




Updated for 8.2.1

## Clustered Data ONTAP<sup>®</sup> 8.2

### Data Protection Guide



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  
Web: [www.netapp.com](http://www.netapp.com)  
Feedback: [doccomments@netapp.com](mailto:doccomments@netapp.com)

Part number: 215-08504\_B0  
February 2014



# Contents

<b>Introduction to data protection .....</b>	<b>7</b>
Methods of protecting data .....	7
Database protection .....	8
What a data loss disaster is .....	8
Tools for protecting against data-loss disasters .....	9
Data protection in a SAN environment .....	10
Types of data protection policies .....	10
<b>Planning your data protection strategy .....</b>	<b>12</b>
Working with Snapshot copies .....	12
What a Snapshot copy is .....	12
User access to Snapshot copies .....	12
Backup and recovery tasks you can perform with Snapshot copies .....	13
Maximum number of Snapshot copies .....	14
Where to find information about Snapshot copies of Infinite Volumes .....	14
Creation of Snapshot copy schedules .....	15
Deleting Snapshot copies automatically .....	17
Viewing settings for the automatic deletion of Snapshot copies .....	18
What Snapshot disk consumption is .....	19
What the Snapshot copy reserve is .....	20
Working with mirroring technology .....	23
Components of a mirror relationship .....	23
Data protection mirror relationships for FlexVol volumes .....	23
Where to find information about data protection mirror relationships for Infinite Volumes .....	24
When a destination volume grows automatically .....	24
Path name pattern matching .....	24
Language setting requirement .....	25
User access to destination volumes .....	25
Guidelines for creating relationships between clusters or SVMs .....	25
Limitations for data protection mirror relationships .....	27
Working with FlexVol volume SnapVault backups .....	28
What a SnapVault backup is .....	28

Which data gets backed up and restored from a SnapVault backup .....	29
Which data does not get backed up to a SnapVault backup .....	30
How a SnapVault backup works .....	31
How SnapVault backups work with data compression .....	32
SnapVault backup limitations .....	32
Data protection for Storage Virtual Machine (SVM) namespace and root information .....	32
Guidelines for planning Snapshot copy schedule and retention for SnapVault backups .....	33
Supported data protection deployment configurations .....	34
What a basic backup deployment is .....	36
Firewall requirements for intercluster SnapMirror relationships .....	36
What source-to-destination-to-tape backup is .....	36
How a mirror-mirror cascade works .....	37
How a mirror-SnapVault cascade works .....	38
How a SnapVault-SnapMirror cascade works .....	39
How a mirror-SnapVault fanout works .....	39
How a multiple-mirrors fanout works .....	40
<b>Protecting data using Snapshot copies .....</b>	<b>42</b>
Managing Snapshot copies .....	42
Commands for managing Snapshot copies .....	42
Managing Snapshot policies .....	43
How Snapshot policies are associated with volumes .....	43
Commands for managing Snapshot policies and schedules .....	44
Strategies for creating a Snapshot copy policy .....	45
Naming convention for scheduled Snapshot copies .....	45
What prefixes are .....	46
Using prefixes to name automatic Snapshot copies .....	46
Restoring files from the Snapshot copy of a FlexVol volume .....	47
Restoring a single file from a Snapshot copy of a FlexVol volume .....	47
Restoring part of a file from a Snapshot copy of a FlexVol volume .....	48
Restoring the contents of a volume from a Snapshot copy .....	48
Snapshot restoration using Shadow Copy Client tools .....	49
Managing Snapshot copy disk space .....	49
Monitoring Snapshot copy disk consumption .....	49
<b>Managing data protection using SnapMirror policies .....</b>	<b>51</b>

Commands for managing mirror and SnapVault policies .....	51
How SnapMirror policies work with clusters and Storage Virtual Machines (SVMs) .....	52
Comparison of what cluster administrators and Storage Virtual Machine (SVM) administrators can manage .....	52
Guidelines for naming Snapmirror and SnapVault policies .....	53
Preserving Snapshot copies for SnapVault relationships after reaching retention limit .....	53
Example of creating a tiered backup policy .....	54
<b>Providing disaster recovery using mirroring technology .....</b>	<b>56</b>
Managing data protection mirror copies for FlexVol volumes .....	56
Creating a data protection mirror copy for FlexVol volumes .....	56
Correcting a SnapMirror initialization failure .....	58
Managing mirror relationships .....	59
Commands for managing SnapMirror relationships .....	59
Using extended queries to operate on many SnapMirror relationships .....	63
What tape seeding is .....	63
Performing tape seeding using SMTape .....	63
Scalability limits for SMTape backup and restore sessions .....	71
Listing the schedule state of a mirror relationship .....	71
Scheduling SnapMirror transfers .....	71
Changing mirror relationship schedules .....	72
Manually updating data protection mirror copies on destination volumes ...	73
Deleting a mirror copy .....	74
Reversing the data protection mirror relationship when disaster occurs .....	75
Converting a data protection mirror destination to a writeable volume .....	79
Testing database applications .....	80
<b>Protecting data on FlexVol volumes by using SnapVault backups .....</b>	<b>81</b>
Creating SnapVault backups on FlexVol volumes .....	81
Guidelines for creating SnapVault relationships on FlexVol volumes .....	81
SnapVault updates fail if destination aggregate is full .....	83
Prepopulated SnapVault secondary scenarios .....	83
Creating a SnapVault backup in an empty FlexVol volume .....	84
Creating the SnapVault relationship of a mirror-SnapVault cascade .....	87
Preserving a Snapshot copy on the primary source volume .....	88
Creating a SnapVault backup in a prepopulated FlexVol volume .....	90

Creating a destination baseline using a tape backup .....	91
Converting a data protection destination to a SnapVault secondary .....	92
Managing backup and restore operations for SnapVault backups .....	93
Backing up from a Snapshot copy that is older than the base Snapshot copy .....	93
Backing up FlexVol volumes that contain the maximum limit of Snapshot copies .....	97
Managing the backup of a copied source volume .....	98
Guidelines for restoring the active file system .....	98
Guidelines for restoring LUNs in SAN environments .....	99
How restore operations work from a SnapVault backup .....	99
Restoring a volume from a SnapVault backup .....	100
Managing a SnapVault-mirror cascade when the SnapVault backup is unavailable .....	101
Managing storage efficiency for SnapVault secondaries .....	104
Guidelines for managing storage efficiency for SnapVault backups .....	105
Enabling storage efficiency on a SnapVault secondary volume .....	106
<b>Copyright information .....</b>	<b>107</b>
<b>Trademark information .....</b>	<b>108</b>
<b>How to send your comments .....</b>	<b>109</b>
<b>Index .....</b>	<b>110</b>

# Introduction to data protection

---

Data protection means backing up data and being able to recover it. You protect the data by making copies of it so that it is available for restoration even if the original is no longer available.

Businesses need data backup and protection for the following reasons:

- To protect data from accidentally deleted files, application crashes, data corruption, and viruses
- To archive data for future use
- To recover from a disaster

## Methods of protecting data

Depending on your data protection and backup needs, Data ONTAP offers a variety of features and methods that enable you to protect data against accidental, malicious, or disaster-induced loss of data.

**Snapshot copies** Enable you to manually or automatically create, schedule, and maintain multiple backups (also called *Snapshot copies*) of data on a volume. Snapshot copies use only a minimal amount of additional volume space, and do not have a performance cost.

If a user accidentally modifies or deletes crucial data on a volume with Snapshot technology enabled, that data can be easily and quickly restored from one of the latest Snapshot copies created. You can also create clones of FlexVol volumes using Snapshot copies.

This method is valid for FlexVol volumes and Infinite Volumes.

**SnapRestore (license required)** Enables you to perform fast, space-efficient, on-request Snapshot recovery from Snapshot copies on an entire volume.

This method is valid for FlexVol volumes and Infinite Volumes.

**Data protection mirror copies (SnapMirror license required)** Provide asynchronous disaster recovery. Data protection mirror relationships enable you to periodically create Snapshot copies of data on one volume; copy those Snapshot copies to a partner volume (the destination volume), usually on another cluster; and retain those Snapshot copies. The mirror copy on the destination volume ensures quick availability and restoration of data from the time of the latest Snapshot copy, if the data on the source volume is corrupted or lost.

If you conduct tape backup and archival operations, you can perform them on the data that is already backed up on the destination volume.

This method is valid for FlexVol volumes and Infinite Volumes.

<b>SnapVault backups (SnapVault license required)</b>	<p>Provide storage-efficient and long-term retention of backups. SnapVault relationships enable you to back up selected Snapshot copies of volumes to a destination volume and retain the backups.</p> <p>If you conduct tape backup and archival operations, you can perform them on the data that is already backed up on the SnapVault secondary volume.</p> <p>This method is valid only for FlexVol volumes.</p>
<b>volume copy</b>	<p>Enables you to perform fast block-copy of data from one volume to another.</p> <p>This method is valid only for FlexVol volumes.</p>
<b>nvfail option to the volume modify command</b>	<p>Provides protection against data corruption by failures of nonvolatile RAM (NVRAM).</p> <p>This method is valid for FlexVol volumes and Infinite Volumes.</p>

## Database protection

If NVRAM problems occur that compromise database validity, Data ONTAP can warn you and automatically rename the database so that it does not restart automatically. You can then ensure that the database is valid before restarting it.

Data ONTAP provides database protection using the `nvfail` option of the `volume create` or `volume modify` command.

For SnapVault relationships in SAN environments, the `nvfail` attribute of the LUN on a SnapVault secondary volume is always off.

For Infinite Volumes, it is recommended to leave the `nvfail` option disabled.

