split commands .....                                                          | 186        |
| clone_split clone_estimate .....                                                    | 186        |
| clone_split clone_start .....                                                       | 187        |
| clone_split clone_status .....                                                      | 187        |
| clone_split clone_stop .....                                                        | 188        |
| <b>Typical SnapDrive configurations .....</b>                                       | <b>189</b> |
| SnapDrive iSCSI configurations .....                                                | 189        |
| Single host direct-attached to a single storage system using iSCSI .....            | 189        |
| Single host attached to a single storage system through a GbE switch .....          | 189        |
| Single host attached to a single storage system through a dedicated<br>switch ..... | 190        |

|                                                                                                                   |            |
|-------------------------------------------------------------------------------------------------------------------|------------|
| Windows cluster connected to a storage system cluster through a<br>dedicated GbE switch .....                     | 191        |
| SnapDrive FC configurations .....                                                                                 | 191        |
| Single host direct-attached to a single storage system using FC .....                                             | 191        |
| Single host attached to a single storage system through an FC switch .....                                        | 192        |
| Windows cluster attached to a storage system active/active<br>configuration through an FC switch .....            | 192        |
| SnapDrive MPIO configurations .....                                                                               | 193        |
| Single host direct-attached to a single storage system using MPIO .....                                           | 193        |
| Windows cluster attached to a storage system active/active<br>configuration through a GbE switch using MPIO ..... | 194        |
| Windows cluster attached to a storage system active/active<br>configuration through an FC switch using MPIO ..... | 195        |
| SAN booting with SnapDrive .....                                                                                  | 195        |
| What SAN booting is .....                                                                                         | 195        |
| How SnapDrive supports SAN booting .....                                                                          | 196        |
| <b>Copyright information .....</b>                                                                                | <b>197</b> |
| <b>Trademark information .....</b>                                                                                | <b>198</b> |
| <b>How to send your comments .....</b>                                                                            | <b>199</b> |
| <b>Index .....</b>                                                                                                | <b>200</b> |

# 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's new in SnapDrive 7.1 for Windows

SDW 7.1 supports Data ONTAP 8.3, including the following highlighted new features.

You should be sure to check "Known limitations" in the *SnapDrive for Windows Release Notes* for information relating to limitations on Data ONTAP 8.3 support and on specific new features.

### New cloning capabilities

#### Clone of clones

When you create SnapManager for Microsoft SQL Server clone volumes, they are temporary volumes; when you disconnect, you lose access to the data in the cloned volume. You can use the new clone of clone capabilities to create a new clone containing your temporary data.

You can initiate the clone of clone operation using the `sdcli snap mount` and `disk connect` commands or from the Connect Disk wizard. Additionally, a new `-cn` parameter in `sdcli snap mount` and `disk connect` and **Volume Clone Name** option in the Disk Connect wizard, optionally specify custom names for the clones. When no custom names are specified, SnapDrive for Windows uses the default clone naming convention.

Creation of a clone on a clone created using Virtual Storage Console earlier than 4.2.2 is not supported.

#### Split clone

You can use the new split clone operations to split the new clone from the parent volume and create a permanent FlexVol volume.

Initiate the split clone operations using the following new `sdcli` commands and PowerShell cmdlets:

| sdcli commands                          | PowerShell cmdlets                          |
|-----------------------------------------|---------------------------------------------|
| <code>clone_split clone_estimate</code> | <code>Get-SdVolumeCloneSplitEstimate</code> |
| <code>clone_split clone_start</code>    | <code>Start-SdVolumeCloneSplit</code>       |
| <code>clone_split clone_status</code>   | <code>Get-SdVolumeCloneSplit</code>         |
| <code>clone_split clone_stop</code>     | <code>Stop-SdVolumeCloneSplit</code>        |

You cannot perform post-split clone operations on cloned VMDK disks created using Virtual Storage Console versions prior to 4.2.2.

### **SIS-clone throttling enhancement**

Starting with clustered Data ONTAP 8.3, your storage system does not accept clone requests if those requests negatively impact storage system client operations or if there is a substantial backlog. This is called *SIS-clone throttling*.

When you schedule a SIS-clone creation operation, you trigger two separate processes: the create operation, and the split operation. The create operation happens immediately, and results in a constant-time clone. The split operation runs in the background and copies blocks from the parent LUN to the child SIS-clone. This can take a significant amount of time and substantially impact your storage system.

If your clone creation operation fails as a result of SIS-clone throttling, you receive the following error message and should attempt your operation later:

```
Error message : Failed to create LUN clone since storage
system is experiencing high load.
```

You must have a FlexClone license for clone operations.

### **Dedicated disk optimization**

In SnapDrive 7.1 for Windows, you need to install SnapDrive only on the cluster node on which you are creating your dedicated disk. This is a change from previous releases, in which you must have SnapDrive for Windows installed on every node in your cluster. This change is applicable to Windows Server Failover Clustering environments.

To ensure this optimization works correctly, you must configure the transport protocol settings on the cluster node on which you created the dedicated disk.

### **Thin LUN provisioning option**

In previous SnapDrive for Windows releases, thin LUN provisioning was possible only when the space guarantee property was set to “none”. Starting in SnapDrive 7.1 for Windows, you can initiate thin provisioning regardless of the space guarantee property value.

You can create a thinly provisioned LUN using the new `-thin` option in the `sdcli disk create` command, or using the Create Disk wizard.

### **MetroCluster support**

MetroCluster in clustered Data ONTAP environments provides cluster-local high availability and cross-site synchronous disaster recovery.

When you are using MetroClusters with VMDK operations, you must restart SnapDrive for Windows service along with Virtual Storage Console SMVI service after every switch over and switch back operation.

For complete MetroCluster information, see the MetroCluster documentation. For information on using MetroCluster support with Virtual Storage Console, see the Virtual Storage Console 4.2.2 product documentation.

### **Microsoft Failover Cluster enhancements in vSphere 5.5 (ESX iSCSI and FCoE)**

SnapDrive for Windows now provides out-of-box iSCSI cluster support for VMware ESXi 5.5.

#### **Related information**

[\*Clustered Data ONTAP 8.3 MetroCluster Installation and Configuration Guide\*](#)

[\*Clustered Data ONTAP 8.3 MetroCluster Management and Disaster Recovery Guide\*](#)

## **Automated storage provisioning and data management using SnapDrive for Windows**

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

SnapDrive for Windows software integrates with Windows Volume Manager so that storage systems can serve as virtual storage devices for application data in Windows Server 2008 and later. You can also use SnapDrive for Windows to provision storage for Windows virtual machines hosted on ESX and Microsoft Hyper-V.

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

SnapDrive for Windows 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 for Windows on multiple hosts.

- It enables support on Microsoft cluster configurations.
- It enables iSCSI session management.

## List of current SnapDrive for Windows limitations

Some functionality is currently not supported in SnapDrive for Windows.

- A LUN managed by SnapDrive for Windows 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 for Windows 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 for Windows.
- SnapDrive for Windows supports LUNs on qtrees, but you cannot manage quotas from SnapDrive for Windows.  
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 for Windows supports the use of SnapMirror to replicate volumes but not individual qtrees.
- SnapDrive for Windows does not support the creation of CSVs with a drive letter or a volume mount point on Windows Server 2008 R2 and later.
- 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 prevent creation of new Snapshot copies until background sharing work is completed.

## Understanding SnapDrive for Windows components

Several components are integrated into the SnapDrive for Windows 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 for Windows workflows, including provisioning; Snapshot copy management; and backup, restore, and mounting operations.

### SnapDrive for Windows “snap-in”

This software module integrates with Microsoft Management Console (MMC) to provide you 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 for Windows command-line interface**

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

- Enter individual commands
- Run management scripts

### **PowerShell cmdlets**

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

SnapDrive for Windows supports PowerShell versions 2.0 and later.

### **Underlying SnapDrive for Windows 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 for Windows 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.

For complete information on using and troubleshooting VSS, see the Microsoft TechNet web site.

### Related information

*[Microsoft TechNet: Volume Shadow Copy Service](#)*

## Recommendations for using SnapDrive for Windows

You should follow certain recommendations whenever you use SnapDrive for Windows.

- Use OnCommand System Manager to set up and configure your storage system and to provision volumes for use with SnapDrive for Windows.
- Use SnapDrive for Windows to create and manage all of the LUNs on your storage system.
- If you want to dedicate all free space on a volume to LUNs, it is a best practice to set the snap reserve setting on the storage system to 0 percent.
- Unless you can be sure that name resolution publishes only the storage system interface you intend, you should 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.

# 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.

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:** You may use a dash (–) in the IQN name; however, you may not use an underscore (\_).

#### **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 must have an iSCSI session to the target on which you plan to manage the LUN.

### **Before you begin**

The iSCSI service must have been started.

### **About this task**

- An iSCSI session 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 for Windows that 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 that you want to establish the iSCSI session with, and then click **Next**.

The IP address that you enter is the storage system management port IP address, not the target portal IP address to which SnapDrive establishes an iSCSI session. You 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 that you want to use.

5. In the lower pane of the **Provide iSCSI HBA** panel, perform the following actions:

- a. Select the target portal with which you want SnapDrive to establish the iSCSI session by clicking the radio button next to the IP address of that portal.
- b. If your target requires authentication, select **Use CHAP**, and then type the user name and password that iSCSI should 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**.

### 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 for Windows support for ESX iSCSI initiators

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

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

- LUN enumeration using either the SnapDrive for Windows 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 ESX iSCSI initiators
- A maximum of 56 LUNs are supported in the guest OS.

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

## ESX limitations

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

- SnapDrive for Windows does not support iSCSI session management using ESX iSCSI initiators. You can add targets from the initiator list during LUN creation and connection.
- 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.
- For RDM LUNs, MPIO is not supported in a guest OS.  
If you are using MPIO within your guest OS using an iSCSI initiator, ensure that you do not have a single point of failure.
- Windows Server failover clustering is not supported using ESX 5.1 and earlier using iSCSI RDM LUNs.
- RDM LUNs larger than 2 TB are not supported in a VMFS 3.0 datastore.
- Windows Server Failover Clustering is not supported using FCoE and ESX iSCSI prior to version 5.5.

## Enabling storage system HTTP communication with SnapDrive

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

### Before you begin

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

### About this task

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

RPC communication is not supported in clustered Data ONTAP.

### Steps

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

#### Example of how to enable HTTP communication with SnapDrive for Windows

This example creates a role called `snaphttp2` with API and login capabilities on Storage1. Then it creates a group called **snapadmins** with **snaphttp2** role capabilities. Finally, it 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 Data ONTAP DSM support for Windows MPIO

You can use Data ONTAP DSM for Windows MPIO to help storage systems integrate with Microsoft MPIO on Windows Server 2008 R2 and later, and to 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 Server R2 x64 and later include support for Microsoft MPIO.

Data ONTAP DSMs for Windows MPIO help storage systems to integrate with Microsoft MPIO on Windows Server 2008 R2 and later, 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 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.

DSM is not required in a heterogenous system configuration.

## Support for Microsoft Cluster Shared Volumes

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

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

### Related information

[\*Microsoft TechNet: Use Cluster Shared Volumes in a Failover Cluster\*](#)

[\*Microsoft TechNet: Cluster Shared Volumes Support for Hyper-V\*](#)

[\*Microsoft TechNet: Requirements for Using Cluster Shared Volumes in a Failover Cluster in Windows Server 2008 R2\*](#)

## 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

In Windows Server 2008, enter the following command for a cluster group called Available Storage:

```
cluster group "Available Storage" /status"
```

### Example

In Windows Server 2012, you can run a PowerShell cmdlet to verify the cluster group owner:

```
get-ClusterGroup
```

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
```

## 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 for Windows, you must keep some rules in mind.

- You should create LUN names using US-ASCII characters only, even when you are using non-ASCII operating systems.
- If you are adding a CSV LUN to a Windows Server 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:
  - Shared disks on Windows Server 2008 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.
  - You should 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, you should 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 for Windows.
- You cannot create a LUN on a SnapDrive for Windows-created clone volume, unless the clone has already been split.

## 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 168 LUNs. By using volume mount points, you can graft, or mount, a target partition to 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.

## Volume mount point requirements

You should be aware of some key volume mount point requirements when creating mount points on cluster Windows Servers.

When creating mount points on cluster 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.

## Creating a dedicated LUN

You can use SnapDrive for Windows to create dedicated Fibre Channel (FC) or iSCSI-accessed LUNs.

### Before you begin

- You have created dedicated volumes to hold your LUNs on the storage system.
- You have started FC or iSCSI services on the storage system.

### 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.
- LUNs should be created using SnapDrive for Windows.
- 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 for Windows displays the message `Checking access control` to indicate that it is checking these access control settings.

- If you are operating in Data ONTAP 8.3 or later, it is sufficient to have SnapDrive for Windows installed only on the cluster node on which you are creating the disk.

### Steps

1. Launch the **Create Disk** wizard from the SnapDrive for Windows **Action** pane.
2. Specify whether you want to create a shared or dedicated disk in the **Select a LUN Type** page.
3. Specify the necessary information in the **Provide a Storage System Name, LUN Path and Name** page.
4. In the **Select LUN Properties** page, provide the following information:
  - Indicate the drive letter of volume mount point.  
The root of the volume mount point must be owned by the node on which you are creating the new disk.  
You can create cascading volume mount points (one mount point mounted on another mount point).
  - Specify whether you want to limit the maximum disk size to accommodate at least one Snapshot copy.  
Limiting the maximum disk size is the recommended option and the default.
  - Select a LUN size.
  - Indicate whether you want to create a thinly provisioned LUN.
5. The **Important Properties of the Storage System Volume** page displays the settings to use for the volume or qtree.

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

- `create_ucose = on`
- `convert_ucose = on`
- `snapshot_schedule = off`

**Note:** The `create_ucose` and `convert_ucose` volume options are no longer used, but they are set to maintain backwards compatibility with earlier versions of SnapDrive for Windows.

6. Select an initiator in the **Select Initiators** page.

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

7. In the **Select Initiator Group Management** page, specify whether you want to use automatic or manual igroup management.

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

8. In the **Completing the Create Disk Wizard** page, verify the settings and click **Finish**.

## Creating a shared LUN

You can use SnapDrive for Windows 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

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

### Steps

1. Launch the **Create Disk** wizard from the SnapDrive for Windows **Action** pane.
2. Specify whether you want to create a shared or dedicated disk in the **Select a LUN Type** page.
3. Specify the necessary information in the **Provide a Storage System Name, LUN Path and Name** page.
4. Review the information about Microsoft cluster services in the **Information About the Microsoft Cluster Services System** and **Specify Microsoft Cluster Services Group** pages.
5. In the **Select LUN Properties** page, provide the following information:
  - Indicate the drive letter of volume mount point.  
The root of the volume mount point must be owned by the node on which you are creating the new disk.  
You can create cascading volume mount points (one mount point mounted on another mount point).
  - Specify whether you want to limit the maximum disk size to accommodate at least one Snapshot copy.  
Limiting the maximum disk size is the recommended option and the default.

- Select a LUN size.
- Indicate whether you want to create a thinly provisioned LUN.

6. Ensure that the properties displayed in the **Important Properties of the Storage System Volume** page are correct.

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

- create\_ucose = on
- convert\_ucose = on
- snapshot\_schedule = off

**Note:** The create\_ucose and convert\_ucose volume options are no longer used, but they are set to maintain backwards compatibility with earlier versions of SnapDrive for Windows.

7. In the **Select Initiators** page, perform the following actions for each host:

- Double-click the cluster group name to display the hosts that belong to the cluster.
- Click the name of a host to select it.
- Select an initiator for the LUN you are creating.

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

8. In the **Select Initiator Group Management** page, specify whether you want to use automatic or manual igroup management.

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

9. In the **Completing the Create Disk Wizard** page, verify the settings and click **Finish**.

## Windows Server failover cluster support

SnapDrive for Windows supports the use of shared and dedicated LUNs in a Windows Server failover cluster for all cluster configuration models using up to 16 nodes in SAN environments and up to 32 nodes in SMB 3.0 environments. However, you cannot perform enumeration and management of offline disk resource tasks in a failover cluster.

For more information about running SnapDrive for Windows in an SMB 3.0 environment, see the *SnapDrive for Windows Administration Guide for SMB 3.0 Environments* and the *SnapDrive for Windows PowerShell Cmdlet Reference Guide*.

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

**Related information**

[\*Microsoft TechNet: Failover Clusters in Windows Server 2008\*](#)

[\*Microsoft TechNet: Failover Clustering Overview\*](#)

[\*Microsoft TechNet: Windows Server 2008 Failover Cluster Overview\*](#)

[\*Microsoft Developer Network: Windows Server Failover Clustering \(WSFC\) with SQL Server\*](#)

**Configuring a Windows Server failover cluster witness disk**

SnapDrive for Windows and Windows Server 2008 and 2012 provide a simpler way of configuring a shared disk as a witness disk. You no longer have to create a shared disk before creating a cluster, or connect that shared disk to the node before that node is added to the cluster.

**Before you begin**

- The Windows Server failover cluster feature must have been installed.
- The failover cluster using the Windows Server MMC snap-in, Failover Cluster Management, must have been created.
- A shared LUN must have been created, ensuring that you select the Microsoft Cluster Services Group named “Cluster Group” to own that disk resource.

**Note:** The shared LUN must have been created on the node that owns “Cluster Group.” To determine which node owns “Cluster Group” in Windows Server 2008, you can type the `cluster group` command at a Windows command prompt. In Windows Server 2012, you can run the `Get-ClusterGroup` PowerShell cmdlet.

**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 **Select Quorum Configuration** panel, select the type of cluster that makes sense for the number of nodes in your cluster.

It is recommended that you select **Node and Disk Majority** for clusters with an even number of Windows cluster node, and that you select Node Majority for clusters with an odd number of Windows cluster nodes.

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

6. In the **Confirmation** panel, click **Next** to configure the cluster quorum settings.
7. In the **Summary** panel, click **Finish** to close the wizard.

#### Related information

*Microsoft TechNet: Understanding Quorum Configurations in a Failover Cluster*

## Support for creating disks on a SVM operating in clustered Data ONTAP

You can use SnapDrive for Windows to create disks on a Storage Virtual Machine (SVM) that is operating in clustered Data ONTAP. You must have SVM 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 SVM in clustered Data ONTAP

You must be aware of the configuration requirements before you create a disk on a Storage Virtual Machine (SVM) in clustered Data ONTAP.

- The SVM administrator (vsadmin) account must be unlocked.  
For more information about how to unlock a vsadmin account, see the *Data ONTAP Cluster-Mode Administration Reference* [Data ONTAP Cluster-Mode Commands: Manual Page Reference](#).
- SAN data protocol must be enabled and configured on the SVM.
- You must use HTTP or HTTPS protocol to communicate with the SVM.  
You must not use RPC protocol, which is not supported in clustered Data ONTAP.
- The SVM 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 the SVM, you must add your storage system name and IP address to the Windows Server etc\host file or add the SVM host name into the Active Directory DNS records.  
This is because the Windows host must be able to resolve the SVM management LIF IP address to the SVM name.

For more information about SVMs and SVM administrator capabilities, see the *Clustered Data ONTAP System Administration Guide for SVM Administrators*.

## GPT partition support

SnapDrive for Windows creates LUNs using the GUID partition table (GPT) format.

Neither SnapDrive for Windows 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 for Windows, 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

#### **Related information**

*[Microsoft Developer Network: Windows and GPT FAQ](#)*

## **Configuring SnapDrive for Windows to create LUNs with virtual Fibre Channel**

Windows Server 2012 and later 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.

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

#### **Related information**

*[Microsoft TechNet: Implement Hyper-V Virtual Fibre Channel](#)*

## **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 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

You can use a LUN managed by SnapDrive for Windows for data storage, and it can be a boot disk. A LUN cannot be a dynamic disk.

SnapDrive for Windows 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.

A LUN managed by SnapDrive for Windows can only have a single partition. SnapDrive for Windows does not support managing LUNs with more than one partition.

## 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.

## 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.

- You cannot simultaneously map one LUN to two different Windows hosts.
- Offline LUNs are visible from the SnapDrive for Windows 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, you should unmap the LUN, bring it online from the storage system, and refresh the SnapDrive for Windows 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.
- SnapDrive 7.1 for Windows does not support operations on dynamic disks.

## 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. Connecting to a LUN allows you to manage it.

### About this task

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

### Steps

1. Expand the SnapDrive instance you want to manage and select **Disks**.
2. Launch the **Connect Disk** wizard from the **Action** pane.
3. In the **Select a LUN Type** page, specify whether the disk is a shared or dedicated disk.
4. If you have specified that you want to connect to a shared disk, review the Cluster Services information in the **Specify Microsoft Cluster Services Group** pane.
5. In the **Provide a Storage System Name, LUN Path and Name** panel, specify the required information.
6. Indicate a drive or mount point for the LUN to which you are connecting.
7. In the **Select Initiators** panel, choose at least one initiator for the LUN.  
  
If the LUN to which you are connecting belongs to a cluster, select at least one initiator for each host in the cluster. If MPIO is installed on the system, you can select multiple FC initiator ports or one iSCSI session.
8. In the **Select Initiator Group Management** page, specify whether you want to use automatic or manual igroup management.  
  
If you select automatic igroup management, SnapDrive for Windows uses existing igroups or, when necessary, creates new igroups. If you select manual igroup management, you manually choose existing igroups or create new ones as needed.
9. In the **Completing the Create Disk Wizard** page, verify the settings and click **Finish**.

## 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.
2. 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.

**Note:** The Change option is unavailable for mount points.

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

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.

## 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.

- You must ensure that the LUN is not in use by any application.
- 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 any monitoring software.
- You must ensure that the LUN you want to disconnect or delete is not being used by a host.
- You can disconnect or delete a shared LUN (that is, a nonquorum disk) only after removing the cluster resource dependencies from the LUN and verifying that all nodes in the cluster are powered on and functioning properly.  
SnapDrive for Windows manages the dependencies to the root disk for volume mount points.
- When disconnecting or deleting LUNs on a Microsoft cluster, you must ensure that all hosts in the cluster are available to SnapDrive for Windows 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 for Windows 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 for Windows, 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, you should rename the volume before disconnecting the last LUN. When you rename the volume, ensure that you change more than just the last integers in the name.
- 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

- Neither Windows Explorer nor any other Windows application can be using or displaying any file on the LUN that you intend to disconnect; otherwise, you must force the disconnect.
- If you are disconnecting a LUN that contains volume mount points, they must be removed first.
- You must 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.
2. 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.

This procedure does 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 remains 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 should delete a folder within a volume mount point by bypassing the Windows Recycle Bin. Failing to bypass the Recycle Bin when deleting volume mount points might cause several errors.

### About this task

When you use 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, in which *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.

Additionally, if you have deleted the entire mount point using Windows Explorer, the mount point displays from the Recycle Bin, and continues to do so until you have deleted the mount point from the recycle bin.

### Steps

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

### Related information

[\*Microsoft Support Article 243514: "Access Denied" When You Delete Folders from a Mounted Drive\*](#)

## 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 for Windows 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: On Microsoft Windows Sever 2008 SP2, Windows 7, and Windows Server 2008 R2, you may not be able 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.

#### 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, perform a disk rescan operation in the Disk Management MMC snap-in so that the resized disk becomes visible.

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

SnapDrive for Windows supports the Windows Server feature that enables you to add or remove a pass-through 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.
- .NET 3.5 SP1 or later must be installed on the Hyper-V parent hosts.
- The version of SnapDrive for Windows installed on the parent hosts and the VMs must be the same.
- If you are using iSCSI or FC for pass-through disk creation, the igroups must already be configured between the SAN host and your storage system.
- 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 for Windows, all SnapDrive for Windows registry key entries are removed from your Windows host.  
If you reinstall SnapDrive for Windows, you must reconfigure pass-through disks.
- When creating pass-through disks in the MMC, SnapDrive for Windows 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 that are 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.
- Pass-through LUN creation fails when Data ONTAP DSM for Windows MPIO 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 not running, the virtual machine to which the disk is assigned cannot be restarted.  
To reboot the virtual machine, you must bring the pass-through disk back online, 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 with different owner nodes, 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 can use SnapDrive for Windows to ensure that you create consistent Snapshot copies if you need to restore a LUN from that copy.

Snapshot operations on a single LUN 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. Therefore, it is important to back up a LUN using a SnapDrive Snapshot copy rather than using other means.

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 OnCommand System Manager or 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, which 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.

## Creating a Snapshot copy

You should always use SnapDrive for Windows to create Snapshot copies of LUNs to ensure that Snapshot copies are consistent at the file system level.

### Before you begin

- You have created Snapshot copies through the SnapDrive for Windows 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 for Windows 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.

- You have created Snapshot copy names using US-ASCII characters only, even when using non-ASCII operating systems.

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

### About this task

While SnapDrive for Windows helps ensure Snapshot copy consistency at the file system level, you should always use SnapManager to create Snapshot backup copies of LUNs to ensure that Snapshot copies are consistent on the application level.

### Steps

- Perform the following actions to get to the **Create Snapshot** menu item:
  - In the left MMC pane, select the instance of SnapDrive for Windows you want to manage.
  - Double-click **Disks**.
  - Double-click the disk for which you want to create a Snapshot copy.
  - Select **Snapshots**.
  - From the menu choices at the top of MMC, navigate to **Action > Create Snapshot**.
- In the **Create Snapshot** text box, perform the following actions:
  - 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

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

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

SnapDrive for Windows 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 volume).

FlexClone volumes do not support traditional volumes.

## Prerequisites for using FlexClone volumes

Several prerequisites must be met 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 FlexClone volumes.
- There must be enough space available on the aggregate to create a non-space-reserved FlexVol volume (volume guarantee=none).
- FlexClone licenses must be installed on your storage system.
- You must have a proper Storage Virtual Machine (SVM) configuration.  
In clustered Data ONTAP, the aggregate must be assigned to the SVM for successful SnapDrive for Windows operations.

## 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.

## Restrictions when working with Snapshot copies

It is useful to keep certain requirements 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 for Windows 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.

## 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 must be enabled.
- You must 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 for Windows that you want to manage, and 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 to which you want to connect.
3. From the menu choices at the top of MMC, navigate to **Action > Connect Disk** to launch the **Connect Disk** wizard and follow the wizard instructions.

4. 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**.
5. In the **Select Initiator Group management** panel, specify whether you want to use automatic or manual igroup management.

## 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.

## How Data ONTAP LUN clone split (rapid LUN restore) works

SnapDrive for Windows uses the LUN clone split (rapid LUN restore) feature of Data ONTAP when restoring a LUN. Because using this feature splits the clone from the backing Snapshot copy in the background, the restored LUN is available to the Windows host for I/O operations within only a few seconds.

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 performed on the clone rather than 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 or until it is being split. 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 is not initiated and no LUN restoration occurs.

**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 for Windows 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.

### About this task

Even before the Snapshot copy limit is reached, a Snapshot copy fails if insufficient reserved space for it remains on the disk. If the Snapshot copy was created by SnapManager, you should delete it with the SnapManager Delete Backup wizard instead of deleting it within SnapDrive for Windows. This helps SnapManager avoid becoming unaware that the backup set to which the Snapshot copy has become an orphan of one or more Snapshot copies.

### Steps

1. Perform the following actions to select the LUN with the Snapshot copy you want to delete:
  - a. In the left MMC pane, select the instance of SnapDrive for Windows 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

In Data ONTAP operating in 7-Mode environments, you can use `storaccl.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.

### Before you begin

- You have run `storaccl.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 `storaccl.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.

For information about how to create thinly provisioned LUNs in clustered Data ONTAP, see the `create disk` command syntax page.

### Steps

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

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 for Windows 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, CSVs, and 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, use the Interoperability Matrix to ensure that you are using the recommended version of Windows Host Utilities (WHU) for the your version of Data ONTAP.

## Related information

*[NetApp Interoperability Matrix Tool](#)*

## 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...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Space to reclaim</b>               | Confirming that the LUN space can be optimized.<br>Continue to the next step.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>No space to reclaim</b>            | Notifying you that you do not need to run Space Reclaimer on the selected disk.<br><br>Click Cancel to exit the Confirm Space Reclamation on Disk window, or go to the next step to continue with space reclamation.<br><br><b>Note:</b> 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.

In Windows Server 2008, 2008R2, and 2012:

```
PS C:\Users\administrator> $HyperVGuestName = "Vm_Name";
PS C:\Users\administrator>
```

```

PS C:\Users\administrator> $VMManagementService = gwmi -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>

```

In Windows Server 2012 R2:

```

PS C:\Users\administrator> $HyperVGuestName = "Vm_Name";
PS C:\Users\administrator>
PS C:\Users\administrator> $VMManagementService = gwmi -Class
"Msvm_VirtualSystemManagementService" -Namespace "root\virtualization
\v2"
PS C:\Users\administrator>
PS C:\Users\administrator> $Vm = gwmi -Namespace "root\virtualization
\v2"
-Query "Select * from Msvm_ComputerSystem Where ElementName=
'$HyperVGuestName'"
PS C:\Users\administrator>
PS C:\Users\administrator> $SettingData = gwmi -Namespace
"root\virtualization\v2" -Query "Associators of {$Vm}
Where ResultClass=MSVM_VirtualSystemSettingData"
PS C:\Users\administrator>
PS C:\Users\administrator> $SettingData.AllowFullSCSICommandSet =
$true
PS C:\Users\administrator>
PS C:\Users\administrator> $VMManagementService.ModifySystemSettings
($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.

## 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.

### Supported VMware guest OS platforms

You can use SnapDrive for Windows to provide LUN provisioning and Snapshot copy management support on x86 and x64 guest operating systems running on VMware ESXi 5.0U3 or higher.

SnapDrive for Windows supports the following VMware guest OS configurations:

- Windows Server 2008 and 2012 and later, on x86 and x64 platforms
- Microsoft cluster configurations of up to a maximum of 16 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, or 42 RDM LUNs with three LSI Logic SCSI controllers on a 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 for Windows is supported on VMware ESX server. Before you use SnapDrive for Windows to perform provisioning and Snapshot copy management operations, you should be aware of some limitations.

- Installing SnapDrive for Windows on a Microsoft cluster on virtual machines using ESX credentials is not supported.  
You should use the vCenter credentials when installing SnapDrive for Windows 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.
- The VMFS datastore housing the RDM descriptor files must be VMFS 5.0 and both ESXi and vCenter must be version 5.0U3 or later, 5.1U2 or later, or 5.5 or later.
- 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 login from SnapDrive for Windows MMC

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

### About this task

Keep the following information in mind when setting vCenter or ESX login from SnapDrive for Windows MMC:

- You cannot disable vCenter or ESX login from SnapDrive for Windows MMC when RDM LUNs are present.  
You must disconnect or delete RDM LUNs before you can disable ESX login.
- To benefit from VMware vMotion technology, ensure that SnapDrive for Windows is configured to connect to vCenter instead of to the ESXi host directly.

### Steps

1. In the left MMC pane, select the instance of SnapDrive for Windows for which you want to enable or disable vCenter or ESX login.
2. 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 login, 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 host name, 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 login, 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 for Windows cannot display WWPNs or IQNs for FC or iSCSI initiators on the ESXi host.

## 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.

## 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 for Windows to create FC-accessed, iSCSI-accessed, or ESX iSCSI-accessed RDM LUNs on a guest OS.

#### Before you begin

- You have created the dedicated volumes to hold your LUNs on the storage system.
- You have verified that the FC or iSCSI service has been started on the storage system.
- You have manually created initiator groups either by using the OnCommand System Manager user interface or at the storage system console.  
You have chosen an initiator for an ESX server to create a single initiator group automatically when you created a shared FC or iSCSI RDM LUN.

#### Steps

1. Select the SnapDrive for Windows instance that you want to manage, and then select **Disks**.
2. Launch the **Create Disk** wizard.
3. In the **Select a LUN Type** page, select **Dedicated**, and then click **Next**.
4. In the **Provide a Storage System Name, LUN Path and Name** page, provide the necessary storage system and LUN information.

5. In the **Select LUN Properties** page, provide the following information:

- Indicate the drive letter of volume mount point.  
The root of the volume mount point must be owned by the node on which you are creating the new disk.  
You can create cascading volume mount points (one mount point mounted on another mount point).
- Specify whether you want to limit the maximum disk size to accommodate at least one Snapshot copy.  
Limiting the maximum disk size is the recommended option and the default.
- Select a LUN size.
- Indicate whether you want to create a thinly provisioned LUN.

6. Check the **Important Properties of the Storage System Volume** page to ensure you have the correct settings to use for the volume or qtree.

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

- create\_unicode = on
- convert\_unicode = on
- snapshot\_schedule = off

**Note:** The create\_unicode and convert\_unicode volume options are no longer used, but they are set to maintain backward compatibility with earlier versions of SnapDrive for Windows.