**Note:** You can use this feature only when there are databases on the storage system.

## What a data loss disaster is

A data loss disaster is a situation in which service from one physical site (for example, a building or a corporate campus) on the network is lost for an extended period of time.

The following are examples of disasters:

- Fire
- Earthquake
- Prolonged power outages at a site
- Prolonged loss of connectivity from clients to the storage system at a site



When a disaster occurs, it can affect all the computing infrastructure including storage systems, application servers, networking connectivity, and client connectivity. When you create a disaster plan, you should take your computing infrastructure into consideration.

## Tools for protecting against data-loss disasters

Data ONTAP provides tools that enable you to back up or replicate data stored at a primary data storage site to an off-site network location. This ensures that you can restore data if data loss is caused by disaster at a primary data storage site.

### SnapVault backups for FlexVol volumes

SnapVault is a Snapshot copy backup and restorability tool on FlexVol volumes. You can locate a SnapVault secondary volume on the same cluster or on a different cluster.

<b>Data recoverability</b>	If a data-loss disaster occurs at a source volume, you can restore data that is backed up to a SnapVault secondary volume. You can restore the data to the source volume after it is running again, or you can restore data to an alternate volume.
<b>Currency of restore data</b>	You can restore data from any Snapshot copy that was replicated to the destination system.
<b>Advantage</b>	A SnapVault backup provides an inexpensive backup solution.

### Data protection mirror copy

A data protection mirror copy is a Snapshot copy replication, availability, and restorability tool. You can locate a data protection mirror destination on the same cluster or on a different cluster.

<b>Data availability</b>	If a source site experiences a data-loss disaster, you can quickly make available data at the data protection mirror copy destination site.
<b>Data recoverability</b>	If a data-loss disaster occurs at a source storage site, you can restore data from a data protection mirror copy destination volume. You can restore the data to the source volume after it is running again, or you can restore data to an alternate volume.
<b>Currency of restore data</b>	You can restore data from the last Snapshot copy that was replicated to the destination volume.
<b>Advantage</b>	Data protection mirror copies provide data protection and availability.

## Data protection in a SAN environment

If FlexVol volumes contain logical units of storage (LUNs) created to enable integration into a storage area network (SAN) environment, the procedures to implement data protection might have to be modified. Infinite Volumes do not support SAN environments or LUNs.

Data protection mirror copies and SnapVault backups are achieved by the use of volume-to-volume relationships. Therefore, to protect data in a LUN, you back up the volume that contains the LUN.

Path-related metadata such as Persistent Reservations, are not replicated to a SnapVault backup. When you restore a volume from a SnapVault secondary volume, the LUNs in the SnapVault secondary volume are exported with a different identity from their counterparts in the source volume. Therefore, you must configure new access controls for the restored LUNs.

For more information about the descriptions of data backup and restore on volumes containing Data ONTAP LUNs, see the *Clustered Data ONTAP SAN Administration Guide*.

## Types of data protection policies

You can assign Snapshot policies to FlexVol volumes and Infinite Volumes, and SnapMirror policies to data protection mirror relationships and SnapVault relationships.

### Snapshot policy

When you assign a Snapshot policy, the policy configures the Snapshot copy creation schedule and retention rules.

You can assign the same Snapshot policy to multiple volumes. For example, you might configure a Snapshot policy to create a Snapshot copy every hour, at the end of every day, and at the end of every week, and then assign that same policy to more than one volume.

You can assign only one Snapshot policy to a volume. You can assign Snapshot policies to FlexVol volumes and Infinite Volumes.

**Note:** You cannot assign a Snapshot policy that contains the `-snapmirror-label` to an Infinite Volume.

### SnapMirror policy

The SnapMirror policy specifies the configuration attributes of a relationship.

A SnapMirror policy can be applied to a data protection mirror relationship or a SnapVault relationship. Whether the SnapMirror policy has rules determines if the policy is applied to a SnapVault relationship or applied to a data protection mirror copy. If the policy has rules that define which Snapshot copies are protected, then that policy can be applied to SnapVault relationships only. If the policy does not have rules, then that policy can be applied to data protection mirror copies only.

**Note:** If no policy is assigned to a relationship, a default policy is assigned. If it is a data protection mirror relationship, the DPDefault policy is assigned. If it is a SnapVault relationship, the XDPDefault policy is assigned.

## Planning your data protection strategy

---

Data ONTAP provides a variety of tools that you can use to build a comprehensive strategy to protect your company's data.

Storage Virtual Machine (SVM) administrators can plan data protection for FlexVol volumes and Infinite volumes within their assigned SVMs. Cluster administrators can plan data protection for FlexVol volumes and Infinite Volumes within their assigned clusters.

## Working with Snapshot copies

Snapshot copies are the first line of defense for data protection. Data ONTAP maintains a configurable Snapshot schedule that creates and deletes Snapshot copies automatically for each FlexVol volume and Infinite Volume. You can also create and delete Snapshot copies, and manage Snapshot schedules based on your requirements.

### What a Snapshot copy is

A Snapshot copy is a read-only image of a FlexVol volume or Infinite Volume that captures the state of the file system at a point in time.

For information about FlexVol volumes, see the *Clustered Data ONTAP Physical Storage Management Guide*.

### User access to Snapshot copies

A Snapshot copy is a copy of a FlexVol volume that represents the volume's contents at a particular point in time. You can view the contents of the Snapshot copy and use the Snapshot copy to restore data that you lost recently.

A Snapshot copy of a volume is located on the parent volume but has read-only access. It represents the contents of the original volume at a particular point in time. A parent volume and a Snapshot copy of it share disk space for all blocks that have not been modified between the creation of the volume and the time the Snapshot copy is made, thereby making Snapshot copies lightweight.

Similarly, two Snapshot copies share disk space for those blocks that were not modified between the times that the two Snapshot copies were created. You can create a chain of Snapshot copies to represent the state of a volume at a number of points in time. Users can access Snapshot copies online, enabling users to retrieve their own data from past copies, rather than asking a system administrator to restore data from tape. Administrators can restore the contents of a volume from a Snapshot copy.

Each volume has a `.snapshot` directory that is accessible to NFS users by using the `ls` command and to CIFS users by double-clicking the `~snapshot` folder. The contents of the `.snapshot` directory are a set of subdirectories, labeled by type, date, and time, resembling the following:

```
$ ls .snapshot
daily.2006-05-14_0013/      hourly.2006-05-15_1306/
daily.2006-05-15_0012/    hourly.2006-05-15_1406/
hourly.2006-05-15_1006/   hourly.2006-05-15_1506/
hourly.2006-05-15_1106/   weekly.2006-05-14_0019/
hourly.2006-05-15_1206/
```

Each subdirectory of the `.snapshot` directory includes a list of the parent volume's files and directories. If users accidentally delete or overwrite a file, they can locate it in the most recent Snapshot directory and restore it to their main read-write volume simply by copying it back to the main directory. The following example shows how an NFS user can locate and retrieve a file named `my.txt` from the `.snapshot` directory:

```
$ ls my.txt
ls: my.txt: No such file or directory
$ ls .snapshot
daily.2006-05-14_0013/      hourly.2006-05-15_1306/
daily.2006-05-15_0012/    hourly.2006-05-15_1406/
hourly.2006-05-15_1006/   hourly.2006-05-15_1506/
hourly.2006-05-15_1106/   weekly.2006-05-14_0019/
hourly.2006-05-15_1206/
$ ls .snapshot/hourly.2006-05-15_1506/my.txt
my.txt
$ cp .snapshot/hourly.2006-05-15_1506/my.txt .
$ ls my.txt
my.txt
```

The `.snapshot` directory is always visible to NFSv2 and NFSv3 clients and available from within the volume, and not visible but still available from any other volume. For NFSv4 clients, the `.snapshot` directory is not visible, but accessible in all paths of a volume.

## Backup and recovery tasks you can perform with Snapshot copies

Snapshot copies enable system administrators and end users to perform important tasks in backup and recovery.

Snapshot copies enable system administrators to perform the following tasks:

- Create instantaneous backups
- Create a clone of a FlexVol volume
- Create a clone of a Data ONTAP LUN

For information about cloning a FlexVol volume, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Snapshot copies enable end users to perform the following tasks:

- Recover older versions or sets of files that were accidentally changed or deleted
- Restore their own files without needing a system administrator to restore files from tape

## Maximum number of Snapshot copies

You should know what the maximum number of Snapshot copies you can accumulate is to minimize the possibility that you do not have Snapshot copies available when you need them.

The maximum number of Snapshot copies follows:

- You can accumulate a maximum of 255 Snapshot copies of a FlexVol volume.
- If the FlexVol volume is in a data protection mirror relationship, the maximum number of Snapshot copies is 254 because one Snapshot copy is reserved for use by the relationship during recovery operations.
- If the FlexVol volume is in a disk to disk backup relationship, the maximum number of Snapshot copies is 251.
- If the Infinite Volume is in a data protection mirror relationship, the maximum number of Snapshot copies is reduced by two because two Snapshot copies are used for the data protection mirror relationship.

Over time, automatically generated hourly, weekly, and monthly Snapshot copies accrue. Having a number of Snapshot copies available gives you a greater degree of accuracy if you have to restore a file.

The number of Snapshot copies can approach the maximum if you do not remove older Snapshot copies. You can configure Data ONTAP to automatically delete older Snapshot copies of volumes as the number of Snapshot copies approaches the maximum.

The following data protection mirror copies affect the maximum number of Snapshot copies available to a volume:

- A FlexVol volume in a data protection mirror relationship
- A FlexVol volume with a load-sharing mirror copy
- An Infinite Volume with one or more namespace mirror constituents  
Each namespace mirror constituent uses two Snapshot copies. By default, a read/write Infinite Volume contains one namespace mirror constituent. If you enable SnapDiff on an Infinite Volume, each additional namespace mirror uses two Snapshot copies.

An Infinite Volume also uses up to four Snapshot copies when technical support runs some commands that require diagnostic privilege. You must keep the number of Snapshot copies far enough below the limit to ensure that technical support can run commands.

## Where to find information about Snapshot copies of Infinite Volumes

Information about Snapshot copies of Infinite Volumes is available in the *Clustered Data ONTAP Infinite Volumes Management Guide*.

## Creation of Snapshot copy schedules

Data ONTAP provides a default Snapshot copy schedule for each FlexVol volume and Infinite Volume. You can create schedules to fit your needs if the default Snapshot copy schedule is not adequate.

For FlexVol volumes, the default Snapshot copy schedule automatically creates one daily Snapshot copy Monday through Saturday at midnight, an hourly Snapshot copy five minutes past the hour, every hour, and a weekly Snapshot copy. Data ONTAP retains the two most recent nightly Snapshot copies and the six most recent hourly Snapshot copies, and deletes the oldest nightly and hourly Snapshot copies when new Snapshot copies are created.

## Types of user-specified Snapshot copy schedules

Data ONTAP contains weekly, daily, and hourly Snapshot copy schedules that you can use to create Snapshot copy policies that retain the number and type of Snapshot copies you want.

The following table describes the available types of Snapshot copy schedules:

Type	Description
Weekly	Data ONTAP creates these Snapshot copies every Sunday at 15 minutes after midnight. Weekly Snapshot copies are named <code>weekly.n</code> , where <code>n</code> is the date in year-month-day format followed by an underscore ( <code>_</code> ) and the time. For example, a weekly Snapshot copy created on 25 November 2012 is named <code>weekly.2012-11-25_0015</code> .
Daily	Data ONTAP creates these Snapshot copies every night at 10 minutes after midnight. Daily Snapshot copies are named <code>daily.n</code> , where <code>n</code> is the date in year-month-day format followed by an underscore ( <code>_</code> ) and the time. For example, a daily Snapshot copy created on 4 December 2012 is named <code>daily.2012-12-04_0010</code> .
Hourly	Data ONTAP creates these Snapshot copies five minutes after the hour. Hourly Snapshot copies are named <code>hourly.n</code> , where <code>n</code> is the date in year-month-day format followed by an underscore ( <code>_</code> ) and the time. For example, an hourly Snapshot copy created on 4 December 2012 at 1:00 (1300) is named <code>hourly.2012-12-04_1305</code> .

### Related concepts

[Commands for managing Snapshot policies and schedules](#) on page 44

## Creating a Snapshot copy schedule

If the default Snapshot copy schedule does not meet your needs, you can create a schedule that does.

### Step

1. Create a Snapshot copy schedule by using the `job schedule cron create` command or the `job schedule interval create` command.

The command you use depends on how you want to implement the schedule. See the man page for each command to determine the command that meets your needs.

## If scheduled Snapshot copy creation fails

Scheduled Snapshot copy creation might fail for various reasons, such as a volume being unavailable. In such cases, Data ONTAP attempts to create a Snapshot copy, when possible, outside the schedule.

If a scheduled Snapshot copy creation fails, Data ONTAP checks the Snapshot copies present in the volume. The checks performed and the actions taken depend on the type of scheduled Snapshot copy creation that failed. The process is described in the following list:

1. When a volume becomes available again for creating a Snapshot copy, Data ONTAP checks whether any Snapshot copies were created during a time period represented by `period_snap`. `period_snap` is a variable representing a time period that depends on the type of Snapshot copy schedule, as shown in the following table:

Type of Snapshot copy schedule	Value of the <code>period_snap</code> variable
Weekly	3 days
Nightly	3 days
Hourly	12 hours

**Note:** You cannot change the value of `period_snap`.

2. The check in the previous step returns one of the following values:

If the check returns...	Then...
Yes (One or more Snapshot copies were created in the <code>period_snap</code> period)	Data ONTAP performs Step 3.
No (Snapshot copies were not created in the <code>period_snap</code> period)	Data ONTAP performs Step 4.

3. Data ONTAP checks whether any scheduled Snapshot copy creations failed after the most recent Snapshot copy. This check returns one of the following values:



If the check returns...	Then...
Yes (One or more scheduled Snapshot copy creations were missed)	Data ONTAP creates a Snapshot copy.
No (No scheduled Snapshot copy creation was missed)	Data ONTAP does not create a Snapshot copy.

4. Data ONTAP checks whether any scheduled Snapshot copy creation have failed in the past 25 minutes. This check returns one of the following values:

If the check returns...	Then...
Yes (A scheduled Snapshot copy creation was missed in the past 25 minutes)	Data ONTAP creates a Snapshot copy.
No (No scheduled Snapshot copy creation was missed in the past 25 minutes)	Data ONTAP does not create a Snapshot copy.

## Deleting Snapshot copies automatically

You can define and enable a policy for automatically deleting Snapshot copies and FlexClone LUNs. Automatically deleting Snapshot copies and FlexClone LUNs can help you manage space utilization.

### About this task

You can automatically delete Snapshot copies from read-write volumes and FlexClone LUNs from read-write parent volumes. You cannot set up automatic deletion of Snapshot copies from Infinite Volumes or from read-only volumes, for example, SnapMirror destination volumes.

### Step

1. You define and enable a policy for automatically deleting Snapshot copies by using the `volume snapshot autodelete modify` command.

See the `volume snapshot autodelete modify man` page for information about the parameters that you can use with this command to define a policy that meets your needs.

### Example

The following command enables the automatic deletion of Snapshot copies and sets the trigger to `snap_reserve` for the `vol3` volume, which is part of the `vs0.example.com` Storage Virtual Machine (SVM):

```
cluster1::> volume snapshot autodelete modify -vserver
vs0.example.com
-volume vol3 -enabled true -trigger snap_reserve
```

**Example**

The following command enables the automatic deletion of Snapshot copies and of FlexClone LUNs for the vol3 volume, which is part of the vs0.example.com Storage Virtual Machine (SVM):

```
cluster1::> volume snapshot autodelete modify -vserver
vs0.example.com
-volume vol3 -enabled true -trigger volume -commitment try -delete-
order
oldest_first -destroy-list lun_clone,file_clone
```

**Viewing settings for the automatic deletion of Snapshot copies**

You can view the settings for the automatic deletion of Snapshot copies to help you when you are deciding if the settings are meeting your needs.

**Step**

1. View the settings for the automatic deletion of Snapshot copies by using the `volume snapshot autodelete show` command.

See the `volume snapshot autodelete show` command man pages for information about parameters shown by this command.

**Example**

The following command displays the automatic deletion settings of Snapshot copies for the vol3 volume, which is part of the vs0.example.com Storage Virtual Machine (SVM):

```
cluster1::> volume snapshot autodelete show -vserver vs0 -volume vol3
```

Vserver	Volume	Option Name	Option Value
vs0	vol3	Enabled	false
		Commitment	try
		Trigger	volume
		Target Free Space	20%
		Delete Order	oldest_first
		Defer Delete	user_created
		Defer Delete Prefix	(not specified)
		Destroy List	none

## What Snapshot disk consumption is


Data ONTAP preserves pointers to all the disk blocks currently in use at the time the Snapshot copy is created. When a file is changed, the Snapshot copy still points to the disk blocks where the file existed before it was modified, and changes are written to new disk blocks.

## How Snapshot copies consume disk space


Snapshot copies minimize disk consumption by preserving individual blocks rather than whole files. Snapshot copies begin to consume extra space only when files in the active file system are changed or deleted. When this happens, the original file blocks are still preserved as part of one or more Snapshot copies.

In the active file system the changed blocks are rewritten to different locations on the disk or removed as active file blocks entirely. As a result, in addition to the disk space used by blocks in the modified active file system, disk space used by the original blocks is still reserved to reflect the status of the active file system before the change.


The following illustration shows disk space usage for a Snapshot copy:



Before any Snapshot copy is created, disk space is consumed by the active file system only.



After a Snapshot copy is created, the active file system and Snapshot copy point to the same disk blocks. The Snapshot copy does not use extra disk space.




After *myfile.txt* is deleted from the active file system, the Snapshot copy still includes the file and references its disk blocks. That is why deleting active file system data does not always free disk space.

- Space used by the active file system
- Space used by the Snapshot copy only
- Space shared by the Snapshot copy and the active file system
- Unused disk space


## How changing file content consumes disk space

A given file might be part of a Snapshot copy. The changes to such a file are written to new blocks. Therefore, the blocks within the Snapshot copy and the new (changed or added) blocks both use space within the volume.

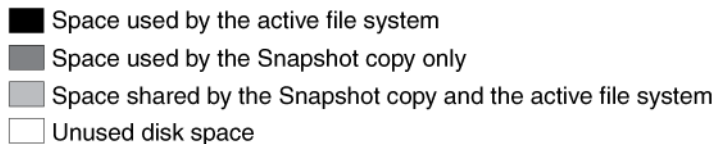
Changing the contents of the `myfile.txt` file creates a situation where the new data written to `myfile.txt` cannot be stored in the same disk blocks as the current contents because the Snapshot copy is using those disk blocks to store the old version of `myfile.txt`. Instead, the new data is written to new disk blocks. As the following illustration shows, there are now two separate copies of `myfile.txt` on disk—a new copy in the active file system and an old one in the Snapshot copy:



After a Snapshot copy is created, the active file system and Snapshot copy point to the same disk blocks, and the Snapshot copy does not use any extra space.



After a change is made to the file, the active file system and Snapshot copy no longer point to the same disk blocks, and the Snapshot copy now uses extra space.