7. In the **Initiator List** page, select an initiator for the LUN that you are creating.

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

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

8. In the **Select Initiator Group Management** page, specify whether you want to use automatic or manual igroup management.

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

9. In the **Select a Datastore** panel, perform the following steps.

| If your virtual machine resides on a... | Then do this...                                                                                                                                                                                                                                                                                        |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>VMFS datastore</b>                   | <ol style="list-style-type: none"> <li>Choose either <b>Store with Virtual Machine</b>, which is the default, or choose <b>Specify datastore</b> to select a different VMFS datastore for your virtual machine disk format (VMDK) file to reside.</li> <li>Click <b>Next</b>.</li> </ol>               |
| <b>NFS datastore</b>                    | <ol style="list-style-type: none"> <li>Choose <b>Specify datastore</b> to select a different VMFS datastore for your VMDK file to reside. The <b>Store with Virtual Machine</b> option is unavailable because a VMDK file cannot be stored on an NFS datastore.</li> <li>Click <b>Next</b>.</li> </ol> |

10. In the **Completing the Create Disk Wizard** page, verify the settings and click **Finish**.

## 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 for Windows 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.

Starting with ESXi 5.5, you can also use ESX iSCSI and FCoE hardware to manage to manage a Microsoft cluster. SnapDrive 7.1 for Windows includes out-of-box support for Microsoft clusters.

### Requirements for using FC RDM LUNs in a Microsoft cluster

SnapDrive for Windows 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 VMs must be running the same Windows Server version, Windows Server 2008 or later.
  - 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.
- SnapDrive for Windows cannot create a single igroup with initiators from multiple hosts. The igroup containing the initiators from all ESXi hosts must be created on the storage controller prior to creating the RDM LUNs that will be used as shared cluster disks.
- 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 prior to ESX 5.5i this feature is not supported.

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

### 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 for Windows MMC snap-in. Connecting to a LUN enables you to manage it.

### Before you begin

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

### Steps

1. Select the instance of SnapDrive for Windows that you want to manage.
2. Launch the **Connect Disk** wizard from **Action > Connect Disk**.
3. In the **Select a LUN Type** page, select **Dedicated**, and then click **Next**.
4. In the **Provide a Storage System Name, LUN Path and Name** page, provide the necessary storage system and LUN information.
5. In the **Select LUN Properties** page, select the drive letter or enter a mount point for the LUN to which you are connecting.

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

6. 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**.  
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.
7. In the **Select Initiator Group Management** panel, specify whether you expect to use automatic or manual igroup management.  
If you select automatic igroup management, SnapDrive for Windows 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.
8. If you selected an FC initiator on an ESX server in the **Select Initiator Group Management** panel, then the **Select a Datastore** page displays.

| If your virtual machine resides on a... | Then do this...                                                                                                                                                                                                                                                                                            |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>VMFS datastore</b>                   | <ol style="list-style-type: none"> <li>Choose either <b>Store with Virtual Machine</b>, which is the default, or choose <b>Specify datastore</b> to select a different VMFS datastore for your virtual machine disk format (VMDK) file to reside.</li> <li>Click <b>Next</b>.</li> </ol>                   |
| <b>NFS datastore</b>                    | <ol style="list-style-type: none"> <li>Choose <b>Specify datastore</b> to select a different VMFS datastore for your VMDK file to reside.<br/>The <b>Store with Virtual Machine</b> option is unavailable because a VMDK file cannot be stored on an NFS datastore.</li> <li>Click <b>Next</b>.</li> </ol> |

- Verify your information in the **Completing the Connect Disk Wizard** panel, and click **Finish**.

The newly connected LUN 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

## 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.

### Requirements for using SnapVault

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

#### System requirements for performing SnapVault operations

- 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 for Windows.
- SnapVault relationships must be configured and initialized.

#### Backup set requirements for SnapVault operations

- 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

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

## Initiating SnapVault backup jobs from SnapDrive for Windows

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

### About this task

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. Launch `sdcli`.
2. Run `sdcli snapvault archive` and specify the machine name, the name of the Snapshot copies on your secondary system, and the list of mount points and Snapshot copies you want to archive:

### Example

```
sdcli snapvault archive
[-m MachineName]
[-force]
-a ArchivalSnapshotName
-DS MountPointandSnapshotList [...]
```

## 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 the SnapDrive for Windows configuration requirements for SnapVault support, and of the support SnapVault operations.

To ensure that SnapDrive for Windows supports SnapVault operations, you must configure both the primary and the secondary storage system login details in the transport protocol settings.

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 copy 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 for Windows checks for any existing rolling Snapshot copies of the LUN containing the specified LUN drive. What happens next depends on whether SnapDrive for Windows detects none, one, or two Snapshot copies.

### No Snapshot copies detected

If SnapDrive for Windows does not find any rolling Snapshot copies containing the LUN image, it creates a rolling Snapshot copy on the SnapMirror source volume. SnapDrive for Windows then initiates a SnapMirror update operation, which replicates the rolling Snapshot copy on the destination volume.

### One Snapshot copy detected

If SnapDrive for Windows finds one rolling Snapshot copy containing the LUN image, it creates a second rolling Snapshot copy and initiates a SnapMirror update.

### Two Snapshot copies detected

If SnapDrive for Windows detects two rolling Snapshot copies containing the LUN image, it deletes the older rolling Snapshot copy and creates a new one to replace it. Then SnapDrive for Windows initiates a SnapMirror update.

When you use SnapDrive for Windows to connect to a LUN in a SnapMirror destination, a Snap Restore operation is initiated from the last Snapshot copy and the restored LUN is connected. 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 for Windows

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 information about how to license and set up SnapMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.  
  
**Note:** In clustered Data ONTAP 8.1.x environments, you must provide cluster credentials in the transport protocol settings before performing SnapMirror operations. In clustered Data ONTAP 8.2 and later, Storage Virtual Machine (SVM) credentials are required.
- 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 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 SVM (vFiler units), you must create your SnapMirror relationship on the vFiler units, 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 for Windows 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 for Windows 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 for Windows 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 for Windows service.  
You should add the domain account to the storage system administrator groups. Additionally, you should set the `waf.l.map_nt_admin_priv_to_root` option 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 for Windows 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 for Windows queues the others until the first task either succeeds or times out.

**Note:** SnapMirror cascaded configurations are not supported.

## Initiating replication manually

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

### Before you begin

You have manually created a Snapshot copy.

### About this task

Manual replication is not monitored by SnapDrive for Windows, 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 for Windows 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 you create the Snapshot copy on the mirrored source volume, SnapDrive for Windows 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

- You have restored the LUN on an asynchronous SnapMirror destination from the most recent SnapDrive for Windows Snapshot copy containing a valid image of that LUN.
- You have performed the restoration using SnapDrive for Windows, as part of the LUN connect operation on an active file system or on a SnapMirror destination volume.
- You have created the most recent Snapshot copy using SnapDrive for Windows to ensure data consistency.

If Data ONTAP creates a Snapshot copy that is more recent than the Snapshot copy created by SnapDrive for Windows, the Data ONTAP Snapshot copy cannot be used by SnapDrive for Windows because it is not consistent.

### About this task

SnapDrive for Windows 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.

### 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 for Windows 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 for Windows can initiate a volume-based Snapshot copy restoration:

- All LUNs on the active file system are consistent in the Snapshot copy you intend to use to restore.
- LUNs on the active file system are of the same size and have the same name as the selected 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**

Specifies the name of the storage system on which the volume resides.

#### **-volume**

Indicates name of the volume on which the restore operation will be performed.

#### **-s**

Specifies 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.

#### **-m**

Specifies 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

The following 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`.

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

## 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 for Windows integrates with OnCommand Unified Manager Core Package data protection capabilities

SnapDrive for Windows 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 for Windows 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 for Windows 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 for Windows, see the Interoperability Matrix for supported Data ONTAP and OnCommand Unified Manager Core Package versions.

For more information, see your SnapManager documentation.

**Related information**

*[NetApp Interoperability Matrix Tool](#)*

## Configuring and managing access control

---

You can configure access control by enabling RBAC with SnapDrive for Windows, and in Data ONTAP operating in a 7-Mode environment, 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 `storac1.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                                                                                                                                                       |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>SD.Config</code>         | Allows users to perform all configuration operations.                                                                                                              |
| <code>SD.Config.Read</code>    | Allows users to only read configuration information on the storage system.                                                                                         |
| <code>SD.Config.Write</code>   | Allows users to create and modify configuration information on the storage system, such as creating and modifying igroups, creating mappings, and setting options. |
| <code>SD.Config.Delete</code>  | Allows users to delete any configuration information.                                                                                                              |
| <code>SD.Storage</code>        | Allows users to perform all storage operations.                                                                                                                    |
| <code>SD.Storage.Read</code>   | Allows users to list storage system objects, such as enumerating and reading aggregates, volumes, qtrees, LUNs, and files.                                         |
| <code>SD.Storage.Write</code>  | Allows users to create and modify storage system objects.                                                                                                          |
| <code>SD.Storage.Delete</code> | Allows users to delete storage objects.                                                                                                                            |

| Operation           | Descriptions                                                                                            |
|---------------------|---------------------------------------------------------------------------------------------------------|
| 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.              |
| 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 | <ul style="list-style-type: none"> <li>• SD.Storage.Write</li> <li>• SD.Storage.Read</li> <li>• SD.Config.Read</li> <li>• SD.Config.Write</li> </ul> | <p>Allows all LUN provisioning and Snapshot copy operations, including create, connect, and map, if the operations are set on the storage system.</p> <p>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 <code>sdc li.exe</code>.</p> |
| SD.Discovery | <ul style="list-style-type: none"> <li>• SD.Config.Read</li> <li>• SD.Storage.Read</li> <li>• SD.Snapshot.Read</li> </ul>                            | Allows all operations for discovering volumes, qtrees, igroups, and Snapshot copies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| SDBackup     | <ul style="list-style-type: none"> <li>• SD.Snapshot.Read</li> <li>• SD.Snapshot.Write</li> </ul>                                                    | Allows create, replicate, and archive backup operations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| SDRestore    | <ul style="list-style-type: none"> <li>• SD.Snapshot.Restore</li> </ul>                                                                              | Allows restore operations from a Snapshot copy or archive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| SDNoAccess   | <ul style="list-style-type: none"> <li>• SD.Access.None</li> </ul>                                                                                   | <p>Denies all access.</p> <p>When you use this role, the storage access control tool does not allow other roles to exist.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## 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 `storac1.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 set
- Add, 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 existing role**

This example adds storage write access to the role called `TESTROLE`:

```
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 accessible to all users**

This example lists the resources that are accessible to all users:

```
user list
```

#### **Remove all access rights to a user**

This example removes all access rights to the domain user usr1:

```
user remove -un mydomain\usr1
```

#### **List storage system resources**

This example lists all volumes on the storage system:

```
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 vol1 -stor System1 -user root
```

When no protocol type is specified, storac1 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:

```
spacereserve set -vol vol1 -val true -stor System1 -user root
```

### Disable space reservation

This example disables space reservation on the volume vol1 on the storage system System1:

```
spacereserve set -vol vol1 -val false -stor 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 Interoperability Matrix .
- 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 ToolChest on the NetApp Support Site 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 for Windows operations map to the DataFabric Manager roles as follows:

| 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                                   |
| disk connect<br>disk disconnect<br>disk disconnect -f | SD.Storage.Write    | Enables LUN connect, disconnect, and forced disconnect when set at the disk resource |

| <b>SnapDrive operation</b> | <b>Role and capability</b> | <b>Description</b>                                                                                      |
|----------------------------|----------------------------|---------------------------------------------------------------------------------------------------------|
| 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<br>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<br>Set on the qtree for LUN clones: SD.Snapshot.Clone<br>Set on the storage system for traditional volume clones: SD.Snapshot.Clone<br>Set on the parent volume for FlexVol volume clones: SD.Snapshot.Clone<br>Set on the parent for FlexVol volume clones that are split: SD.Snapshot.Clone<br>SD.Storage.Write | Enables Snapshot copy mounting when set on the specified resources                                        |

| <b>SnapDrive operation</b> | <b>Role and capability</b>                                                                                                               | <b>Description</b>                                                                                  |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| snapshot unmount           | For LUN clones in a volume or qtree: SD.Snapshot.Clone on the volume<br><br>For volume clones:<br>SD.Snapshot.Clone on the parent volume | Enables unmounting of LUN clones or volume clones on the indicated resources                        |
| 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**

C:

**cd \Program Files\NetApp\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 <code>sdcli</code> cannot find the drive letter specified through the <code>-d</code> 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     | <p>The initiator name.</p> <ul style="list-style-type: none"> <li>• For FC, the initiator name is the WWPN (World Wide Port Name) for the initiator, which takes the form <code>hh:hh:hh:hh:hh:hh:hh:hh</code>.</li> <li>• For iSCSI, the initiator name takes the form <code>iqn.iSCSI</code> qualified name. For more information on iSCSI node names, see the <i>Data ONTAP SAN Administration Guide for 7-Mode</i>.</li> </ul> |

| Switch | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| - I    | <p>The list of hosts and initiators.</p> <p>Separate the character strings that specify hosts and initiators with spaces.</p> <p>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.</p> <p>To specify the initiator, type the appropriate WWPN, which you can determine through the <code>lputilnt.exe</code> utility supplied with your FC HBA Attach Kit. After you launch <code>lputilnt.exe</code>, navigate to <b>Main Menu &gt; Adapter &gt; Configuration Data</b> and select “16 - World-Wide Name” in the Region field. The available WWPNs appear in the list box directly beneath the Region field.</p> <p>When MPIO is running, you can specify up to four node-initiator pairs. The first <code>NodeMachineName</code> in the cluster applies to two of the available initiator WWPNs; the other <code>NodeMachineName</code> applies to the remaining pair of initiator WWPNs.</p> |
| - IG   | <p>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).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| - m    | <p>The host on which the operation is executed. You can use an IP address or a machine name to identify the host.</p> <p><b>Note:</b> When you use an IP address to identify a host, you should use only those IP addresses displayed in the output from the <code>ipconfig /all</code> command rather than those displayed in the output from the <code>sdcli sysconfig list</code> command.</p> <p><b>Note:</b> Do not specify the <code>-m</code> switch when running an <code>sdcli</code> command on the local host.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| - n    | <p>The name and description of an MSCS cluster resource group to be created as part of the associated command.</p> <p>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.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| - np   | <p>The IP address and port of the network portal on the iSCSI connection target.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| - p    | <p>The storage system path to the location of the LUN on the storage system. This string takes the following form: <code>storagesystemname:/vol/volname/[qtree]/lun</code> for a storage system path.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| - z    | <p>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.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**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.

### **sysconfig list**

You can use `sysconfig list` to display SnapDrive for Windows configuration information for your host.

```
sdcli sysconfig list
```

#### **Description**

You can use `sysconfig list` to display SnapDrive for Windows configuration information for your host.

**Note:** When you are performing this operation in a Microsoft Cluster, `sysconfig list` only works if SnapDrive for Windows is present on each node of the cluster.

## 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 NetApp 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.

#### -i

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
```

### 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.

**-an**

Specifies the application on which you want to create the new dataset.

**-av**

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

#### -dn

Specifies the name of the dataset you want to restore.

#### -bv

Specifies the version of the dataset backup from which you want to restore. You can specify either the backup version or the backup ID.

#### -i

Specifies the ID of the dataset backup from which you want to restore. You can specify either the backup version or the backup ID.

#### SAP

Specifies the source-appointed list for the dataset backup from which you want to restore.

#### -S

Specifies 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 for Windows license operations.

### license set

You can use `license set` to provide the license key for the SnapDrive for Windows module you specify.

```
sdcli license set
-module ModuleName
-key LicenseKey
```

**Description**

You can use `license set` to provide the license key for the SnapDrive for Windows module you specify.

**Parameters****-module**

Specifies the SnapDrive for Windows module you want to license.

**-key**

Provides your licence key for the SnapDrive for Windows module you want to license.

**Example**

```
sdcli license set
-module LPSM
-key ABCDEFGHIJKLMN
```

## license list

You can use `license list` to display all the SnapDrive for Windows licenses you have installed.

```
sdcli license list
```

### Description

You can use `license list` to display all the SnapDrive for Windows licenses you have installed.

## license remove

The `license remove` command enables you to delete a SnapDrive for Windows license.

```
sdcli license remove  
-module ModuleName
```

### Description

You can use the `license remove` command to delete a SnapDrive for Windows license.

### Parameters

#### **-module**

Specifies the SnapDrive for Windows module for which you want to delete the license.

### Example

```
sdcli license remove  
-module LPSM
```

## Initiator group management commands

The `sdcli` utility provides command-line support for initiator group management.

### igroup list

You can use the `igroup list` command to display all the igroups on a storage system that have initiators on your local host or, if specified, a remote host.