## What the Snapshot copy reserve is

The Snapshot copy reserve sets a specific percent of the disk space for Snapshot copies. For FlexVol volumes, the default Snapshot copy reserve is set to 5 percent of the disk space. By default, the Snapshot copy reserve is 5 percent of the disk space for a FlexVol volume and 0 percent for aggregates.

The active file system cannot consume the Snapshot copy reserve space, but the Snapshot copy reserve, if exhausted, can use space in the active file system.

## How Data ONTAP uses deleted active file disk space

When enough disk space is available for Snapshot copies in the Snapshot copy reserve, deleting files in the active file system frees disk space for new files, while the Snapshot copies that reference those files consume only the space in the Snapshot copy reserve.

If Data ONTAP created a Snapshot copy when the disks were full, deleting files from the active file system does not create any free space because everything in the active file system is also referenced by the newly created Snapshot copy. Data ONTAP has to delete the Snapshot copy before it can create any new files.

The following example shows how disk space being freed by deleting files in the active file system ends up in the Snapshot copy:

If Data ONTAP creates a Snapshot copy when the active file system is full and there is still space remaining in the Snapshot reserve, the output from the `df` command—which displays statistics about the amount of disk space on a volume—is as follows:

```
Filesystem              kbytes   used   avail  capacity  Mounted
on      Vserver
/vol/vol0/              3000000 3000000      0      100%
--                   vs1
/vol/vol0/.snapshot    1000000  500000  500000      50%
--                   vs1
```

If you delete 100,000 KB (0.1 GB) of files, the disk space used by these files is no longer part of the active file system, so the space is reassigned to the Snapshot copies instead.

Data ONTAP reassigns 100,000 KB (0.1 GB) of space from the active file system to the Snapshot reserve. Because there was reserve space for Snapshot copies, deleting files from the active file system freed space for new files. If you enter the `df` command again, the output is as follows:

```
Filesystem              kbytes   used   avail  capacity  Mounted
on      Vserver
/vol/vol0/              3000000 2900000  100000      97%
--                   vs1
/vol/vol0/.snapshot    1000000  600000  400000      60%
--                   vs1
```

### Example of what happens when Snapshot copies exceed the reserve

Because there is no way to prevent Snapshot copies from consuming disk space greater than the amount reserved for them, it is important to reserve enough disk space for Snapshot copies so that the active file system always has space available to create new files or modify existing ones.

Consider what happens in the following example if all files in the active file system are deleted. Before the deletion, the `node run -node nodename df` output is as follows:

```
Filesystem      kbytes  used  avail  capacity
/vol/vol0/      3000000 3000000 0      100%
/vol/vol0/./snapshot 1000000 500000 500000 50%
```

After the deletion, the `node run -node nodename df` command generates the following output:

```
Filesystem      kbytes  used  avail  capacity
/vol/vol0/      3000000 2500000 500000 83%
/vol/vol0/./snapshot 1000000 3500000 0      350%
```

The output shows that the entire 3,000,000 KB (3 GB) in the active file system is still being used by Snapshot copies in addition to the 500,000 KB (0.5 GB) that was used by Snapshot copies before the deletion. Therefore, a total of 3,500,000 KB (3.5 GB) is being used by Snapshot copy data, which is 2,500,000 KB (2.5 GB) more than the space reserved for Snapshot copies. This means that 2.5 GB of space that would be available to the active file system is now unavailable to it. The post-deletion output of the `node run -node nodename df` command lists this unavailable space as used even though no files are stored in the active file system.

### Recovery of disk space for file system use

Whenever Snapshot copies consume more than 100% of the Snapshot reserve, they begin to occupy the active file system space. This process is called Snapshot spill. When the Snapshot copies continue to occupy the active file system space, the system is in danger of becoming full. If the system becomes full due to Snapshot spill, you can create files only after you delete enough Snapshot copies.

If 500,000 KB (0.5 GB) of data is added to the active file system, a `node run -node nodename df` command generates the following output:

```
Filesystem      kbytes  used  avail  capacity
/vol/vol0       3000000 3000000 0      100%
/vol/vol0/./snapshot 1000000 3500000 0      350%
```

As soon as Data ONTAP creates a new Snapshot copy, every disk block in the file system is referenced by some Snapshot copy. Therefore, no matter how many files you delete from the

active file system, there is still no room to add any more. The only way to recover from this situation is to delete enough Snapshot copies to free more disk space.

## Working with mirroring technology

Before using mirroring technology, you should understand the components of a mirror relationship, types of mirror copies, where mirror copies are located, path naming and language requirements, and what mirror relationships are not intended to do.

### Components of a mirror relationship

In its simplest configuration, a mirror relationship is between a source volume and a destination volume and data is replicated to the destination volume using Snapshot copies.

Typically, the source volume is a read-write volume that clients can access and modify. The destination volume is a read-only volume that exports a Snapshot copy to clients for read-only access. The only time the source volume is not a read-write volume is in a cascade configuration where the source volume is a destination of one mirror relationship and the source of another mirror relationship.

Snapshot copies are used by the source volume to update destination volumes. Snapshot copies are transferred from the source volume to the destination volume using an automated schedule or manually; therefore, mirror copies are updated asynchronously. You use the set of `snapmirror` commands to create and manage mirror relationships.

#### Related concepts

[Supported data protection deployment configurations](#) on page 34

### Data protection mirror relationships for FlexVol volumes

You can create a mirror relationship to a destination within a cluster to protect your data or, for greater disaster protection, you can create a mirror relationship to a destination in a different cluster in a different location.

A data protection mirror configuration consists of a source volume that can be replicated to one or more destination volumes. Each data protection mirror relationship is independent from the other.

**Note:** The version of Data ONTAP that is running on the destination volume must be the same or a later version than the one running on the source volume.

You can create data protection mirror relationships to destinations on the same aggregate as the source volume, and on the same Storage Virtual Machine (SVM) or on a different SVM. For greater protection, you can create the relationships to destinations on a different aggregate, which enables you to recover from the failure of the source volume's aggregate. Neither of these two configurations protects against a cluster failure.

To protect against a cluster failure, you can create a data protection mirror relationship in which the source volume is on one cluster and the destination volume is on a different cluster. If the cluster on which the source volume resides experiences a disaster, you can direct user clients to the destination volume on the cluster peer until the source volume is available again.

You can also use mirror relationships for limited disaster recovery, off-loading tape backup, data distribution, and making offline copies of production data for research, such as data mining.

## Where to find information about data protection mirror relationships for Infinite Volumes

Information about creating and managing data protection mirror relationships for Infinite Volumes and recovering Infinite Volumes is available in the *Clustered Data ONTAP Infinite Volumes Management Guide*.

## When a destination volume grows automatically

During a data protection mirror transfer, the destination volume grows to ensure the success of the transfer.

At the start of a data protection mirror transfer, the destination volume grows in size if the source volume has grown. This occurs irrespective of any automatic growth setting on the destination volume. The automatic growth of the destination volume occurs as long as there is available space in the aggregate that contains the destination volume. You cannot prevent Data ONTAP from growing or limiting its growth.

## Path name pattern matching

You can use pattern matching when you use `snapmirror` commands to have the command work on selected mirroring relationships.

The `snapmirror` commands use fully qualified path names in the following format: `vserver:volume`. You can abbreviate the path name by not entering the Storage Virtual Machine (SVM) name. If you do this, the `snapmirror` command assumes the local SVM context of the user.

Assuming that the SVM is called “vserver1” and the volume is called “vol1”, the fully qualified path name is `vserver1:vol1`.

You can use the asterisk (\*) in paths as a wildcard to select matching, fully qualified path names. The following table provides examples of using the wildcard to select a range of volumes.

<code>*</code>	Matches all paths.
<code>vs*</code>	Matches all SVMs and volumes with SVM names beginning with <code>vs</code> .
<code>*:*src*</code>	Matches all SVMs with volume names containing the <code>src</code> text.
<code>*:vol*</code>	Matches all SVMs with volume names beginning with <code>vol</code> .



```
vs1::> snapmirror show -destination-path *:*dest*
```

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Healthy	Progress Last Updated
vs1:sm_src2	DP	vs2:sm_dest1	Snapmirrored	Idle	-	true	-

## Language setting requirement

The source and destination FlexVol volumes or Infinite Volumes of a mirror relationship must have the same language setting; otherwise, NFS or CIFS clients might not be able to access data.

For FlexVol volumes, it is not a problem if the source and destination volumes are located on the same Storage Virtual Machine (SVM) because the language is set on the SVM. For FlexVol volumes and Infinite Volumes with mirror relationships between volumes on two different SVMs, the language setting on the SVMs must be the same.

## User access to destination volumes

Users have read-only access to the active file system on the destination FlexVol volume or Infinite Volume. The active file system on the destination volume is an exported Snapshot copy of the active file system from the source volume.

For information about user access to destination Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

## When clients can access the active file system on the destination FlexVol volume

The active file system on a destination volume is available to clients after the system transfers Snapshot copies of the source volume to the destination volume. When the active file system is available differs between FlexVol volumes in mirror relationships and SnapVault relationships.

For a FlexVol volume in a mirror relationship, the storage system automatically directs clients to use the active file system in the latest Snapshot copy on the destination FlexVol volume. For FlexVol volumes that are secondary volumes of a SnapVault relationship, the active file system on the secondary volume is available after the baseline transfer. Attributes of the file system, such as the number of files or the amount of space consumed, are refreshed after the Snapshot copy for the volume is transferred.

## Guidelines for creating relationships between clusters or SVMs

Before you create a mirror or SnapVault relationship between volumes in different clusters or Storage Virtual Machines (SVMs), you should ensure that the relationship follows the supported

configurations. Mirror relationships are supported by FlexVol volumes and Infinite Volumes. SnapVault relationships are supported only by FlexVol volumes.