```
sdcli igroup list
-f StorageSystem
[-m MachineName]
```

#### Description

You can use the `igroup list` command to display all the igroups on a storage system that have initiators on your local host or, if specified, a remote host.

#### Parameters

**-f**

Specifies the storage system name or IP address for which you want to list igroups.

**-m**

Specifies the remote host for which you want to list igroups. When you do not specify the `-m` parameter, the default is to specify igroups for your local host.

#### Examples

The following command displays the `igroup list` for the storage system with the IP address 172.17.167.45:

```
sdcli igroup list
-f 172.17.167.45
```

The following command displays the `igroup list` for the storage system `v34filer`, which has initiators on the remote host `server3`:

```
sdcli igroup list
-m server3
-f v34filer
```

## igroup create

`igroup create` enables you to create a new igroup.

### Syntax

```
sdcli igroup create  
[-m MachineName]  
-f StorageSystem  
-I NodeName InitiatorName  
-ig igroupName
```

### Description

`igroup create` enables you to create a new igroup.

### Parameters

#### **-m**

Optionally specifies the remote host on which you want to create an igroup. When you do not specify the `-m` parameter, the default is to specify igroups for your local host.

#### **-f**

Specifies the storage system on which you are creating your new igroup.

#### **-I**

Lists the machine name and initiator name in pairs. 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

The following command creates a new igroup called `v3group1`, on a storage system called `v34filer`, for two initiators, `10:00:00:00:c9:48:c9:5d` and `10:00:00:00:c9:48:c9:5e`, on a host called `server3`:

```
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
```

## igroup rename

The `igroup rename` command enables you to rename an existing igroup.

```
sdcli igroup rename  
[-m MachineName]  
-f StorageSystem  
-ig groupName  
-igNew groupNameNewName
```

### Description

`igroup rename` enables you to rename an existing igroup.

### Parameters

#### **-m**

Optionally specifies the remote host for which you want to list igroups. When you do not specify the `-m` parameter, the default is to specify igroups for your local host.

#### **-f**

Specifies the storage system on which the igroup you are renaming is located.

#### **-ig**

Specifies the name of the existing igroup you are renaming.

#### **-igNew**

Specifies the new name for your existing igroup.

### Example

The following command renames an igroup from `v3group1` to `v3group1fc` on a storage system with the IP address `172.17.167.45`:

```
sdcli igroup rename  
-f 172.17.167.45  
-ig v3group1  
-igNew v3group1fc
```

## igroup modify

`igroup modify` enables you edit an existing igroup.

### Syntax

```
sdcli igroup modify  
-m machine_name  
-f storage_system  
-ig igroupName
```

### Description

In Windows Server 2012 and 2012 R2 environments, you can use `sdcli igroup modify` to update your igroup with both sets of initiators for virtual Fibre Channel (vFC) environments. This new SDCLI command enables you to perform live migration operations in environments where vFC is enabled.

### Parameters

#### **-m**

Optionally specifies the remote host for which you want to list igroups. When you do not specify the `-m` parameter, the default is to specify igroups for your local host.

#### **-f**

Specifies the storage system on which the igroup you are modifying is located.

#### **-ig**

Specifies the name of the igroup you want to modify.

## igroup delete

The `igroup delete` command enables you to delete an existing igroup if there is no LUN mapped to it.

```
sdcli igroup delete  
[-m MachineName]  
-f StorageSystem  
-ig groupName
```

### Description

`igroup delete` enables you to delete an existing igroup if there is no LUN mapped to it.

**Parameters****-m**

Optionally specifies the remote host on which the igroup you want to delete is located. When you do not specify the **-m** parameter, the default is to specify igroups for your local host.

**-f**

Specifies the name of the storage system on which the igroup you want to delete is located.

**-ig**

Specifies the name of the igroup you want to delete.

**Example**

The following command deletes the igroup named v3group1fc from the storage system with the IP address 172.17.167.45:

```
sdcli igroup delete
-f 172.17.167.45
-ig v3group1fc
```

## Fractional space reservation monitoring commands

The `sdcli` utility provides command-line support for fractional space reservation monitoring.

### spacemon list

You can use `spacemon get` to display the space reservation monitoring settings for the specified host.

```
sdcli spacemon list
{-m MachineName}
```

**Description**

You can use `spacemon get` to display the space reservation monitoring settings for the specified host.

**Parameters****-m**

*MachineName* is the name of the machine on which you want to execute the command. If you do not specify a machine name, the command is executed on your local machine.

## spacemon set

You can use `spacemon set` to set the space reservation monitoring settings for your specified host.

```
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
```

### Description

You can use `spacemon set` to set the space reservation monitoring settings for your specified host specified host.

### Parameters

#### **-mi**

Specifies the frequency, in minutes, at which you want to monitor your available fractional space.

#### **-f**

Specifies the name of the storage system on which your LUNs reside.

#### **-vn**

Specifies the name of the volume you want to monitor.

#### **-m**

Specifies the name of the machine on which you want to set space reservation monitoring. If you do not specify a machine name, the default is your local host.

#### **-rap**

Specifies the point at which you want to be warned of a low space reservation condition.

#### **-roc**

Specifies 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.

#### **-css**

Specifies whether you want to monitor whether a Snapshot copy can be created.

**Example**

The following 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.

```
sdcli spacemon set
-mi 30
-f controller1
-vn testvol
-rap 90
-roc 500mb
-ccs true
```

**spacemon snap\_delta**

You can use the `spacemon snap_delta` command to display the rate of change between two Snapshot copies or between a Snapshot copy and the active file system of the storage system volume.

```
sdcli spacemon snap_delta
-f StorageSystem
-vn VolumeName
-s1 snapshot1
-s2 snapshot2
{-m MachineName}
```

**Description**

You can use the `spacemon snap_delta` command to display the rate of change between two Snapshot copies or between a Snapshot copy and the active file system of the storage system volume.

**Parameters****-f**

Specifies the name of the storage system on which the volume exists.

**-vn**

Specifies the name of the volume for which you want to display the rate of change between two Snapshot copies, or a Snapshot copy and the active file system of a storage system volume.

**-s1**

Indicates the name of the Snapshot copy you want to compare with either a second Snapshot copy or with the active file system.

**-s2**

Indicates the name of the second Snapshot copy or the active file system.

**-m**

Provides the name of the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

## spacemon snap\_reclaimable

You can use `spacemon snap_reclaimable` to display the space that can be reclaimed by deleting a Snapshot copy.

```
sdcli spacemon snap_reclaimable
-f StorageSystem
-vn VolumeName
-s snapshot
```

### Description

You can use `spacemon snap_reclaimable` to display the space that can be reclaimed by deleting a Snapshot copy.

### Parameters

**-f**

Specifies the name of the storage system on which the volume exists.

**-vn**

Specifies the name of the volume on which the Snapshot copy resides.

**-s**

Provides the name of the Snapshot copy for which you want to view reclaimable space.

## spacemon vol\_info

You can use `spacemon vol_info` to get information about fractional space reserved volumes.

```
sdcli spacemon vol_info
{-m MachineName}
```

### Description

You can use `spacemon vol_info` to get information about fractional space reserved volumes. The output for `spacemon vol_info` is displayed in XML format.

**Parameters****-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**spacemon delete**

The `spacemon delete` command enables you to delete the fractional space reservation monitor settings for the specified storage system volume.

```
sdcli spacemon delete  
-f StorageSystem  
-vn VolumeName  
{-m MachineName}
```

**Description**

The `spacemon delete` command enables you to delete the fractional space reservation monitor settings for the specified storage system volume.

**Parameters****-f**

Specifies is the name of the storage system on which the volume exists.

**-vn**

Indicates the name of the volume from which you want to delete fractional space reservation settings.

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

## Space reclamation commands

The `sdcli` utility provides command-line support for space reclamation operations.

### **spacereclaimer start**

The `spacereclaimer start` command enables you to start 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 (seven days).

#### **Example**

The following example starts the space reclamation process by specifying the CSV reparse point of a CSV disk:

```
sdcli spacereclaimer start  
-d C:\ClusterStorage\Volume8
```

## spacereclaim stop

The `spacereclaim stop` command enables you to stop the space reclamation process.

### Syntax

```
sdcli spacereclaim stop  
[-m MachineName]  
-d MountPoint
```

### Description

The `spacereclaim stop` command enables you to stop the space reclamation process.

### Parameters

**-m**

Specifies the *MachineName* on which you want to start the space reclamation process. If you do not specify a machine name, the default is your local system.

**-d**

Specifies the LUN mount point, volume name, or CSV reparse point on which you want to stop Space Reclaim.

### Example

The following example stops the space reclamation process on a CVS volume.

```
sdcli spacereclaim stop  
-d \\?\Volume{0944ca87-be05-45ad-8606-ba13ee7388a0}\
```

## spacereclaim analyze

You can use `spacereclaim analyze` to check whether space reclamation is needed for the LUN specified.

### Syntax

```
sdcli spacereclaim analyze  
[-m MachineName]  
-d MountPoint
```

## Description

You can use `spacereclai mer analyze` to check whether space reclamation is needed for the LUN specified.

## Parameters

**-m**

Specifies the *MachineName* on which you want to start the space reclamation process. If you do not specify a machine name, the default is your local system.

**-d**

Specifies the LUN mount point, volume name, or CSV reparse point on which you want to analyze.

## spacereclai mer status

You can use `spacereclai mer status` to display the space reclamation status for the LUN specified.

## Syntax

```
sdcli spacereclai mer status  
[-m MachineName]  
[-D MountPointList]
```

## Description

You can use `spacereclai mer status` to display 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.

## Parameters

**-m**

Specifies the *MachineName* on which you want to start the space reclamation process. If you do not specify a machine name, the default is your local system.

**-D**

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.

## Virtual Storage Console commands

The `sdcli` utility provides command-line support for managing communication between SnapDrive for Windows and Virtual Storage Console.

### **vsc\_config list**

You can use `vsc_config list` to display Virtual Storage Console configuration settings.

#### **Syntax**

```
sdcli vsc_config list
```

#### **Description**

You can use `vsc_config list` to display Virtual Storage Console configuration settings.

### **vsc\_config set**

You can use `vsc_config set` to set the Virtual Storage Console server IP address and port.

#### **Syntax**

```
sdcli vsc_config set  
-host host  
[-vscport VSC_port_number]  
[-port SnapDrive port]  
[-m MachineName]
```

#### **Description**

You can use `vsc_config set` to set the Virtual Storage Console server IP address and port.

#### **Parameters**

##### **-host**

Specifies the server name or IP address of the host where Virtual Storage Console is registered to the vSphere Client.

##### **-vscport**

Specifies the service port number with which you communicate with 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.

**-m**

Specifies the *MachineName* on which you want to configure Virtual Storage Console settings. If you do not specify a machine name, the default is your local system.

## vsc\_config delete

You can use `vsc_config delete` to delete the specified Virtual Storage Console server entry.

### Syntax

```
sdcli vsc_config delete  
[-m MachineName]  
[-port SnapDrive port]
```

### Description

You can use `vsc_config delete` to delete the specified Virtual Storage Console server entry.

### Parameters

**-m**

Specifies the *MachineName* from which you want to delete the Virtual Storage Console connection credentials. If you do not specify a machine name, the default is your local system.

**-port**

Specifies the Web service port number you use to communicate with SnapDrive. The default port is 808.

## Preferred IP address commands

The `sdcli` utility provides command-line support for managing preferred storage system IP addresses.

### preferredIP set

You can use `preferredIP set` to set a SnapDrive preferred IP address for the specified storage system management traffic.

#### Syntax

```
sdcli preferredIP set  
-f StorageSystem  
-IP PreferredIPAddress
```

#### Description

You can use `preferredIP set` to set a SnapDrive preferred IP address for the specified storage system management traffic.

#### Parameters

**-f**

Specifies the storage system for which you want to set a preferred IP address.

**-IP**

Specifies the IP address you want to use for storage system management traffic.

#### Example

The following example sets the SnapDrive preferred IP address for the storage system named `Storage1` to `172.28.53.94`:

```
sdcli preferredIP set  
-f Storage1  
-IP 172.18.53.94
```

## preferredIP list

The `preferredIP list` command enables you to display the storage system names and IP addresses that you set as the preferred IP addresses for SnapDrive to use for management traffic.

### Syntax

```
sdcli preferredIP list
```

### Description

The `preferredIP list` command enables you to display the storage system names and IP addresses that you set as the preferred IP addresses for SnapDrive to use for management traffic.

## preferredIP delete

You can use `preferredIP delete` to delete the preferred IP address for the specified storage system.

### Syntax

```
sdcli preferredIP delete  
-f StorageSystem
```

### Description

You can use `preferredIP delete` to delete the preferred IP address for the specified storage system.

### Parameters

**-f**

Specifies the storage system for which you want to delete the preferred IP address.

## iSCSI connection commands

The `sdcli` utility provides command-line support for managing connections to iSCSI targets.

### `iscsi_target disconnect`

You can use `iscsi_target disconnect` to disconnect the specified iSCSI initiator from the specified iSCSI target on all portals.

#### Syntax

```
sdcli iscsi_target disconnect  
-t TargetName
```

#### Description

You can use `iscsi_target disconnect` to disconnect the specified iSCSI initiator from the specified iSCSI target on all portals.

#### Parameters

`-t`

#### Example

The following example disconnects the specified iSCSI initiator from the specified iSCSI target:

```
sdcli iscsi_target disconnect  
-t iqn.1992.08.com.netapp:sn.33604307
```

### `iscsi_target list`

You can use `iscsi_target list` to display a list of all iSCSI targets. For each target, you can see all the portals through which the target is available or to which the target is connected.

#### Syntax

```
sdcli iscsi_target list  
{-f Storage_System | -i InitiatorPortName}
```

## Description

You can use `iscsi_target list` to display a list of all iSCSI targets. For each target, you can see all the portals through which the target is available or to which the target is connected.

## Parameters

**-f**

Displays all targets on the specified storage system

**-i**

### Example

The following example lists all the iSCSI targets on the Storage2 storage system, as well as all portals those targets are available through or connected to:

```
sdcli iscsi_target list
-f Storage2
```

## iSCSI initiator commands

The `sdcli` utility provides command-line support for managing iSCSI initiators.

### iscsi\_initiator list

You can use `iscsi_initiator list` to display a list of all iSCSI sessions on the specified machine.

## Syntax

```
sdcli iscsi_initiator list
{-m MachineName}
-s
```

## Description

You can use `iscsi_initiator list` to display a list of all iSCSI sessions on the specified machine.

## Parameters

**-m**

Specifies the name of the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-s**

Indicates that you want to list the iSCSI sessions.

**Example**

## iscsi\_initiator establish\_session

You can use `iscsi_initiator establish_session` to establish a session with a target using the specified host bus adapter (HBA).

### Syntax

```
sdcli iscsi_initiator establish_session
{-m MachineName}
{-h HBA_ID}
{-hp HBA_Portal_ID}
-t TargetName
-np IPAddress IPPort
{-c CHAPName CHAPPassword}
```

### Description

You can use `iscsi_initiator establish_session` to establish a session with a target using the specified HBA.

### Parameters

**-m**

Specifies the name of the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-h**

Establishes the iSCSI session. You can obtain the HBA ID by using `sysconfig list`.

**-hp**

Specifies the portal on the iSCSI HBA you want to use to establish the iSCSI session. You can obtain the HBA portal ID using `sysconfig list`.

**-t**

Specifies the name of the iSCSI target.

**-np**

Specifies the IP address and IP port of the network portal on the target. You can obtain the IP port by using `iscsi_initiator list`.

**-c**

Specifies the CHAP name and password. The parameter is required if the target requires CHAP authentication of the initiator.

### Example

The following example establishes an iSCSI session with the specified target using the specified HBA ID:

```
sdcli iscsi_initiator establish_session
-h 0
-t iqn.1992-8.com.netapp:maya
-np 172.18.53.94 3260
```

## iscsi\_initiator terminate\_session

You can use `iscsi_initiator terminate_session` to terminate your iSCSI initiator session.

### Syntax

```
sdcli iscsi_initiator terminate_session
{-m MachineName}
-s Session_ID
```

### Description

You can use `iscsi_initiator terminate_session` to terminate your iSCSI initiator session.

### Parameters

**-m**

Specifies the name of the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-s**

Specifies the session ID for the session you want to terminate.

### Example

The following example terminates the specified iSCSI session on your local machine:

```
sdcli iscsi_initiator terminate_session
-s 0xfffffffff868589cc-0x4000013700000006
```

## LUN commands

The `sdcli` utility provides command-line support for managing LUNs in SnapDrive.

### disk create

You can use `disk create` to create a new LUN.

#### Syntax

```
sdcli disk create
[-m MachineName]
[-p LUNpath | -d MountPoint
-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
-thin
```

#### Description

You can use `disk create` to create a new LUN.

#### Parameters

##### -m

Specifies the name of the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

##### -p

Specifies the path of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

##### -d

Specifies the mount point of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

##### -z

Specifies the drive size. Allowed unit specification are MB, GB, or TB; the default is MB.

##### -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.

**-I**

Specifies a list of machine node name and initiator name, in pairs.

**-IG**

Specifies machine node name and the igroup name, in pairs.

Specifies the existing igroups. LUNs are mapped to these Initiator groups. One pair required for dedicated disks. Two or more pairs (at least one pair for each node in the cluster) are required for shared disks. In case of FC RDM shared disks, provide the igroup having initiators from all the nodes and specify the same igroup name for all the nodes.

**-dtype**

Specifies the type of disk you want to create. Valid values are shared and dedicated.

**-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.

**-port**

Specifies the Web Service port number used to communicate with SnapDrive. The default is 808.

**-passthrough**

Enables you to dynamically add a pass-through disk. This option is applicable only when a Hyper-V virtual machine is present.

**-thin**

Specifies that you want to create a thinly provisioned LUN. This parameter is set to no by default.

**Example 1**

The following example creates a dedicated, 1-GB LUN named `mktnng.lun` in the `storage2` volume named `sd_vds_only` and then connects this LUN to the host as drive `R`:

```
sdcli disk create
-dtype dedicated
-z 1024
-p storage2:/vol/sd_vds_only/mktnng.lun
-d R
-I host3 10:00:00:00:C9:2B:FD:12
```

**Example 2**

The following example creates a shared, 4-GB LUN on `host4` (the local machine running the `SDCLI` command) and maps it to drive `R`, using a pair of initiators:

```
sdcli disk create
-p \\133.25.61.62\sd_vds_only\mktnng.lun
-d r
-z 4096
-dtype shared
-e "mktnng"
-I host4 10:00:00:00:C9:2B:FD:12 host4 10:00:00:00:C9:2B:FD:11
host5 10:00:00:00:C9:2B:FC:12 host5 10:00:00:00:C9:2B:FC:11
```

**Example 3**

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 "datastore1"
```

## disk connect

You can use `disk connect` to connect 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
-cn VolumeCloneName
```

### Description

You can use `disk connect` to connect a LUN to a host by mapping the LUN to a Windows drive letter.

#### 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.

### Parameters

#### -m

Specifies the name of the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

#### -p

Specifies the path of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

#### -d

Specifies the mount point of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-I**

Specifies a list of machine node name and initiator name, in pairs.

**-IG**

Specifies machine node name and the igroup name, in pairs.

Specifies the existing igroups. LUNs are mapped to these Initiator groups. One pair required for dedicated disks. Two or more pairs (at least one pair for each node in the cluster) are required for shared disks. In case of FC RDM shared disks, provide the igroup having initiators from all the nodes and specify the same igroup name for all the nodes.

**-dtype**

Specifies the type of disk to which you want to connect. Valid values are shared and dedicated.

**-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**

Specifies an existing resource group name.

**-n**

Specifies a new resource group name and description.

**-csv**

Enables you to add a shared disk to a cluster shared volume.

**-c**

Specifies a cluster name. This parameter is required only if the LUN is SCSI reserved.

**-port**

Specifies the Web Service port number used to communicate with SnapDrive. The default is 808.

**-passthrough**

Enables you to dynamically connect a pass-through disk. This option is applicable only when a Hyper-V virtual machine is present.

**-cn**

Specifies the name of the volume clone you want to mount.

**Example 1**

The following example connects a LUN in 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 10:00:00:00:C9:2B:FD:1B host3 10:00:00:00:C9:2B:FD:1C
host4 10:00:00:00:C9:2B:FD:12 host4 10:00:00:00:C9:2B:FD:11
-e "tech_mktng"
-c "mktng"
```

**Example 2**

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 "datastore1"
```

**Example 3**

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.netapp.com
CLAB-A9-7 iqn.1991-05.com.microsoft:clab-a9-7.sddev.netapp.com
-dtype shared
-csv
```

## disk delete

You can use `disk delete` to delete a LUN. The LUN must be mapped to a Windows drive letter or mount point for the command to succeed.

### Syntax

```
sdcli disk delete
[-m MachineName]
{-p LUNpath | -d MountPoint}
```

### Description

You can use `disk delete` to delete a LUN. The LUN must be mapped to a Windows drive letter or mount point for the command to succeed.

### Parameters

**-m**

Specifies the name of the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-p**

Specifies the path of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-d**

Specifies the mount point of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

#### Example 1

The following example deletes the LUN `mktnng.lun` from the `sd_vds_only` volume on the storage system identified by the IP address `133.25.61.62`:

```
sdcli disk delete
-p \\133.25.61.62\sd_vds_only\mktnng.lun
```

#### Example 2

The following example specifies the CSV volume name to delete a CSV disk:

```
sdcli disk delete
-d \\?\Volume{f1816466-f1d8-4b96-b547-2ce12415aee4}\
```

**Example 3**

The following example specifies the CSV repare point to delete a CSV disk:

```
sdcli disk delete
-d C:\ClusterStorage\Volume8
```

**disk disconnect**

You can use `disk disconnect` to disconnect a LUN from the host. You must have the LUN mapped to a Windows drive letter for this command to succeed.

**Syntax**

```
sdcli disk disconnect
[-m MachineName]
{-p LUNpath | -d MountPoint}
[-f]
```

**Description**

You can use `disk disconnect` to disconnect a LUN from the host. You must have the LUN mapped to a Windows drive letter for this command to succeed.

You must make sure that the LUN you are disconnecting is not monitored by Windows Performance Monitor (perfmon).

**Parameters**

**-m**

Specifies the name of the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-p**

Specifies the path of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-d**

Specifies the mount point of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-f**

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 1**

The following example disconnects the LUN mapped to the drive letter Z on the SnapDrive host running the `sdcli` command:

```
sdcli disk disconnect
-d
-z
```

**Example 2**

The following example forces disconnection of the LUN `mktng.lun`, which is in the `sd_vds_only` volume on `storage2`. Because you are using the `-z` switch, all open files in the LUN might be lost or corrupted.

```
sdcli disk disconnect
-p \\storage2\sd_vds_only\mktng.lun
-f
```

**disk resize**

You can use `disk resize` to increase or decrease the size of your disk. You can specify any size, as long as that figure falls within the SnapDrive-specified minimum and maximum values.

**Syntax**

```
sdcli disk resize
[-m MachineName]
{-p LUNpath | UNCPath | -d MountPoint}
-z DriveSizeIncrement
[-s SnapshotName]
[-x LunSnapshot]
```

**Description**

You can use `disk resize` to increase or decrease the size of your disk. You can specify any size, as long as that figure falls within the SnapDrive-specified minimum and maximum values.

**Parameters**

**-m**

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**

Specifies the LUN path including the storage system name or the UNC path to the location of the LUN file on the storage system.

**-d**

Specifies the mount point volume name or CSV repare point of the LUN.

**-z**

Specifies 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 your disk increase or decrease size in increments of MB, GB, or TB. If you do not provide a unit of measurement, the default is MB.

**-s**

Specifies the name of the Snapshot copy you want to use when you are making a Snapshot copy for restore operations.

**-x**

Specifies that you want to create a Snapshot copy 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.

### Example 1

The following example shrinks the size of the disk mapped to G by 500 MB:

```
sdcli disk resize  
-d G:  
-z -500
```

### Example 2

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 resize  
-d P:  
-z 1GB
```

**Example 3**

The following example increases a CSV disk by 100 MB using the CSV volume name:

```
sdcli disk resize
-d \\?\Volume{f1816466-f1d8-4b96-b547-2ce12415aee4}\
-z +100
```

**disk expand**

You can use `disk expand` to expand your disk. You can specify any size, as long as that figure falls within the minimum and maximum values required by SnapDrive.

**Syntax**

```
sdcli disk expand
[-m MachineName]
{-p LUNpath | -d MountPoint}
-z DriveSizeIncrement
[-rs Reserve Snapshot copy space {y|n}]
```

**Description**

You can use `disk expand` to expand your disk. You can specify any size, as long as that figure falls within the minimum and maximum values required by SnapDrive.

**Parameters****-m**

Specifies the name of the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-p**

Specifies the path of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-d**

Specifies the mount point of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-z**

Specifies the amount by which you want to expand the disk. You can expand your disk in increments of MB, GB, or TB. If you do not specify an increment size, 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 1

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
```

### Example 2

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

```
sdcli disk add_initiator
[-m MachineName]
{-p LUNpath | -d MountPoint}
-i InitiatorPortName
```

### Description

You can use the `disk add_initiator` command to add a new initiator to a LUN.

### Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-p**

Specifies the path of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-d**

Specifies the mount point of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-i**

Specifies the port name for the initiator you want to remove.

### 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
```

### Example 2

The following example adds an initiator to a LUN located at the path `\\sdwatf2\sdwatf2_vol1\sdwath2_EEE.lun` on the SnapDrive host from which the `sdcli` command was executed:

```
sdcli disk add_initiator
-p \\sdwatf2\sdwatf2_vol1\sdwath2_EEE.lun
-i 21:00:00:e0:8b:85:19:ba
```

### Example 3

The following example specifies the CSV repare 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

You can use `disk remove_initiator` to remove an initiator from the specified LUN.

### Syntax

```
sdcli disk remove_initiator  
[-m MachineName]  
{-p LUNpath | -d MountPoint}  
-i InitiatorPortName
```

### Description

You can use `disk remove_initiator` to remove an initiator from the specified LUN.

### Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-p**

Specifies the path of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-d**

Specifies the mount point of the LUN containing the initiator you want to remove. You can provide either a LUN path or a mount point.

**-i**

Specifies the port name for the initiator you want to remove.

#### Example 1

The following example removes an initiator from a LUN mounted on drive letter E on the SnapDrive host running your SDCLI command:

```
sdcli disk remove_initiator  
-d E  
-i 21:00:00:e0:8b:85:19:ba
```

**Example 2**

The following example removes an initiator from a LUN located at the path \\\sdwatf2\\sdwatf2\_vol1\\sdwath2\_EEE.lun on the SnapDrive host from which your SDCLI command was executed:

```
sdcli disk remove_initiator
-p \\\sdwatf2\\sdwatf2_vol1\\sdwath2_EEE.lun
-i 21:00:00:e0:8b:85:19:ba
```

**Example 3**

The following example specifies the CSV volume name \\?\Volume{24cd94c5-d201-4474-9246-2ad491b155ea}\ to remove an initiator from a CSV disk:

```
sdcli disk remove_initiator
-d \\?\Volume{24cd94c5-d201-4474-9246-2ad491b155ea}\
-i 50:0a:09:80:85:f4:69:37
```

**disk list**

You can use `disk list` to display a list of all the LUNs connected to the host.

**Syntax**

```
sdcli disk list
[-m MachineName]
```

**Description**

You can use `disk list` to display a list of all the LUNs connected to the host. You can also view the following information for each LUN:

- LUN path (storage system name, share name, virtual disk file name, and possibly 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)
- 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.

### Example

The following example lists all the SnapDrive LUNs mapped to drive letters on your local host:

```
sdcli disk list
```

## disk add\_mount

disk add\_mount enables you to add a volume mount point.

### Syntax

```
sdcli disk add_mount
{-m MachineName}
-vn Volume_Name
-mp Volume_Mount_Point
{-create_folder}
```

### Description

disk add\_mount enables you to add a volume mount point.

### Parameters

#### -m

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

#### -vn

Specifies 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.

#### -mp

Specifies the location to which 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

The following example creates mount point

**G:\mount\_vol1**

on a volume named

**\\?\Volume{db6160d8-1f14-11da-8ef3-000d5671229b}**

, and also indicates that you want to create a folder for your new mount point:

```
sdcli disk add_mount
-vn \\?\Volume{db6160d8-1f14-11da-8ef3-000d5671229b}
```

```
-mp G:\mount_vol1  
-create_folder
```

## disk remove\_mount

You can use `disk remove_mount` to remove a volume mount point or drive letter. This command does not delete the folder that was created at the time you added the volume mount point. After you remove a mount point, an empty folder remains with the same name as the mount point you removed.

### Syntax

```
sdcli disk remove_mount  
{-m MachineName}  
-vn VolumeName  
-mp Volume_Mount_Point
```

### Description

You can use `disk remove_mount` to remove a volume mount point or drive letter.

This command does not delete the folder that was created at the time you added the volume mount point. After you remove a mount point, an empty folder remains with the same name as the mount point you removed.

### Parameters

#### **-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

#### **-vn**

Specifies the volume for which you want to delete the mount point.

#### **-mp**

Indicates the mount point you want to delete.

**Example****disk rename\_flexclone**

The `disk rename_flexclone` command enables you to rename the FlexClone volume from the default.

**Syntax**

```
sdcli disk rename_flexclone  
-d MountPoint  
-n NewName
```

**Description**

The `disk rename_flexclone` command enables you to rename the FlexClone volume from the default.

**Note:** You cannot rename VMDK volumes in SnapDrive for Windows because Virtual Storage Console is not supported.

**Parameters****-d**

Specifies the name of the LUN in the FlexClone volume.

**-n**

Specifies the new name that you assign for the FlexClone volume. After you rename the FlexClone volume, you can find that name the output from the `disk list` command.

### Example

## Snapshot copy commands

The `sdcli` utility provides command-line support for managing Snapshot copies of SnapDrive LUNs.

### snap create

You can use `snap create` to create a new Snapshot copy of the LUNs you specify on the SnapDrive system you designate.

#### Syntax

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

#### Description

You can use `snap create` to create a new Snapshot copy of the LUNs you specify on the SnapDrive system you designate.

#### Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-s**

Names your new Snapshot copy

**-D**

Specifies the mount point, volume name, or CSV reparse points to which you want to save your new Snapshot copy.

**-x**

Indicates that you want 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.

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 1

The following 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 Jun_13_03
-D j k l
```

### Example 2

The following example creates a Snapshot copy named csvLUN11\_051109 on the CSV reparse point C:\ClusterStorage\Volume8:

```
sdcli snap create
-s csvLun11_051109
-D C:\ClusterStorage\Volume8
```

## snap delete

You can use `snap delete` to delete an existing Snapshot copy.

### Syntax

```
sdcli snap delete
[-m MachineName]
-D MountPointList [. . .]
-s SnapshotName
```

### Description

You can use `snap delete` to delete an existing Snapshot copy.

You must ensure that the LUN whose Snapshot copy you are deleting is not being monitored by Windows Performance Monitor (perfmon).

### Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-d**

Provides a list of mount points at which the Snapshot copy you want to delete is located.

**-s**

Specifies the Snapshot copy that you want to delete.

### Example 1

The following example deletes the Snapshot copy named Jun\_14\_09 that is associated with the LUN mapped to drive K on the local host:

```
sdcli snap delete
-D k
-s Jun_14_09
```

### Example 2

The following example deletes the Snapshot copy named csvLun11\_051109\_new from the CSV volume \\?\Volume{239889f5-3a36-4993-b957-0a85f56cab45}\:

```
sdcli snap delete
-d \\?\Volume{239889f5-3a36-4993-b957-0a85f56cab45}\
-s csvLun11_051109_new
```

## snap list

You can use `snap list` to list all the Snapshot copies that exist for the specified LUN.

### Syntax

```
sdcli snap list
[-m MachineName]
-d MountPoint
```

### Description

You can use `snap list` to list all the Snapshot copies that exist for the specified LUN.

### Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-d**

Specifies a drive letter, mount point path, volume name, or CSV reparse point for the LUN for which you want to view Snapshot copies.

**Example**

The following example displays all the Snapshot copies that exist for the volume containing the LUN mapped to drive J on the local host:

```
sdcli snap list
-d j
```

## snap mirror\_list

You can use `snap mirror_list` to display the SnapMirror relationships and other information associated with your SnapMirror source volume.

### Syntax

```
sdcli snap mirror_list
-d MountPoint
```

### Description

You can use `snap mirror_list` to display the SnapMirror relationships associated with your SnapMirror source volume. This information includes 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 relationship.

### Parameters

**-d**

Specifies a drive letter, mount point path, volume name, or CSV reparse point of the SnapMirror source drive.