### **Relationships between volumes in different clusters**

Before you can create a relationship between volumes in different clusters, there must be a cluster peer relationship established between the two clusters and an SVM peer relationship established between the two SVMs.

### **Mirror relationships between clusters running different versions of Data ONTAP**

The version of Data ONTAP that is running on the destination volume must be the same or a later version than the one running on the source volume.

For example, you can create a mirror relationship between a source volume running Data ONTAP 8.1.x and a destination volume running Data ONTAP 8.2.x but not vice versa. Data ONTAP 8.1 commands are supported only for creating and managing these relationships, and you must specify the cluster names.

The `snapmirror show` command displays mixed-version relationships in addition to same-version relationships.

SnapVault relationships are supported only between clusters running Data ONTAP 8.2 or later.

For more information about mirror relationships between clusters running different versions of Data ONTAP, see the *Clustered Data ONTAP Upgrade and Revert/Downgrade Guide*.

### **SnapVault relationships between clusters running different versions of Data ONTAP**

A mixed cluster has at least one node that is running Data ONTAP 8.1.x and other nodes running Data ONTAP 8.2.x. data protection mirror relationships created in Data ONTAP 8.1.x are supported in Data ONTAP 8.2.x, but only the cluster administrator can manage and modify them. Only Data ONTAP 8.1 commands are supported for managing these data protection mirror relationships.

For SnapVault relationships, the version of Data ONTAP that is running on the primary and secondary volumes must be Data ONTAP 8.2 or later. You cannot create a SnapVault relationship with a secondary volume that is running a later version of Data ONTAP than the source volume.

### **Relationships between volumes in different SVMs**

Before you can create a relationship between volumes in different SVMs, there must be a peer relationship established between the two SVMs. You can only establish an SVM peer relationship between SVMs with unique names. You should use unique, fully qualified domain names for each SVM.

You can create a peer relationship between two SVMs with FlexVol volumes or between two SVMs with Infinite Volume, but you cannot create a peer relationship between an SVM with FlexVol volume and an SVM with Infinite Volume.

## Limitations for data protection mirror relationships

When working with data-protection mirror relationships, you should be aware that there are limitations to data protection mirror relationships.

The following limitations apply to data-protection mirror relationships:

- Snapshot copies cannot be deleted on destination volumes.
- An empty junction path on a destination FlexVol volume is not accessible from CIFS clients.
- A volume can have a maximum of 255 Snapshot copies.
- A FlexClone volume cannot be the source of a data-protection mirror relationship.

## Supported number of destination volumes in fanout SnapMirror relationships

When you are planning the number and types of SnapMirror relationships for a single source volume, you should remember that the source volume supports a certain number of destination volumes.

The number of destination volumes you can fan out depends on the type of SnapMirror relationship that you want to fan out from a single source volume:

- For load-sharing mirror relationships, you can fan out a maximum of one destination volume on a node for a single source volume.  
The maximum number of nodes within a cluster depends on the platform model and licensed protocols. For details about cluster size limits, see the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com).
- For data protection mirror relationships, you can fan out a maximum of eight destination volumes from a single source volume.
- A single source volume can have both one load-sharing destination volume on a node and eight data protection destination volumes.

## Snapshot copies cannot be deleted automatically on destination volumes

You cannot automatically delete old Snapshot copies on destination FlexVol volumes or Infinite Volumes of mirror relationships because the destination volume is a read-only version of the source volume and should contain the same data as the source.

This is not true of Snapshot copies on destination FlexVol volumes of SnapVault relationships. You can delete old Snapshot copies on SnapVault secondary volumes.

**Note:** Using the `snap autodelete` command to automatically delete Snapshot copies from a destination volume to remove older Snapshot copies will fail.

## Empty junction path on a destination volume is not accessible from CIFS clients

If internally mounted FlexVol volumes form a namespace and you have a mirror relationship, CIFS clients on a destination volume that attempt to view mirrored volumes not at the highest level of the namespace are denied access.

This occurs when you create a namespace using more than one volume, in which one volume is the source volume of a mirror relationship and the other volumes are members of the namespace. For

example, assume that you have two volumes: vol x, which has a junction path /x, and vol y, which has a junction path /x/y. When a SnapMirror transfer occurs, a directory under vol x is created for vol y on the destination volume. From an NFS client, you can see that the directory is empty, but from a CIFS client, you get the following message:

```
access is denied.
```

## Maximum number of Snapshot copies for volumes that are mirrored

The maximum number of Snapshot copies that a FlexVol volume in a mirror relationship can contain is 251. The maximum number of Snapshot copies that an Infinite Volume in a data protection mirror relationship can contain is 250.

For FlexVol volumes, whenever an update to a data protection mirror copy or set of load-sharing mirror copies occurs, Data ONTAP creates one new Snapshot copy. For Infinite Volumes, whenever an update to a data protection mirror copy occurs, Data ONTAP creates one new Snapshot copy. You should consider this as you manage the number of Snapshot copies on the source volume. You must keep the number of Snapshot copies far enough below the limit that updates to the mirror copy do not exceed the limit.

### Related concepts

[Maximum number of Snapshot copies](#) on page 14

## Working with FlexVol volume SnapVault backups

Before using SnapVault technology, you should understand how SnapVault backups work, where SnapVault volumes are located, and what SnapVault relationships are not intended to do.

### What a SnapVault backup is

A SnapVault backup is a collection of Snapshot copies on a FlexVol volume that you can restore data from if the primary data is not usable. Snapshot copies are created based on a Snapshot policy. The SnapVault backup backs up Snapshot copies based on its schedule and SnapVault policy rules.

A SnapVault backup is a disk-to-disk backup solution that you can also use to offload tape backups. In the event of data loss or corruption on a system, backed-up data can be restored from the SnapVault secondary volume with less downtime and uncertainty than is associated with conventional tape backup and restore operations.

The following terms are used to describe SnapVault backups:

<b>baseline transfer</b>	An initial complete backup of a primary storage volume to a corresponding volume on the secondary system.
--------------------------	---

<b>secondary volume</b>	A volume to which data is backed up from a primary volume. Such a volume can be a secondary or tertiary (and onward) destination in a cascade or fanout backup configuration. The SnapVault secondary system maintains Snapshot copies for long-term storage and possible restore operations.
<b>incremental transfer</b>	A follow-up backup to the secondary system that contains only the changes to the primary data since the last transfer action.
<b>SnapMirror label</b>	An attribute that identifies Snapshot copies for the purpose of selection and retention in SnapVault backups. Each SnapVault policy configures the rules for selecting Snapshot copies on the primary volume and transferring the Snapshot copies that match a given SnapMirror label.
<b>Snapshot copy</b>	<p>The backup images on the source volume that are created manually or automatically as scheduled by an assigned policy. Baseline Snapshot copies contain a copy of the entire source data being protected; subsequent Snapshot copies contain differential copies of the source data. Snapshot copies can be stored on the source volume or on a different destination volume in a different Storage Virtual Machine (SVM) or cluster.</p> <p>Snapshot copies capture the state of volume data on each source system. For SnapVault and mirror relationships, this data is transferred to destination volumes.</p>
<b>primary volume</b>	A volume that contains data that is to be backed up. In cascade or fanout backup deployments, the primary volume is the volume that is backed up to a SnapVault backup, regardless of where in the chain the SnapVault source is. In a cascade chain configuration in which A has a mirror relationship to B and B has a SnapVault relationship to C, B serves as the source for the SnapVault backup even though it is a secondary destination in the chain.
<b>SnapVault relationship</b>	A backup relationship, configured as a SnapVault relationship, between a primary volume and a secondary volume.

## Which data gets backed up and restored from a SnapVault backup

You create SnapVault relationships to back up and restore volumes. You can select the Snapshot copies that the SnapVault relationship uses to backup and restore volumes.

The SnapVault operation backs up a specified volume on the primary system to the associated volume on the SnapVault secondary system. If necessary, data is restored from the SnapVault secondary volume back to the associated primary volume or to a different volume.

The Snapshot policy assigned to the source volume specifies when Snapshot copies are performed. The SnapVault policy assigned to the SnapVault relationship specifies which of the source volume Snapshot copies are replicated to the SnapVault backup.

If the destination volume is a FlexClone volume, the volume retains two more Snapshot copies than the number you configure in the policy. This occurs because the volume retains the FlexClone Snapshot copy and an exported Snapshot copy. For example, if your policy specifies to retain three

Snapshot copies, five Snapshot copies are retained (three specified Snapshot copies, one FlexClone Snapshot copy, and one exported Snapshot copy).

In SAN environments, LUN identifiers are preserved on the SnapVault secondary volume.

The secondary system uses slightly more disk space and directories than the source system.

### Related concepts

*Which data does not get backed up to a SnapVault backup* on page 30

*Guidelines for restoring the active file system* on page 98

## Which data does not get backed up to a SnapVault backup

If you back up an entire Storage Virtual Machine (SVM) to a SnapVault backup by establishing a SnapVault relationship for each volume in the SVM, namespace and root information is not backed up. To protect namespace and root information for an SVM, you must manually create the namespace and root on the SnapVault secondary volume. When backing up LUNs to a SnapVault secondary volume, not all LUN information is replicated.

In SAN environments, the following LUN attributes are not replicated to the secondary volume:

- Path  
The LUN in the SnapVault secondary volume can be in a different SVM or volume from the source LUN. Path-related metadata, such as persistent reservations, are not replicated to the SnapVault primary volume.
- Serial number
- Device ID
- UUID
- Mapped status
- Read Only state  
The Read Only state is always set to `true` on the destination LUN.
- NVFAIL attribute  
The NVFAIL attribute is always set to `false` on the destination LUN.