**Example**

The following 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.

```
C:\Program Files\NetApp\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.

## snap mount

You can use `snap mount` to mount a Snapshot copy of a LUN. Snapshot copies are always mounted in read/write mode.

### Syntax

```
sdcli snap mount
[-m MachineName]
[-r LiveMachineName]
-k LiveMountPoint
-s SnapshotName
-d MountPoint
-cn VolumeCloneName
```

### Description

You can use `snap mount` to mount a Snapshot copy of a LUN. Snapshot copies are always mounted in read/write mode.

### Parameters

#### -m

Specifies machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

#### -r

Specifies the name of the host connected to the LUN in the active file system. When left unspecified, -r defaults to the local host.

When you are 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.

#### -k

Specifies the drive letter, mount point, volume name, or CSV reparse point assigned to the LUN in the active file system.

**-s**

Indicates the name of the Snapshot copy you want to mount.

**-d**

Provides the mount point where your Snapshot copy is located.

**-cn**

Specifies the name of the volume clone you want to mount.

### Example

The following 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 drive J on host3.

```
sdcli snap mount
-r host3
-k j
-s Jun_14_09
-d t
```

## snap rename

The `snap rename` command enables you to change the name of an existing Snapshot copy.

### Syntax

```
sdcli snap rename
[-m MachineName]
-d MountPoint
-o OldSnapshotName
-n NewSnapshotName
```

### Description

The `snap rename` command enables you to change the name of an existing Snapshot copy.

### Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-d**

Specifies drive letter, mount point path, volume name, or CSV reparse point on which the Snapshot copy you want to rename is located.

**-o**

Indicates the old name of your Snapshot copy.

**-n**

Provides the new name for your Snapshot copy.

### Example 1

The following example changes the name of the June\_14\_09 Snapshot copy associated with the J drive to last\_known\_good:

```
sdcli snap rename
-d j
-o Jun_14_09
-n last_known_good
```

### Example 2

The following example changes the name of the csvLun11\_051109 Snapshot copy associated with CSV reparse point C:\ClusterStorage\Volume8 to csvLun11\_051109\_new:

```
sdcli snap rename
-d C:\ClusterStorage\Volume8
-o csvLun11_051109
-n csvLun11_051109_new
```

## snap restore

You can use `snap restore` to replace your current LUN image in the active file system with the point-in-time image captured by a Snapshot copy that you specify. You can also perform file-level restore operations, restoring only the files you specify.

### Syntax

```
sdcli snap restore
[-m MachineName]
-d MountPoint
-flr
[-copy]
-s SnapshotName
-files filepath[...]
```

## Description

You can use `snap restore` to replace your current LUN image in the active file system with the point-in-time image captured by a Snapshot copy that you specify. You can also perform file-level restore operations, restoring only the files you specify.

You must make sure that the LUN you are disconnecting is not being monitored by Windows Performance Monitor (perfmon).

## Parameters

### **-m**

Specifies machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

### **-d**

Specifies a drive letter, mount point path, volume name, or CSV reparse point to which you want to restore your Snapshot or individual files.

### **-flr**

Specifies that you are performing a file-level restore operation.

### **-copy**

Forces a copy-restore operation if other restore options fail. This is an option setting.

### **-s**

Specifies the Snapshot copy from which you want to perform your restore operation.

### **-files**

Lists the individual files you want to restore, including their paths and drive letters or volume mount points.

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 are not recursively restored.

File-level wildcard entries are not supported; therefore, you cannot combine the hyphen and asterisk with a file name at the end of the path you want to restore.

**Example 1**

The following example restores the LUN mapped to drive L on your local host to its state when the Snapshot copy named Jun\_14\_09 was made:

```
sdcli snap restore
-d l
-s Jun_14_09
```

**Example 2**

The following example restores a single file called v1\_d.vhd in the volume mount point c:\mnt1\, restores 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:

```
sdcli snap restore
-flr
-s st_10_18_2009
-files c:\mnt1\v1\v1_d.vhd f:\vm1\vm1.vhd "f:\vm1\Virtual Machines
\"
```

**snap unmount**

You can use `snap unmount` to disconnect a Snapshot copy of a LUN that is mounted as a LUN.

**Syntax**

```
sdcli snap unmount
[-m MachineName]
-d MountPoint
[-f]
```

**Description**

You can use `snap unmount` to disconnect a Snapshot copy of a LUN that is mounted as a LUN.

You must ensure that the LUN whose Snapshot copy you are disconnecting is not being monitored by Windows Performance Monitor (perfmon).

**Attention:** If you unmount a LUN on a FlexClone volume that SnapDrive 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 automatically deleting a 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 `new_sdwvol_0`, and not `sdw_cl_myvol_20`. If you rename only the last integers in the volume name, SnapDrive still recognizes that it created that volume and deletes the volume when you disconnect the last LUN.

## Parameters

**-m**

Specifies machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-d**

Specifies the mount point from which you want to unmount your Snapshot copy.

**-f**

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 1

The following example disconnects the Snapshot copy mapped to drive K on the local host:

```
sdcli snap unmount
-d k
```

### Example 2

The following example forces disconnection of the Snapshot copy mapped to the K drive on the local host:

```
sdcli snap unmount
-d k
-f
```

## snap update\_mirror

You can use `snap update_mirror` to update the LUN relationship to a SnapMirror destination volume residing on the same or a different storage system.

## Syntax

```
sdcli snap update_mirror
[-m MachineName]
-d MountPoint
```

## Description

You can use `snap update_mirror` to update the LUN relationship to a SnapMirror destination volume residing on the same or a different storage system.

## Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-d**

Specifies the drive letter or mount point of the LUN for which you want to update the SnapMirror destination.

### Example

The following 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.

```
sdcli snap update_mirror
-d l
```

## snap restore\_volume\_check

You can use `snap restore_volume_check` to verify whether you can perform the restore operation on a specified volume.

## Syntax

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

## Description

You can use `snap restore_volume_check` to verify whether you can perform the restore operation on a specified volume.

## Parameters

**-f**

Specifies the name of the storage system on which the volume resides.

**-volume**

Indicates the name of the volume on which the restore operation is to be performed.

**-s**

Indicates the name of the Snapshot copy from which the volume is to be restored.

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**Example**

The following 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`:

```
sdcli snap restore_volume_check
-f clpubs-storage1
-volume vol3
-s my_snap
```

**snap restore\_volume**

You can use `snap restore_volume` to restore a storage system volume from a Snapshot copy you specify.

**Syntax**

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

**Description**

You can use `snap restore_volume` to restore a storage system volume from a Snapshot copy you specify.

**Parameters****-f**

Specifies the name of the storage system on which the volume resides.

**-volume**

Indicates the name of the volume on which the restore operation is to be performed.

**-s**

Indicates the name of the Snapshot copy from which you want to restore the volume.

**-force**

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

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**Example**

The following 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`:

```
sdcli snap restore_volume
-f clpubs-storage1
-volume vol3
-s my_snap
```

## SnapVault commands

The `sdcli` utility provides command-line support for SnapVault management using SnapDrive. You can only perform SnapVault management operations in 7-Mode environments.

### snapvault verify\_configuration

The `snapvault verify_configuration` command enables you to check the SnapVault configuration to ensure that it is correct.

#### Syntax

```
sdcli snapvault verify_configuration
[-m MachineName]
{-D MountPoint | -G GUIDlist}
```

#### Description

The `snapvault verify_configuration` command enables you to check the SnapVault configuration to ensure that it is correct.

**Parameters****-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.

**Example****snapvault snapshot\_rename**

The `snapvault snapshot_rename` command enables you to rename an existing Snapshot copy on a secondary system.

**Syntax**

```
sdcli snapvault snapshot_rename
[-m MachineName]
-o OldName
-n NewName
{-d MountPoint | -G GUIDlist}
```

**Description**

The `snapvault snapshot_rename` command enables you to rename an existing Snapshot copy on a secondary system.

**Parameters****-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-o**

Specifies the old name of the Snapshot copy that you want to change on the secondary system.

**-n**

Specifies the new name of the Snapshot copy on the secondary system. The new name must not yet exist.

**-d**

Specifies the mount point that identifies the disk on the primary system.

**-G**

Specifies a list of GUIDs of disks on the primary system.

### Example

## snapvault snapshot\_delete

You can use `sdcli snapvault snapshot_delete` to delete an existing Snapshot copy on a SnapVault secondary system.

### Syntax

```
sdcli snapvault snapshot_delete
[-m MachineName]
{-D MountPoint | -G GUIDlist}
-a ArchivalSnapshotName
```

### Description

You can use `sdcli snapvault snapshot_delete` to delete an existing Snapshot copy on a SnapVault secondary system.

### Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-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.

**-a**

Specifies the name of the Snapshot copy that you want to delete

### Example

## snapvault archive

You can use `snapvault archive` to archive a backup set to a secondary storage system.

### Syntax

```
sdcli snapvault archive
[-m MachineName]
[-force]
-a ArchivalSnapshotName
-DS MountPointandSnapshotList [...]
```

### Description

You can use `snapvault archive` to archive a backup set to a secondary storage system.

### Parameters

#### -m

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

#### -force

Forces the secondary Snapshot copy to be made, regardless of the possible failure of some qtrees updates.

#### -a

Specifies the name of the Snapshot copy on the secondary system. This Snapshot copy name must not already exist.

#### -DS

Specifies a list of mount points and Snapshot copies to be archived.

### Example

## snapvault relationship\_status

You can use `snapvault relationship_status` to display the relationship status of the primary storage system for a disk you specify.

### Syntax

```
sdcli snapvault relationship_status
[-m MachineName]
{-D MountPoint | -G GUIDlist}
```

`-G guidlist s`

## Description

You can use `snapvault relationship_status` to display the relationship status of a primary storage system for the disk you specify.

## Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-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.

### Example

## snapvault snap\_list

You can use `snapvault snap_list` to display the Snapshot copies on a volume you specify using a mount point or GUID on the SnapVault secondary system.

## Syntax

```
sdcli snapvault snap_list
[-m MachineName]
{-D MountPoint | -G GUIDlist}
```

## Description

You can use `snapvault snap_list` to display the Snapshot copies on a volume you specify using a mount point or GUID on the SnapVault secondary system.

## Parameter

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-D**

Specifies a list of mount points of disks on your primary system.

**-G**

Specifies a list of GUIDs of disks on the primary system.

#### Example

## OnCommand commands

The `sdcli` utility provides command-line support for managing OnCommand Manager credentials after SnapDrive has been installed.

### `oncommand_config list`

You can use `oncommand_config list` to display a list of already-configured OnCommand Unified Manager servers.

#### Syntax

```
sdcli oncommand_config list
```

#### Description

You can use `oncommand_config list` to display a list of already-configured OnCommand Unified Manager servers.

#### Example

### `oncommand_config set`

The `oncommand_config set` command enables you to set OnCommand Unified Manager server credentials.

#### Syntax

```
sdcli oncommand_config set  
-host Host  
-user Username  
-pwd Password  
[-port Port]
```

## Description

The `oncommand_config set` command enables you to set OnCommand Unified Manager server credentials.

## Parameters

### **-host**

Specifies the host name or IP address of the host running a OnCommand Unified Manager server.

### **-user**

Specifies the user name for the OnCommand Unified Manager server.

### **-pwd**

Specifies the password to be used for the OnCommand Unified Manager server.

### **-port**

Specifies a new TCP port. The default port 8088 is used if a new port is not specified.

### Example

## **oncommand\_config delete**

The `oncommand_config delete` command enables you to remove an OnCommand Unified Manager server from the SnapDrive Unified Manager server list.

## Syntax

```
sdcli oncommand_config delete  
-host Host
```

## Description

The `oncommand_config delete` command enables you to remove an OnCommand Unified Manager server from the SnapDrive Unified Manager server list.

## Parameters

### **-host**

Specifies the name or IP address of the host running the Unified Manager server you want to remove from the list.

### Example

## oncommand\_config rbaccache

You can use `oncommand_config rbaccache` to enable SnapDrive to cache RBAC operations for use when the OnCommand Unified Manager server is down and unavailable for fewer than 24 hours.

### Syntax

```
sdcli oncommand_config rbaccache  
[-m MachineName]  
-rc Enable| Disable
```

### Description

You can use `oncommand_config rbaccache` to enable SnapDrive to cache RBAC operations for use when the OnCommand Unified Manager server is down and unavailable for fewer than 24 hours.

### Parameters

#### -m

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

#### -rc

Enable or disable OnCommand Unified Manager RBAC caching. Valid values are enable and disable.

When `dfm_config rbaccache` is enabled, updates to the RBAC cache occur automatically whenever RBAC information changes on the storage system. If the OnCommand Unified Manager server is unavailable for more than 24 hours, SnapDrive no longer operates.

### Example

## Transport protocol commands

The SDCL utility provides command-line support for managing transport protocols used by SnapDrive.

### transport\_protocol list

You can use `transport_protocol list` to display the transport protocol configuration settings SnapDrive uses on the storage system.

#### Syntax

```
sdcli transport_protocol list  
[-m MachineName]
```

#### Description

You can use `transport_protocol list` to display the transport protocol configuration settings SnapDrive uses on the storage system.

#### Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

#### Example

The following example indicates that HTTP is the default transport protocol setting on the local SnapDrive system:

```
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.
```

## transport\_protocol set

You can use `transport_protocol set` to set or modify your 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]
```

### Description

You can use `transport_protocol set` to set or modify your transport protocol on the storage system.

When you are configuring your transport protocol settings in a clustered Data ONTAP environment, you should be aware of the following additional information:

- When you use clustered Data ONTAP 8.1 or later with SnapDrive, you must 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.

### Parameters

#### **-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

#### **-f**

Specifies the storage system name or IP address.

#### **-default**

Indicates that you want to use the default protocol.

#### **-type**

Specifies the protocol type to be used. Protocol type is either HTTP, HTTPS, or RPC. RPC is not supported in clustered Data ONTAP environments.

#### **-port**

Specifies the port number the protocol is to use. The default port for HTTPS is 443. The default port for HTTP is 80.

**-user**

Specifies the user with permission on the storage system. A user name is required if the protocol type is HTTP or HTTPS.

**-pwd**

Specifies the user password. A password is required if the protocol type is HTTP or HTTPS. You are prompted for a password if one is not specified.

**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.
```

**transport\_protocol delete**

You can use `transport_protocol delete` to delete a transport protocol from a storage system.

**Syntax**

```
sdcli transport_protocol delete
[-m MachineName]
-f StorageSystem | -default
```

**Description**

You can use `transport_protocol delete` to delete a transport protocol from a storage system.

**Parameters****-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-f**

Specifies the storage system name or IP address.

**-default**

Indicates that you want to use the default transport protocol 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.

### vsconfig list

You can use `vsconfig list` to display virtual server configuration settings.

#### Syntax

```
sdcli vsconfig list
```

#### Description

You can use `vsconfig list` to display virtual server configuration settings.

### vsconfig set

The `vsconfig set` command enables you to set the virtual server configuration.

#### Syntax

```
sdcli vsconfig set  
-ip IP Address  
-user User Name  
-pwd Password  
-m  
-port
```

#### Description

The `vsconfig set` command enables you to set the virtual server configuration.

You can change the IP address of a virtual server that you have previously configured during SnapDrive installation without reentering the user name and password. A valid IP address, user name, and password must already exist in the Windows registry; otherwise, an error message is displayed, indicating that the user credentials are not set.

#### Parameters

##### -ip

Specifies the IP address of the virtual server.

**-user**

Specifies the virtual machine user name.

**-pwd**

Specifies the password for the virtual machine.

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-port**

Specifies the Web service port number you use to communicate with SnapDrive. The default port is 808.

## vsconfig dslist

You can use `vsconfig dslist` to display datastores available on your ESX servers.

### Syntax

```
sdcli vsconfig dslist
-m MachineName
-port "Port Number"
-SharedDS
```

### Description

You can use `vsconfig dslist` to display datastores available on your ESX servers.

### Parameters

**-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

**-port**

Specifies the Web service port number you use to communicate with SnapDrive. The default port is 808.

**-SharedDS**

Specifies that you only want to list the shared datastores.

## vsconfig delete

You can use `vsconfig delete` to disable ESX server and vCenter Server settings.

### Syntax

```
sdcli vsconfig delete
-m machine_name
-port port_number
```

### Parameters

#### **-m**

Specifies the machine on which you want to execute the command. If no machine name is specified, the command is executed on the local machine.

#### **-port**

Specifies the Web service port number you use to communicate with SnapDrive. The default port is 808.

### Description

You can use `vsconfig delete` to disable ESX server and vCenter Server settings.

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

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.

### hyperv\_config list

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, to discover if they belong to Cluster Shared Volumes.

### Syntax

```
sdcli hyperv_config list
```

## Description

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, to discover if they belong to Cluster Shared Volumes.

### Example

## hyperv\_config set

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

```
sdcli hyperv_config set
-host hostname
-IP IP_address
[-port port_number]
```

## Description

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.

## Parameters

### -host

Specifies the Hyper-V parent node host name.

### -IP

Specifies the parent node IP address.

### -port

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
```

## hyperv\_config delete

The `hyperv_config delete` command enables you to delete a Hyper-V parent node configuration.

### Syntax

```
sdcli hyperv_config delete  
-host hostname
```

### Description

The `hyperv_config delete` command enables you to delete a Hyper-V parent node configuration.

### Parameters

#### **-host**

Specifies the name of the Hyper-V parent node from which you want to delete your configuration information.

## Clone split commands

You can use the `sdcli` utility to manage your clone splitting operations using the command line. Clone splitting operations enable you to split your clone volume from the parent volume and turn it into an independent FlexClone volume. If you want the FlexClone volume to have its own disk space, rather than using that of its parent, you can split it from its parent.

## clone\_split clone\_estimate

You can use `clone_split clone_estimate` to estimate the amount of disk space you have available before splitting a clone.

### Syntax

```
sdcli clone_split clone_estimate  
-D MountPoints
```

### Description

You can use `clone_split clone_estimate` to estimate the amount of disk space you have available before splitting a clone.

**Parameters****-D**

Specifies the clone volume mount points.

**clone\_split clone\_start**

You can use `clone_split clone_start` to split a clone volume from its parent volume and create an independent FlexVol volume.

**Syntax**

```
sdcli clone_split clone_start  
-D MountPoints
```

**Description**

You can use `clone_split clone_start` to split a clone volume from its parent volume and create an independent FlexVol volume.

**Parameters****-D**

Specifies the clone volume mount points.

**clone\_split clone\_status**

You can use `clone_split clone_status` to get information about the status of your clone splitting operation.

**Syntax**

```
sdcli clone_split clone_status  
-D MountPoints
```

**Description**

You can use `clone_split clone_status` to get information about the status of your clone splitting operation.

**Parameters****-D**

Specifies your clone volume mount points.

## **clone\_split clone\_stop**

You can use `clone_split clone_stop` to halt a clone splitting operation.

### **Syntax**

```
sdcli clone_split clone_stop  
-D MountPoints
```

### **Description**

You can use `clone_split clone_stop` to halt a clone splitting operation.

### **Parameters**

**-D**

Specifies your clone volume mount points.

# 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

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 NetApp Support Site. 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 NetApp Support Site.

### 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



## 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 NetApp Support Site.

## Copyright information

---

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S.

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. 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

---

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

## How to send your comments

---

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to [doccomments@netapp.com](mailto:doccomments@netapp.com). To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

# Index

7-Mode limitations

SnapVault [70](#)

7-Mode SAN environment

initiating SnapVault backup jobs in [71](#)

## A

access control

operations list and descriptions [81](#)

AccessControl.xml file

operations [81](#)

using to manage storage system access control [80](#)

adding a backup

using dataset backup\_add [99](#)

adding metadata to a dataset

using the dataset set\_metadata command [115](#)

adding objects to a dataset

using dataset add\_members [98](#)

assigning roles using RBAC [91](#)

automatic storage provisioning

using SnapDrive for Windows [13](#)

## B

backing up a dataset

using the dataset backup\_start command [104](#)

backup copy location

of datasets [108](#)

backup jobs

using SnapVault with SnapDrive in 7-Mode SAN [71](#)

backup retention policy information

obtaining using dataset get\_retention\_info [110](#)

backup sets

SnapVault [70](#)

backup version numbers

converting to timestamps [105](#)

backups

excluding independent VMDK disks [68](#)

best practices

for creating LUNs when using SnapDrive for Windows [25](#)

for using SnapDrive for Windows [16](#)

## C

changing dataset backup metadata [103](#)

CHAP authentication

about using with SnapDrive [19](#)

clone support

included in SnapDrive [11](#)

clone\_split clone\_estimate command

syntax for [186](#)

clone\_split clone\_status command

syntax for [187](#)

clone\_split clone\_stop command

syntax for [188](#)

clonesplit start command

syntax for [187](#)

cluster group

change owner [24](#)

verify owner [24](#)

Cluster Shared Volumes

LUN provisioning with [23](#)

SnapDrive for Windows support for [23](#)

Snapshot copy management with [23](#)

using with Hyper-V and Windows Server [23](#)

clustered Data ONTAP

creating disks in [31](#)

SnapVault operations supported in [71](#)

command

disk resize [148](#)

command-line interface

spacereclaimer start command [129](#)

commands

clone\_split clone\_status [187](#)

clone\_split clone\_stop [188](#)

clone\_split clone)\_estimate [186](#)

clonesplit start [187](#)

dataset [98](#)

dataset add\_members [98](#)

dataset backup\_add [99](#)

dataset backup\_change\_retention\_type [100](#)

dataset backup\_delete [100](#)

dataset backup\_end [101](#)

dataset backup\_get\_metadata [102](#)

dataset backup\_list [102](#)

dataset backup\_set\_metadata [103](#)

dataset backup\_start [104](#)

dataset backup\_status [104](#)

dataset backup\_version\_convert [105](#)

- dataset create [105](#)
- dataset create\_local\_backup [106](#)
- dataset delete [106](#)
- dataset dfm\_request [107](#)
- dataset get\_available\_policies [107](#)
- dataset get\_backup\_location [108](#)
- dataset get\_backup\_version\_info [108](#)
- dataset get\_metadata [109](#)
- dataset get\_policy [109](#)
- dataset get\_retention\_info [110](#)
- dataset info [110](#)
- dataset initiate\_conformance [111](#)
- dataset list\_members [111](#)
- dataset mount\_backup [112](#)
- dataset protect [113](#)
- dataset remove\_members [113](#)
- dataset restore [114](#)
- dataset restore\_status [114](#)
- dataset set\_metadata [115](#)
- dataset set\_policy [116](#)
- dataset transfer\_now [116](#)
- dataset vss\_backup\_end [117](#)
- dataset vss\_backup\_prepare [117](#)
- disk add\_initiator [151](#)
- disk add\_mount [156](#)
- disk connect [143](#)
- disk delete [146](#)
- disk disconnect [147](#)
- disk expand [150](#)
- disk list [154](#)
- disk remove\_initiator [153](#)
- disk remove\_mount [157](#)
- disk rename\_flexclone [158](#)
- hyperv\_config delete [186](#)
- hyperv\_config set [185](#)
- igroup create [121](#)
- igroup list [120](#)
- igroup modify [123](#)
- igroup rename [122](#)
- iscsi\_initiator establish\_session [138](#)
- iscsi\_initiator list [137](#)
- iscsi\_initiator terminate\_session [139](#)
- iscsi\_target disconnect [136](#)
- iscsi\_target list [136](#)
- license list [119](#)
- license remove [119](#)
- license set [118](#)
- oncommand\_config delete [177](#)
- oncommand\_config rbaccache [178](#)
- oncommand\_config set [176](#)
- preferredIP delete [135](#)
- preferredIP list [135](#)
- preferredIP set [134](#)
- sending a request to Protection Manager [107](#)
- snap create [159](#)
- snap delete [160](#)
- snap list [161](#)
- snap mirror\_list [162](#)
- snap mount [163](#)
- snap rename [164](#)
- snap restore\_volume [170](#)
- snap restore\_volume\_check [169](#)
- snap unmount [167](#)
- snap update\_mirror [168](#)
- snapvault archive [174](#)
- snapvault relationship\_status [174](#)
- snapvault snap\_list [175](#)
- snapvault snapshot\_delete [173](#)
- snapvault snapshot\_rename [172](#)
- snapvault verify\_configuration [171](#)
- spacemon delete [128](#)
- spacemon list [124](#)
- spacemon set [125](#)
- spacemon snap\_delta [126](#)
- spacemon snap\_reclaimable [127](#)
- spacemon vol\_info [127](#)
- spacereclaimer analyze [130](#)
- spacereclaimer status [131](#)
- spacereclaimer stop [130](#)
- sysconfig list [98](#)
- transport\_protocol delete [181](#)
- transport\_protocol list [179](#)
- transport\_protocol set [180](#)
- vsc\_config delete [133](#)
- vsc\_config list [132](#)
- vsc\_config set [132](#)
- vsconfig delete [184](#)
- vsconfig dslist [183](#)
- vsconfig list [182](#)
- vsconfig set [182](#)
- commands, hyperv\_config list [184](#)
- commandsigroup delete [123](#)
- comments
  - how to send feedback about documentation [199](#)
- communication using HTTP
  - enabling [22](#)
- components, SnapDrive
  - described [14](#)
- configuration
  - of a failover cluster witness disk [30](#)

- configuration information
  - displaying for your host [98](#)
- configuration requirements
  - for creating disks in clustered Data ONTAP [31](#)
  - for SnapVault support [71](#)
  - SnapVault [70](#)
- configurations
  - FC, single host, dedicated switch [192](#)
  - FC, single host, direct-attached [191](#)
  - FC, Windows cluster [192](#)
  - iSCSI, single host, dedicated switch [190](#)
  - iSCSI, single host, direct-attached [189](#)
  - iSCSI, single host, GbE switch [189](#)
  - iSCSI, Windows cluster [191](#)
  - MPIO, single host, direct-attached [193](#)
  - MPIO, Windows cluster, FC switch [195](#)
  - MPIO, Windows cluster, GbE switch [194](#)
- configuring
  - RBAC on SnapDrive [89](#)
  - space reservation monitoring [52](#)
- conformance check
  - initiating using dataset initiate\_conformance [111](#)
- Connect Disk wizard
  - using to connect your host to a LUN [33](#)
- connecting LUNs
  - list of guidelines for [33](#)
- connecting to a disk
  - using the Connect Disk wizard [33](#)
- connecting to a LUN
  - using the Connect Disk wizard [33](#)
- copies, Snapshot
  - reasons for creating [42](#)
- copy management
  - provided to a guest OS by ESX iSCSI initiators [21](#)
- creating
  - LUNs [26](#)
- creating a local dataset backup
  - using dataset create\_local\_backup [106](#)
- creating a new dataset
  - using the dataset create command [105](#)
- creating disks
  - on a virtual storage server [31](#)
- creating LUNs
  - rules for [25](#)
- creating RDM LUNs
  - on a guest OS [60](#)
- creating roles using RBAC [90](#)
- CSVs
  - See* Cluster Shared Volumes

## D

- data management
  - using SnapDrive for Windows [13](#)
- Data ONTAP DSM
  - using with SnapDrive [23](#)
- data protection capabilities
  - of OnCommand Unified Manager Core Package [78](#)
- Data protection capabilities
  - integration with SnapDrive [78](#)
- dataset
  - creating a local backup [106](#)
  - transferring a [116](#)
  - viewing members of [111](#)
- dataset commands
  - usage [98](#)
- dataset operations
  - using dataset commands [98](#)
- datasets
  - adding objects to [98](#)
  - changing the retention type [100](#)
  - creating a new [105](#)
  - deleting [106](#)
  - getting a backup copy location [108](#)
  - getting available retention policies [107](#)
  - getting backup metadata [102](#)
  - getting backup retention policy information [110](#)
  - getting information about [110](#)
  - getting information about the storage policy [109](#)
  - getting metadata for [109](#)
  - getting your backup version number [108](#)
  - initiating a conformance check [111](#)
  - listing backup copies [102](#)
  - mounting a backup of [112](#)
  - preparing for a VSS backup [117](#)
  - removing objects [113](#)
  - restoring from a backup [114](#)
  - restoring the status of [114](#)
  - scheduling a backup [99](#)
  - securing [113](#)
  - setting a storage policy [116](#)
  - setting backup metadata [103](#)
  - starting a backup [104](#)
  - stopping a VSS backup [117](#)
  - viewing a backup status [104](#)
- dedicated disk optimization
  - included in SnapDrive [11](#)
- dedicated LUN
  - creating [26](#)
- Delete Disk [37](#)

- deleting
  - a LUN [37](#)
  - folders within volume mount points [38](#)
  - volume mount point [34](#)
- deleting a dataset
  - using the dataset delete command [106](#)
- deleting a dataset backup
  - using dataset backup\_delete [100](#)
- deleting a remote backup
  - using dataset backup\_delete [100](#)
- disabling ESX login
  - from SnapDrive for Windows MMC [59](#)
- disabling vCenter login
  - from SnapDrive for Windows MMC [59](#)
- disconnecting
  - a LUN [36](#)
  - iSCSI session [20](#)
  - iSCSI target [19](#)
- disk add\_initiator command
  - syntax for [151](#)
- disk add\_mount command
  - syntax for [156](#)
- disk connect command
  - syntax for [143](#)
- disk delete command
  - syntax for [146](#)
- disk disconnect command
  - syntax for [147](#)
- disk expand command
  - syntax for [150](#)
- disk list command
  - syntax for [154](#)
- disk remove\_initiator command
  - syntax for [153](#)
- disk remove\_mount command
  - syntax for [157](#)
- disk rename\_flexclone command
  - syntax for [158](#)
- disk resize command
  - syntax for [148](#)
- disks
  - creating in clustered Data ONTAP [31](#)
  - creating on a virtual storage server [31](#)
  - expanding [38](#)
  - resizing guidelines [38](#)
  - shrinking [38](#)
- documentation
  - how to send feedback about [199](#)

## E

- enabling ESX login
  - from SnapDrive for Windows MMC [59](#)
- enabling vCenter login
  - from SnapDrive for Windows MMC [59](#)
- ending dataset backup operations
  - using dataset backup\_end [101](#)
- ESX iSCSI initiators
  - supported by SnapDrive for Windows [21](#)
- ESX iSCSI-accessed RDM LUNs
  - creating on a guest OS [60](#)
- ESX login
  - enabling and disabling from SnapDrive for Windows MMC [59](#)
- ESX server
  - limitations to SnapDrive support [58](#)
- ESX support
  - limitations [21](#)
- examples
  - storage system access control commands [85](#)
- expanding
  - a LUN [39](#)
- expanding LUNs
  - about [38](#)

## F

- failover cluster witness disks
  - configuring [30](#)
- failover clusters
  - support for [29](#)
- FC
  - configurations [191](#)
- FC RDM LUN
  - requirements for using in Microsoft clusters [63](#)
- FC-accessed RDM LUNs
  - creating on a guest OS [60](#)
- features, new
  - included in this release of SnapDrive [11](#)
- feedback
  - how to send comments about documentation [199](#)
- Fibre Channel sessions
  - managing [17](#)
- file-level restore operation
  - described [50](#)
  - support for [50](#)
- FlexClone volumes
  - about using in SnapDrive [45](#)
  - prerequisites for using with SnapDrive [45](#)

- forced disconnection
  - LUN [36](#)
- fractional space reservation monitoring
  - about [52](#)
- functionality
  - currently unsupported in SnapDrive for Windows [14](#)

## G

- getting backup [102](#)
- getting dataset information
  - using the dataset info command [110](#)
- getting retention policy information
  - using dataset get\_retention\_info [110](#)
- GPT partition
  - support [31](#)
- guest OS
  - providing LUN provisioning and copy management to [21](#)

## H

- high availability
  - providing [23](#)
- hot add and remove feature
  - requirements for using [40](#)
- HTTP communication
  - enabling [22](#)
- Hyper-V pass-through disk support
  - list of limitations [41](#)
- Hyper-V pass-through disks
  - requirements for dynamically adding and removing [40](#)
- hyperv\_config command,syntax for [184](#)
- hyperv\_config delete command
  - syntax for [186](#)
- hyperv\_config set command
  - syntax for [185](#)

## I

- igroup create command
  - syntax for using [121](#)
- igroup delete command
  - syntax for using [123](#)
- igroup modify command
  - parameters [123](#)
  - purpose [123](#)
  - syntax for using [123](#)
- igroup rename command