You can set persistent reservations for LUNs on the SnapVault secondary volume.

### Related concepts

*Which data gets backed up and restored from a SnapVault backup* on page 29

## How a SnapVault backup works

Backing up volumes to a SnapVault backup involves starting the baseline transfers, making scheduled incremental transfers, and restoring data upon request.

### Baseline transfers

In general, baseline transfers work as follows:

A baseline transfer occurs when you initialize the SnapVault relationship. When you do this, Data ONTAP creates a new Snapshot copy. Data ONTAP transfers the Snapshot copy from the primary volume to the secondary volume. This Snapshot copy is the baseline of the volume at the time of the transfer and is a complete transfer, not an incremental transfer. As a result, none of the other Snapshot copies on the primary volume are transferred as part of the initial SnapVault transfer, regardless of whether they match rules specified in the SnapVault policy.

### Incremental transfers

The source system creates incremental Snapshot copies of the source volume as specified by the Snapshot policy that is assigned to the primary volume. Each Snapshot copy for a specific volume contains a label that is used to identify it.

The SnapVault secondary system selects and retrieves specifically labeled incremental Snapshot copies, according to the rules that are configured for the SnapVault policy that is assigned to the SnapVault relationship. The Snapshot label is retained to identify the backup Snapshot copies.

Snapshot copies are retained in the SnapVault backup for as long as is needed to meet your data protection requirements. The SnapVault relationship does not configure a retention schedule, but the SnapVault policy does specify number of Snapshot copies to retain.

### SnapVault backup updates

At the end of each Snapshot copy transfer session, which can include transferring multiple Snapshot copies, the most recent incremental Snapshot copy in the SnapVault backup is used to establish a new common base between the primary and secondary volumes and is exported as the active file system.

### Data restore

If data needs to be restored to the primary volume or to a new volume, the SnapVault secondary transfers the specified data from the SnapVault backup.

## How SnapVault backups work with data compression

SnapVault relationships preserve storage efficiency when replicating data from the source to the SnapVault secondary volume except when additional data compression is enabled.

If additional compression is enabled on the SnapVault secondary volume, storage efficiency is affected as follows:

- Storage efficiency is not preserved for data transfers between the primary and secondary volumes.
- You do not have the option of returning to replicating data while preserving storage efficiency.

### Related tasks

[Managing storage efficiency for SnapVault secondaries](#) on page 104

## SnapVault backup limitations

When planning SnapVault relationships, you must keep in mind what is supported and what is not supported.

The following limitations apply to SnapVault backups:

- 32-bit aggregates are not supported.  
Clustered Data ONTAP systems do not support the SnapVault backup feature for volumes in 32-bit aggregates.
- A SnapVault secondary volume cannot be the secondary volume for multiple primary volumes. A volume can be the secondary for one SnapVault relationship only. However, that same volume can be the source for other relationships.
- SnapVault backups are not supported on Infinite Volumes.

## Data protection for Storage Virtual Machine (SVM) namespace and root information

Backups to secondary volumes in SnapVault relationships between FlexVol volumes replicate only volume data, not the SVM namespace or root information.

SnapVault relationships replicate only volume data. If you want to back up an entire SVM to a SnapVault secondary volume, you must first create SnapVault relationships for every volume in the SVM.

To provide data protection of the SVM namespace information, you must manually create the namespace on the SnapVault secondary immediately after the first data transfer is completed for all of the volumes in the SVM and while the source SVM volumes are still active. When subsequent changes are made to the namespace on the source SVM, you must manually update the namespace on the destination SVM.



You cannot create the namespace for an SVM on a SnapVault secondary volume if only a subset of the SVM volumes are in a SnapVault relationship, or if only a subset of the SVM volumes have completed the first data transfer.

## Guidelines for planning Snapshot copy schedule and retention for SnapVault backups

It is important to plan the Snapshot copy transfer schedule and retention for your SnapVault backups.

When planning SnapVault relationships, consider the following guidelines:

- Before you create a SnapVault policy, you should create a table to plan which Snapshot copies you want replicated to the SnapVault secondary volume and how many of each you want to keep. For example:
  - Hourly (periodically throughout the day)
 

Does the data change often enough throughout the day to make it worthwhile to replicate a Snapshot copy every hour, every two hours, or every four hours?
  - Nightly
 

Do you want to replicate a Snapshot copy every night or just workday nights?
  - Weekly
 

How many weekly Snapshot copies is it useful to keep in the SnapVault secondary volume?
- The primary volume should have an assigned Snapshot policy that creates Snapshot copies at the intervals you need, and labels each Snapshot copy with the appropriate `snapmirror-label` attribute name.
- The SnapVault policy assigned to the SnapVault relationship should select the Snapshot copies you want from the primary volume, identified by the `snapmirror-label` attribute name, and specify how many Snapshot copies of each name you want to keep on the SnapVault secondary volume.

<b>Sample transfer schedule and retention</b>			
<b>snapmirror-label attribute value</b>	<b>Source volume: Snapshot copy schedule</b>	<b>Primary volume: Snapshot copies retained</b>	<b>SnapVault secondary volume: Snapshot copies retained</b>
weekly	Every Saturday at 19:00	4	8
nightly	Every Monday through Friday at 19:00	10	60
hourly	Every hour from 07:00 through 18:00	11	120

snapmirror-label attribute value	Source volume: Snapshot copy schedule	Primary volume: Snapshot copies retained	SnapVault secondary volume: Snapshot copies retained
Total	n/a	25	188

## Supported data protection deployment configurations

A simple data protection deployment consists of a FlexVol volume or Infinite Volume in a single mirror relationship or a FlexVol volume in a SnapVault relationship. More complex deployment configurations that provide additional data protection consist of a cascade chain of relationships between FlexVol volumes or a set of fanout relationships for a FlexVol volume or Infinite Volume.

Although a single volume-to-volume relationship does provide data protection, your data protection needs might require the additional protection that is provided by more complex cascade and fanout configurations.

An example of a cascade chain is an A to B to C configuration. In this example, A is the source that is replicated to B as a data protection mirror, and B is the primary that is backed up to C as a SnapVault backup. Cascade chains can be more complex than A to B to C, but the more relationships that are involved in the chain, the greater the number of temporary locks on volumes while replication or update operations are in progress.

An example of a fanout is an A to B and A to C backup or mirror replication configuration. In this example, A is the primary source that is replicated to both B (either in a mirror or SnapVault relationship) and C.

**Note:** Only one SnapVault relationship is supported in a cascade chain configuration, but many SnapVault relationships are supported in a fanout configuration; multiple mirror relationships are supported.

**Attention:** The longer you configure a chain of relationships or the more you add fanout destinations, the greater the risk of Snapshot copies being locked on the source. Depending on the update schedule, the worst case is when one Snapshot copy is locked for each cascade or fanout destination.

The types of supported deployment configurations are as follows:

- Basic data protection configuration (for FlexVol volumes and Infinite Volumes)  
A FlexVol volume or Infinite Volume is in a single relationship with another volume as the source or the destination of mirror replication operations, or a FlexVol volume is in a single relationship with another volume as the primary or the secondary of SnapVault operations
- Cascade (one-to-one-to-one relationship)  
The three types of cascade chain relationships that you can configure are as follows:

- **Mirror-mirror cascade (for FlexVol volumes only)**  
A chain of at least two mirror relationships in which a volume is the source for replication operations to a secondary volume, and the secondary volume is the source for replication operations to a tertiary volume. This configuration might be described as follows: A mirror to B mirror to C.
- **Mirror-SnapVault cascade (for FlexVol volumes only)**  
A chain of a mirror relationship followed by a SnapVault relationship in which a volume is the source for replication operations to a secondary volume, and the secondary volume is the primary for SnapVault operations to a tertiary volume. This configuration might be described as follows: A mirror to B SnapVault backup to C.
- **SnapVault-mirror cascade (for FlexVol volumes only)**  
A chain of a SnapVault relationship followed by a mirror relationship in which a volume is the primary for SnapVault operations to a secondary volume, and the secondary volume is the source for replication operations to a tertiary volume. This configuration might be described as follows: A SnapVault backup to B mirror to C.

A load-sharing mirror source volume or destination volume cannot be a part of any cascade relationship. See the *Clustered Data ONTAP Logical Storage Management Guide* for information about load-sharing mirror relationships.

- **Fanout (one-to-many relationship)**  
In a fanout relationship structure, the source is replicated to multiple destinations, which can be mirror or SnapVault destinations. Only one SnapVault relationship is allowed in a fanout.
  - **Mirror-SnapVault fanout (for FlexVol volumes only)**  
A volume is the source for replication operations to a secondary volume and also the source for SnapVault operations to a different secondary volume. This configuration might be described as follows: A mirror to B and A also SnapVault backup to C.
  - **Multiple-mirrors fanout (for FlexVol volumes and Infinite Volumes)**  
A volume is the source for replication operations to a destination volume and also the source for replication operations to another, different destination volume. This configuration might be described as follows: A mirror to B and A also mirror to C.

## Related concepts

[Components of a mirror relationship](#) on page 23

[What source-to-destination-to-tape backup is](#) on page 36

[How a mirror-mirror cascade works](#) on page 37

[How a mirror-SnapVault cascade works](#) on page 38

[How a SnapVault-SnapMirror cascade works](#) on page 39

[How a mirror-SnapVault fanout works](#) on page 39

[How a multiple-mirrors fanout works](#) on page 40


## What a basic backup deployment is

A basic data protection deployment consists of two volumes, either FlexVol volumes or Infinite Volumes, in a one-to-one, source-to-destination relationship. This deployment backs up data to one location, which provides a minimal level of data protection.