- syntax for using [122](#)
- igroups
  - command to display list of [120](#)
- information
  - how to send feedback about improving documentation [199](#)
- initiating
  - SnapMirror replication manually [75](#)
  - Update Mirror operations [73](#)
- initiating a conformance check
  - using dataset initiate\_conformance [111](#)
- initiator groups
  - See* igroups
- integration
  - SnapDrive with OnCommand Unified Manager Core Package data protection [78](#)
- iSCSI
  - configurations [189](#)
  - disconnecting target from Windows host [19](#)
  - examining session details [20](#)
- iSCSI session
  - disconnecting from target [20](#)
  - establishing with a target [17](#)
- iSCSI sessions
  - managing [17](#)
- iSCSI Software Initiator
  - node naming standards [17](#)
- iscsi\_initiator establish\_session command
  - syntax for [138](#)
- iscsi\_initiator list command
  - syntax for [137](#)
- iscsi\_initiator terminate\_session command
  - syntax for [139](#)
- iscsi\_target disconnect command
  - syntax for [136](#)
- iscsi\_target list command
  - syntax for [136](#)
- iSCSI-accessed RDM LUNs
  - creating on a guest OS [60](#)

## L

- license
  - viewing all installed [119](#)
- license operations
  - command-line options [118](#)
  - for SnapDrive modules [118](#)
- license remove command
  - syntax for using [119](#)
- limitations

- SnapVault [70](#)
  - to current SnapDrive for Windows functionality [14](#)
  - to SnapDrive support of ESX [21](#)
  - to SnapDrive support of Hyper-V pass-through disks [41](#)
  - VMDK and Snapshot functionality interoperability in SnapDrive [67](#)
- listing backups
  - using dataset backup\_list [102](#)
- listing members of a dataset
  - using dataset list\_members [111](#)
- login, vCenter or ESX
  - enabling and disabling from SnapDrive for Windows MMC [59](#)
- LUN
  - protocols [33](#)
- LUN clone split feature
  - advantages of [48](#)
  - how it works [48](#)
- LUN provisioning
  - provided to a guest OS by ESX iSCSI initiators [21](#)
- LUN restoration
  - how LUN clone splitting works [48](#)
  - how rapid LUN restore works [48](#)
- LUN restore
  - checking status [49](#)
- LUNs
  - about disconnecting from a FlexClone volume [35](#)
  - about disconnecting or deleting [35](#)
  - adding, removing, or changing a drive letter or path [34](#)
  - capabilities [33](#)
  - connecting to by using the Connect Disk wizard [33](#)
  - connecting to in Snapshot copies [46](#)
  - creating [26](#)
  - creating and using in VMware environments [58](#)
  - creating and using SnapDrive in Microsoft environments [23](#)
  - creating in VMware environments [60](#)
  - creating shared [28](#)
  - creating with virtual Fibre Channel [32](#)
  - deleting [37](#)
  - disconnecting [36](#)
  - expanding [38](#), [39](#)
  - FC RDM [64](#)
  - limitations [33](#)
  - managing [23](#), [32](#), [58](#)
  - managing in VMware environments [64](#)
  - Microsoft cluster support using RDM LUNs [64](#)
  - restoring on a volume from one Snapshot copy [77](#)

- LUNS
  - on a SnapMirror destination volume [76](#)
- LUNs, connecting
  - list of guidelines for [33](#)
- LUNs, creating
  - rules for [25](#)
- LUNs, FC RDM
  - requirements for using in a Microsoft cluster [63](#)

## M

- managing
  - rolling Snapshot copies [73](#)
- managing LUNs [32](#)
- managing space [52](#)
- manually initiating
  - SnapMirror replication [75](#)
- metadata
  - adding to a dataset [115](#)
  - obtaining for specified datasets [109](#)
- MetroClusters
  - support for [11](#)
- Microsoft MPIO
  - integrating with Data ONTAP DSM for Windows MPIO [23](#)
- mirroring operations
  - performing by integrating SnapDrive with SnapMirror [70](#)
- monitoring
  - space reservation [52](#)
- mount points, volume
  - about [26](#)
- mounting a dataset backup
  - using dataset mount\_backup [112](#)
- MPIO
  - configurations [193](#)

## N

- naming requirements
  - iSCSI Software Initiator [17](#)
- new features
  - included in this release of SnapDrive [11](#)
- nodes
  - naming standards for the iSCSI Software Initiator [17](#)

## O

- oncommand\_config delete command
  - syntax for [177](#)

- oncommand\_config rbacache command
  - syntax for [178](#)
- oncommand\_config set command
  - syntax for [176](#)
- operations
  - on datasets [98](#)

## P

- pass-through disks
  - requirements for dynamically adding and removing [40](#)
- pass-through LUNs
  - running Space Reclaimer on [56](#)
- path-failover methods
  - using [23](#)
- performance optimization
  - and spaceclaimer start command [129](#)
- preferredIP delete command
  - syntax for [135](#)
- preferredIP list command
  - syntax for [135](#)
- preferredIP set command
  - syntax for [134](#)
- preparing for a VSS backup
  - using dataset vss\_backup\_prepare [117](#)
- Protection Manager
  - sending requests to for backup version information [107](#)
- protocols
  - for accessing LUNs [33](#)
- provisioning, storage
  - automatic, using SnapDrive for Windows [13](#)
- provisioning, thin LUN
  - support included in SnapDrive [11](#)

## R

- rapid LUN restore
  - advantages of [48](#)
  - how it works [48](#)
- RBAC
  - configuring SnapDrive to use [89](#)
  - enabling on a storage system [89](#)
  - enabling to configure access control [80](#)
  - support for [87](#)
  - See also* RBAC
- RBAC (role-based access control)
  - assigning roles [91](#)
  - creating roles [90](#)

- RDM LUN
  - requirements for FC support [63](#)
  - shared [64](#)
- RDM LUNs
  - creating on a guest OS [60](#)
  - troubleshooting creation of [62](#)
- RDM operations
  - vCenter privileges for [60](#)
- RDMs
  - removing stale [35](#)
- Recycle Bin, Windows
  - bypassing when deleting folders within volume mount points [38](#)
- removal, LUN
  - unexpected [36](#)
- removing a dataset backup
  - using dataset backup\_delete [100](#)
- removing objects from a dataset
  - using dataset remove\_members [113](#)
- requirements
  - for space reclamation in VMDK files in NFS datastores [69](#)
  - hot add and remove [40](#)
- requirements, configuration
  - for creating disks in clustered Data ONTAP [31](#)
- restrictions
  - on creating Snapshot copies [42](#)
- restoration, LUN
  - how LUN clone splitting works [48](#)
  - how rapid LUN restore works [48](#)
- restoring
  - a volume on a SnapMirror destination [77](#)
- restoring a dataset status
  - using dataset restore\_status [114](#)
- restoring from a dataset backup
  - using dataset restore [114](#)
- restrictions
  - to current SnapDrive for Windows functionality [14](#)
- retention policy information
  - obtaining using dataset get\_retention\_info [110](#)
- role mappings
  - SnapDrive to DataFabric Manager server [91](#)
- role-based access control
  - See* RBAC
- roles
  - storage system access control [82](#)
- rolling Snapshot copies
  - checking for [73](#)
  - managing [73](#)

## S

- SAN booting
  - SnapDrive support for [196](#)
- SAN environment
  - initiating SnapVault backup jobs in [71](#)
- SAN environments
  - SnapDrive for Windows support for [13](#)
- schedules
  - Snapshot copy [44](#)
- scheduling a backup
  - using dataset backup\_add [99](#)
- SDCLI
  - dataset commands [98](#)
  - license commands [118](#)
  - sysconfig list command [98](#)
- sdcli spacereclaimer start command
  - syntax and parameter descriptions [129](#)
- session, iSCSI
  - establishing with a target [17](#)
- setting a dataset storage policy
  - using the dataset set\_policy command [116](#)
- settings, space reservation
  - managing by using the storage access control tool [53](#)
- shared FC RDM LUN
  - creating [64](#)
- shared LUNs
  - creating [28](#)
- shared RDM LUN
  - using in a Microsoft cluster [64](#)
- SMB 3.0 environments
  - SnapDrive for Windows support for [13](#)
- snap create command
  - syntax for [159](#)
- snap delete command
  - syntax for [160](#)
- snap list command
  - syntax for [161](#)
- snap mirror\_list command
  - syntax for [162](#)
- snap mount command
  - syntax for [163](#)
- snap rename command
  - syntax for [164](#)
- snap restore\_volume command
  - syntax for [170](#)
- snap restore\_volume\_check command
  - syntax for [169](#)
- snap unmount command
  - syntax for [167](#)
- snap update\_mirror command
  - syntax for [168](#)
- SnapDrive
  - configurations [189](#)
  - integrating with OnCommand Unified Manager Core Package [78](#)
  - integration with data protection capabilities [78](#)
  - new features in this release [11](#)
  - not listing all VMDK disks [68](#)
  - using in Microsoft environments [23](#)
  - using in VMware environments [58](#)
- SnapDrive CLI commands
  - igroup list [120](#)
- SnapDrive components
  - described [14](#)
- SnapDrive for Windows
  - capabilities [13](#)
  - in SMB 3.0 environments [13](#)
  - integration with Windows Volume Manager [13](#)
  - recommendations for using [16](#)
- SnapDrive to DataFabric Manager server role mappings [91](#)
- SnapManager for SQL
  - and SnapVault support in cluster environments [71](#)
- SnapMirror
  - integrating with SnapDrive to perform mirroring operations [70](#)
  - requirements for using with SnapDrive [74](#)
- Snapshot backup copies
  - managing using SnapDrive in Microsoft environments [23](#)
  - managing using SnapDrive in VMware environments [58](#)
- Snapshot backups
  - managing in VMware environments [66](#)
- Snapshot copies
  - backing up to a secondary storage system using SnapVault [70](#)
  - connecting to LUNs in [46](#)
  - creating [43](#)
  - creating in VMDKs [66](#)
  - deleting [50](#)
  - deleting in VMDKs [66](#)
  - reasons for creating [42](#)
  - restrictions on creating [42](#)
  - scheduling [44](#)
  - support with VMDKs on NFS and VMFS datastores [66](#)
- Snapshot copy management
  - provided to a guest OS by ESX iSCSI initiators [21](#)

- Snapshot copy support
  - interoperability with VMDKs in SnapDrive [67](#)
- SnapVault
  - 7-Mode limitations [70](#)
  - backup sets [70](#)
  - configuration requirements [70](#)
  - initiating backup jobs from SnapDrive 7-Mode SAN [71](#)
  - integrating with SnapDrive to perform vaulting operations [70](#)
  - limitations [70](#)
  - system requirements [70](#)
  - usage described [70](#)
- snapvault archive command
  - syntax for [174](#)
- snapvault relationship\_status command
  - syntax for [174](#)
- snapvault snap\_list command
  - syntax for [175](#)
- snapvault snapshot\_delete command
  - syntax for [173](#)
- snapvault snapshot\_rename command
  - syntax for [172](#)
- SnapVault support
  - by SnapDrive for Windows in cluster environments [71](#)
- snapvault verify\_configuration command
  - syntax for [171](#)
- space
  - managing [52](#)
- space optimization
  - and spacereclaimer start command [129](#)
- Space Reclaimer
  - about [54](#)
  - enabling [56](#)
  - guidelines for using [54](#)
  - reasons for SnapDrive to automatically stop [56](#)
  - running on pass-through LUNs [56](#)
  - starting [55](#)
  - stopping [56](#)
  - support with VMDK [69](#)
- space reclamation
  - requirements for VMDK files in NFS datastores [69](#)
- space reservation monitoring [52](#)
- space reservation settings
  - managing by using the storage access control tool [53](#)
- spacemon delete command
  - syntax for [128](#)
- spacemon list command
  - syntax for [124](#)
- spacemon set command
  - syntax for using [125](#)
- spacemon snap\_delta command
  - syntax for [126](#)
- spacemon snap\_reclaimable command
  - syntax for [127](#)
- spacemon vol\_info command
  - syntax for [127](#)
- spacereclaimer analyze
  - syntax for [130](#)
- spacereclaimer start command
  - syntax and parameter descriptions [129](#)
- spacereclaimer status
  - syntax for [131](#)
- spacereclaimer stop command
  - syntax for [130](#)
- starting
  - Space Reclaimer [55](#)
- starting a backup
  - using the dataset backup\_start command [104](#)
- starting a conformance check
  - using dataset initiate\_conformance [111](#)
- stopping
  - Space Reclaimer [56](#)
- stopping a VSS backup
  - using dataset vss\_backup\_end [117](#)
- storacl.exe
  - using to enable thinly provisioned LUNs [53](#)
- storage access control tool
  - using to enable thinly provisioned LUNs [53](#)
- storage policy information
  - getting using the dataset get\_policy command [109](#)
- storage provisioning, automatic
  - using SnapDrive for Windows [13](#)
- storage system
  - configuring access control [80](#)
  - managing access control [80](#)
- storage system access
  - allowing using AccessControl.xml operations [81](#)
  - restricting using AccessControl.xml operations [81](#)
- storage system access control
  - command examples [85](#)
  - described [80](#)
  - roles, described [82](#)
  - support for [80](#)
- storage system access control tool (storacl.exe)
  - using [80](#)
- storage system capabilities
  - adding for HTTP users [22](#)
- storage systems

- access control operations described [81](#)
- access control tool [84](#)
- controlling user activity on [84](#)
- enabling RBAC on [89](#)
- establishing a connection to [17](#)
- managing space on using SnapDrive in Microsoft environments [23](#)
- managing space on using SnapDrive in VMware environments [58](#)

Storage Virtual Machine

*See* SVM

suggestions

- how to send feedback about documentation [199](#)

support, clone

- included in SnapDrive [11](#)

surprise removal

- LUN [36](#)

SVM

- configuration requirements in clustered Data

- ONTAP [31](#)

system requirements

- SnapVault [70](#)

## T

thin LUN provisioning

- support included in SnapDrive [11](#)

thinly provisioned LUNs

- enabling by using the storage access control tool [53](#)

timestamps

- converting backup version numbers to [105](#)

transferring a dataset

- using dataset transfer\_now [116](#)

transport protocol settings

- managing [17](#)

transport\_protocol delete command

- syntax for [181](#)

transport\_protocol list command

- syntax for [179](#)

transport\_protocol set command

- syntax for [180](#)

troubleshooting

- RDM LUN creation [62](#)

- VMDKs [67](#), [68](#)

## U

Update mirror operations

- initiating [73](#)

using the storage system access control tool (storacl.exe)

- to determine operations enabled for users [80](#)

- to determine user roles [80](#)

## V

vaulting operations

- performing by integrating SnapDrive with SnapVault [70](#)

vCenter login

- enabling and disabling from SnapDrive for Windows MMC [59](#)

vCenter privileges

- for performing RDM operations [60](#)

- minimum [60](#)

verifying

- cluster group owner [24](#)

viewing backup copies

- using dataset backup\_list [102](#)

virtual Fibre Channel

- configuring SnapDrive for Windows [32](#)

- creating LUNs with [32](#)

Virtual Storage Console [132](#)

VMDK

- disk partition style [68](#)

- inconsistent disk enumeration [67](#)

- interoperability with Snapshot functionality and SnapDrive [67](#)

- Snapshot copy support [66](#)

- Space Reclaimer support [69](#)

- troubleshooting [67](#), [68](#)

VMDKs

- not listed in SnapDrive enumeration [68](#)

- troubleshooting [67](#), [68](#)

VMware

- support for ESX iSCSI initiators [58](#)

- support for FC HBAs [58](#)

- support for iSCSI HBAs [58](#)

- support for iSCSI initiators [58](#)

VMware ESX

- support for [58](#)

VMware ESX server

- limitations to SnapDrive support [58](#)

VMware guest OS platforms

- SnapDrive for Windows support for [58](#)

volume mount points

- about [26](#)

- adding [34](#)

- changing [34](#)

- deleting folders within [38](#)

- removing [34](#)

- requirements [26](#)
- volume restore
  - performing [77](#)
- Volume Shadow Copy Service
  - using with SnapDrive for Windows [15](#)
- volume-based Snapshot copy restoration
  - about [50](#)
- volume-level SnapVault support
  - by SnapDrive for Windows in cluster environments [71](#)
- volumes, FlexClone
  - about using in SnapDrive [45](#)
- vsc\_config delete
  - syntax for [133](#)
- vsc\_config list command
  - syntax for [132](#)
- vsc\_config set
  - syntax for [132](#)
- vsconfig delete command
  - syntax for [184](#)
- vsconfig dslist command

- syntax for [183](#)
- vsconfig list command
  - syntax for [182](#)
- vsconfig set command
  - syntax for [182](#)
- VSS
  - See* Volume Shadow Copy Service
- VSS backups
  - preparing for [117](#)

## W

- Windows Recycle Bin
  - bypassing when deleting folders within volume mount points [38](#)
- Windows Server failover cluster
  - support for [29](#)
- Windows Volume Manager
  - integration with SnapDrive for Windows [13](#)