In a data protection configuration, source volumes are the data objects that need to be replicated. Typically, users can access and write to source volumes.

Destination volumes are data objects to which the source volumes are replicated. Destination volumes are read-only. Destination FlexVol volumes are usually placed on a different Storage Virtual Machine (SVM) from the source SVM. Destination Infinite Volumes must be placed on a different SVM from the source SVM. Destination volumes can be accessed by users in case the source becomes unavailable. The administrator can use SnapMirror commands to make the replicated data at the destination accessible and writable.

The following illustration depicts a basic data protection deployment:



## Firewall requirements for intercluster SnapMirror relationships


SnapMirror relationships that have source volumes on one cluster and destination volumes on another cluster require certain ports on the intercluster network.

SnapMirror relationships use port 11104 and port 11105 on the intercluster network to replicate data from their source volumes to their destination volumes. Clustered Data ONTAP uses port 11104 to manage intercluster communication sessions and uses port 11105 to transfer data.

## What source-to-destination-to-tape backup is

A common variation of the basic data protection backup deployment adds a tape backup of a destination FlexVol volume. By backing up to tape from the destination volume, you do not subject the heavily accessed source volume to the performance degradation and complexity of a direct tape backup.

The following illustration depicts a data protection chain deployment with a tape backup:



NDMP is required for this configuration, and Infinite Volumes do not support NDMP. However, you can use other methods to create a tape backup of an Infinite Volume. For more information, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

## How a mirror-mirror cascade works


A mirror-mirror cascade deployment is supported on FlexVol volumes and consists of a chain of mirror relationships in which a volume is replicated to a secondary volume and the secondary is replicated to a tertiary volume. This deployment adds one or more additional backup destinations without degrading performance on the source volume.

By replicating source A (as shown in the following illustration) to two different volumes (B and C) in a series of mirror relationships in a cascade chain, you create an additional backup. The base for the B-to-C relationship is always locked on A to ensure that the backup data in B and C always stay synchronized with the source data in A.

If the base Snapshot copy for the B-to-C relationship is deleted from A, the next update operation from A to B fails and an error message is generated that instructs you to force an update from B to C. The forced update establishes a new base Snapshot copy and releases the lock, which enables subsequent updates from A to B to finish successfully.

If the volume on B becomes unavailable, you can synchronize the relationship between C and A to continue protection of A without performing a new baseline transfer. After the resynchronize operation finishes, A is in a direct mirror relationship with C, bypassing B.

The following illustration depicts a mirror-mirror cascade chain:



## How a mirror-SnapVault cascade works

A mirror-SnapVault cascade deployment, which is supported on FlexVol volumes, consists of a chain of relationships in which a volume is replicated to a destination volume, and then the destination volume becomes the primary for a SnapVault backup on a tertiary volume. This deployment adds a SnapVault backup, which fulfills more strict protection requirements.


In a typical mirror-SnapVault cascade, only the exported Snapshot copies from the mirror destination are transferred to the SnapVault secondary when the SnapVault update occurs. These exported Snapshot copies are created by Data ONTAP and have a prefix of “snapmirror” and a “sm\_created” SnapMirror label.

If the default SnapVault policy is used, the SnapVault destination will accumulate up to 251 “sm\_created” Snapshot copies. After this limit is reached, when a newer “sm\_created” Snapshot copy is transferred, the oldest one is rotated out. This retention and rotation behavior can be managed by adding a rule for the “sm\_created” SnapMirror label to the SnapVault policy.

For example, if a rule is added with a `-snapmirror-label` of “sm\_created” and with a `-keep` value of 40, then only 40 “sm\_created” Snapshot copies are retained on the SnapVault destination. If the `-preserve` value for this rule is set to `true`, then no rotation will occur and “sm\_created” Snapshot copy transfers will halt when the SnapVault destination reaches a count of 40 “sm\_created” Snapshot copies. If the `-preserve` value for this rule is set to `false`, then “sm\_created” Snapshot copy transfers will occur after 40 Snapshot copies with the oldest copy rotating out for the newest copy.

**Note:** A cascade chain can contain multiple mirror relationships but only one SnapVault relationship. The SnapVault relationship can occur anywhere in the chain, depending on your data protection requirements.

The following illustration depicts a mirror-SnapVault cascade chain:



### Related references


[Creating the SnapVault relationship of a mirror-SnapVault cascade](#) on page 87

## How a SnapVault-SnapMirror cascade works

A SnapVault-SnapMirror cascade consists of a chain of relationships in which a volume has a SnapVault backup on a secondary volume, and then that secondary volume data is replicated to a tertiary volume. In effect, this deployment provides two SnapVault backups.

A SnapVault-SnapMirror cascade deployment is only supported on FlexVol volumes. The first leg of the cascade consists of a SnapVault backup. A cascade chain in which the first leg is a SnapVault relationship behaves in the same manner as does a single leg SnapVault relationship. The updates to the SnapVault backup include the Snapshot copies that are selected in conformance with the SnapVault policy assigned to the relationship. In a typical SnapVault-SnapMirror cascade, all Snapshot copies up to the latest one are replicated from the SnapVault backup to the SnapMirror destination.

The following illustration depicts a SnapVault-to-SnapMirror cascade chain:



### Related concepts


[Managing a SnapVault-mirror cascade when the SnapVault backup is unavailable](#) on page 101

## How a mirror-SnapVault fanout works

A mirror-SnapVault fanout deployment is supported on FlexVol volumes and consists of a source volume that has a direct mirror relationship to a secondary volume and also a direct SnapVault relationship to a different secondary volume.

**Note:** A fanout deployment might not provide as much data protection as a cascade chain.

The following illustration depicts a mirror and SnapVault fanout:



## How a multiple-mirrors fanout works


A multiple-mirrors fanout deployment is supported on FlexVol volumes and Infinite Volumes, and consists of a source volume that has a direct mirror relationship to multiple secondary volumes.

The volume on A (as shown in the following illustration) always contains the base Snapshot copies for both B and C. Because updates to B or C automatically include the base Snapshot copy of the other relationship, B and C always have a common Snapshot copy.

**Note:** A fanout deployment might not provide as much data protection as a cascade chain.



The following illustration depicts a mirror and mirror fanout:



## Protecting data using Snapshot copies

You can use Snapshot copies to restore data that is lost because of accidental deletion to FlexVol volumes and Infinite Volumes.

Data ONTAP maintains a configurable Snapshot schedule that creates and deletes Snapshot copies automatically for each volume. You can also create and delete Snapshot copies, and manage Snapshot schedules based on your requirements.

If you lose data due to a disaster, you use data protection mirror copies to restore data.

## Managing Snapshot copies

You can create multiple schedules that create and delete Snapshot copies, as desired.

### Commands for managing Snapshot copies

Cluster administrators can use the `volume snapshot` commands to create and manage all Snapshot copies. Storage Virtual Machine (SVM) administrators can use the same commands to create and manage Snapshot copies within SVMs.

If you want to...	Use this command...
Display information about Snapshot copies	<code>volume snapshot show</code>
Create a Snapshot copy of a volume If you are using Infinite Volumes, you must ensure that the Infinite Volume is in an online state. You cannot create a Snapshot copy if the Infinite Volume is in a mixed state due to a offline constituent.	<code>volume snapshot create</code>
Modify the attributes of a Snapshot copy If you are using Infinite Volumes, you cannot change the comment or name associated with a Snapshot copy of an Infinite Volume.	<code>volume snapshot modify</code>
Rename a Snapshot copy of a FlexVol volume You cannot rename a Snapshot copy that is created as a reference copy during execution of the <code>volume copy</code> or <code>volume move</code> commands. If you are using Infinite Volumes, you cannot rename a Snapshot copy of an Infinite Volume.	<code>volume snapshot rename</code>

If you want to...	Use this command...
<p>Delete a Snapshot copy</p> <p>If you are using Infinite Volumes, the Infinite Volume must be online. You cannot delete a Snapshot copy of an Infinite Volume when the Infinite Volume is in a mixed state without assistance from technical support.</p> <p>If you are using SnapMirror, base Snapshot copies must exist and at least one common Snapshot copy must exist between the source and destination volume to use the <code>snapmirror resync</code> command.</p>	<pre>volume snapshot delete</pre>

See the man page for each command for more information.

## Managing Snapshot policies

Snapshot policies automatically manage Snapshot copy schedules and retention on FlexVol volumes or Infinite Volumes. You must be a cluster administrator or Storage Virtual Machine (SVM) administrator to perform most of the Snapshot policy commands.

### How Snapshot policies are associated with volumes

Unless you specify a Snapshot policy when you create a FlexVol volume or an Infinite Volume, a volume inherits the Snapshot policy associated with its containing Storage Virtual Machine (SVM).

When you create the SVM, you can specify a Snapshot policy. If you do not specify a Snapshot policy when you create the SVM, a default Snapshot policy is associated with the SVM. The default Snapshot policy for an SVM with FlexVol volume is named `default`, and the default Snapshot policy for an SVM with Infinite Volume is named `default-1weekly`.

**Note:** When you upgrade the SVM with Infinite Volume from Data ONTAP 8.1.x, the default Snapshot policy changes from `default` to `default-1weekly`.

When you create a volume, you can specify a Snapshot policy. If you do not specify a Snapshot policy when you create a volume, the volume inherits the Snapshot policy associated with its containing SVM.

**Note:** A Snapshot policy is not associated with each constituent in an Infinite Volume, and you cannot associate a Snapshot policy with constituents. A Snapshot policy is only associated with an Infinite Volume.

## Commands for managing Snapshot policies and schedules

Cluster administrators can use the `volume snapshot policy` commands to create and manage all Snapshot copy policies. Storage Virtual Machine (SVM) administrators can use the same commands to create and manage Snapshot policies within SVMs.

If you want to...	Use this command...
Display information about Snapshot copy policies	<code>volume snapshot policy show</code>
Create a new Snapshot copy policy	<code>volume snapshot policy create</code>
Create a schedule that can be used in Snapshot copy policies	<code>job schedule cron create</code>
Add a schedule to an existing Snapshot copy policy A Snapshot policy can have up to five schedules. If you are using Infinite Volumes, scheduled Snapshot copies cannot occur more often than hourly. Cluster administrator only	<code>volume snapshot policy add-schedule</code>
Remove a schedule from a Snapshot copy policy Cluster administrator only	<code>volume snapshot policy remove-schedule</code>
Modify the maximum number of Snapshot copies for a Snapshot copy policy schedule Cluster administrator only	<code>volume snapshot policy modify-schedule</code>
Modify the description of a Snapshot copy policy	<code>volume snapshot policy modify</code>
Dissociate a Snapshot copy policy from a volume Cluster administrator only	<code>volume modify</code>
Delete a Snapshot copy policy Cluster administrator only	<code>volume snapshot policy delete</code>

See the man page for each command for more information.

## Strategies for creating a Snapshot copy policy

You should create a Snapshot copy policy that meets the needs of your organization and users.

Following are some strategies for using policies and schedules to schedule and retain Snapshot copies:

- If users rarely lose files or typically notice lost files right away, you can use the default Snapshot copy policy.  
This policy uses the weekly schedule to create two weekly Snapshot copies; the daily schedule to create a Snapshot copy every day and keeps two; and the hourly schedule to create hourly Snapshot copies and keeps six.
- If users commonly lose files or do not typically notice lost files right away, you should delete the Snapshot copies less often than you would if you used the default policy.

Following is the recommended Snapshot copy policy for this situation. It uses the weekly schedule to keep two weekly Snapshot copies, the daily schedule to keep six daily Snapshot copies, and the hourly schedule to keep eight hourly Snapshot copies:

```
snapshot policy create -vserver vs1.example.com -policy keep-more-
snapshot -enabled true -schedule1 weekly -count1 2 -prefix1 weekly -
schedule2 daily -count2 6 -prefix2 daily -schedule3 hourly -count3 8 -
prefix3 hourly
```

On many systems, only 5 to 10 percent of the data changes each week, so the Snapshot copy schedule of six daily and two weekly Snapshot copies consumes 10 to 20 percent of disk space. Considering the benefits of Snapshot copies, it is worthwhile to reserve this amount of disk space for Snapshot copies.

- You can create different Snapshot copy policies for different volumes on a Vserver.  
On a very active volume, you should schedule Snapshot copies every hour and keep them for just a few hours, or turn off Snapshot copies. For example, the following schedule creates a Snapshot copy every hour and keeps the last three:

```
snapshot policy create -vserver vs1.example.com -policy hourly-keep-3 -
enabled true -schedule1 hourly -count1 3 -prefix1 hourly
```

- When you create a new volume, the new volume inherits the Snapshot copy schedule from the root volume.

After you use the volume for a while, you should check how much disk space the Snapshot copies consume and how often users need to recover lost files, and then adjust the schedule as necessary.

## Naming convention for scheduled Snapshot copies

The scheduled Snapshot copy name is composed of an optional prefix or the schedule name specified in the Snapshot policy, and the timestamp. Snapshot names cannot be longer than 255 characters.

If prefix is specified, the Snapshot name is made up of prefix and the timestamp.

If you do not specify the prefix, by default, the schedule name is prepended with the timestamp to form a Snapshot name.

## What prefixes are

A prefix is an optional string that you can specify to be used in creating automatic Snapshot copies. Using prefixes in Snapshot names provides more flexibility than using schedule names in naming automatic Snapshot copies.

Prefix names must be unique within a policy. The length of the prefix cannot exceed the maximum allowable length for Snapshot names; that is, Snapshot names cannot be longer than 255 characters. Prefix names must follow the character encoding rules used by Snapshot names.

If a prefix is specified in the Snapshot schedule, then the schedule name is not used to name Snapshot copies. If the prefix is not specified for a Snapshot schedule within a Snapshot policy, then the schedule name is used.

## Using prefixes to name automatic Snapshot copies

You can use prefixes to provide flexibility to the naming convention for scheduled Snapshot copies. It removes the dependency on using the schedule names in naming scheduled Snapshot copies.

### About this task

- A schedule cannot have more than one prefix.
- Prefixes within a policy should be unique.

### Step

1. You can specify prefixes when you create a Snapshot policy or when you add a schedule to the Snapshot policy.

### Example

The following command creates a Snapshot policy “test”, which contains the schedule named “5min” having the temp prefix:

```
cluster1::> volume snapshot policy create -policy test -enabled true  
-schedule1 5min -count1 2 -prefix1 temp
```

### Example

The following command adds the “6min” schedule with the “test” prefix to the default policy:

```
cluster1::> volume snapshot policy add-schedule -policy default  
-schedule 6min -count 4 -prefix test
```

## Restoring files from the Snapshot copy of a FlexVol volume

You might have to restore a file from the Snapshot copy of a FlexVol volume if the file was accidentally erased or corrupted. You can use the SnapRestore feature to automatically restore files from the Snapshot copy of a FlexVol volume.

### Steps

1. If the original file still exists and you do not want it overwritten by the file in a Snapshot copy, then use your UNIX or Windows client to rename the original file or move it to a different directory.
2. Locate the Snapshot copy that contains the version of the file that you want to restore.
3. Copy the file from the `.snapshot` directory to the directory in which the file originally existed.

## Restoring a single file from a Snapshot copy of a FlexVol volume

You can restore a single file to the required version from a Snapshot copy of a FlexVol volume.

### Before you begin

- The volume to which you want to restore the file should be online and writeable.
- The volume to which you want to restore the file should have enough space for the restore operation to be completed successfully.

### About this task

The restored file can replace an existing file with the same name in the active file system or become a new file if there is data in the existing file that you want to retain. You can also restore LUNs, but you cannot restore a single file from a Snapshot copy of an Infinite Volume.

If you are restoring an existing LUN, a LUN clone is created and is backed up in the form of a Snapshot copy. During the restore operation, you can read to and write from the LUN.

### Step

1. To restore a single file or LUN, use the `volume snapshot restore-file` command.

The restore operation might take a long time, depending on the size of the file or LUN that you are restoring.

If you want to display the number of in-progress single file restore operations, use the `volume snapshot restore-file-info` command.

## Restoring part of a file from a Snapshot copy of a FlexVol volume

You can restore a range of data from a file in a Snapshot copy to an existing file in the active file system. Partial file restores can only be used to restore specific pieces of a LUN, and NFS or CIFS container files.

### Before you begin

- You must understand the metadata of the host LUN or container file so that you know which bytes belong to the object that you want to restore.
- Write operations are not allowed on the object that you are restoring. Otherwise, it might result in an inconsistent data.
- The volume where the LUN or the container file is to be restored must be online and writable.

### Steps

1. To restore part of a file, use the `volume snapshot partial-restore-file` command.

To get the settings for partial file restore on a cluster, use the `volume snapshot partial-restore-file-list-info` command.

After the restore is complete, you must purge operating system or application buffers so that the stale data is cleaned.

2. After the restore is complete, purge operating system or application buffers so that the stale data is cleaned.

## Restoring the contents of a volume from a Snapshot copy

You can restore the contents of a FlexVol volume or Infinite Volume from a Snapshot copy to quickly recover lost or damaged data.

### Before you begin

- You must have the advanced privilege level or higher to run the command.
- If you are working with a Snapshot copy of an Infinite Volume, the Snapshot copy must be valid and the Infinite Volume must be online.
- You must not have any I/O traffic running on the volume.

### Steps

1. If the volume is an Infinite Volume, use the `volume unmount` command to unmount it.
2. Use the `volume snapshot restore` command to restore the contents of a volume from a Snapshot copy.



### Example

The following command restores data to a volume named `src_os` from a Snapshot copy named `src_os_snap_3` on the Storage Virtual Machine (SVM) named `vs0`:

```
vs1::*> volume snapshot restore -vserver vs0.example.com  
-volume src_os -snapshot src_os_snap_3
```

3. If the volume is an Infinite Volume, use the `volume mount` command to mount it.
4. If the volume has SnapMirror relationships, manually replicate all mirror copies of the volume immediately after you restore from a Snapshot copy.

Not doing so can result in unusable mirror copies that must be deleted and re-created.

## Snapshot restoration using Shadow Copy Client tools

You can access and restore Data ONTAP Snapshot files using the Windows Shadow Copy Client. The Shadow Copy Client provides a Previous Versions tab in the Properties menu from which you can view and restore Data ONTAP Snapshot images.

The Shadow Copy Client software for Windows 2003 is called the Previous Versions Client. Downloads available from Microsoft allow you to use Shadow Copy client tools on most older versions of Windows. For more information about Shadow Copy Client or Previous Versions Client software, consult the Microsoft documentation.

## Managing Snapshot copy disk space

The data referenced by a Snapshot copy cannot be accidentally deleted because of the Snapshot feature's design.

## Monitoring Snapshot copy disk consumption

You can monitor Snapshot copy disk consumption using the `df` command, which displays the amount of free space on a disk.

### About this task

For an Infinite Volume, the `df` command displays information about all of the data constituents, not about the Infinite Volume as a whole.

### Step

1. To display information about Snapshot copy disk consumption, use the `df` command.

**Example**

The `df` command treats Snapshot copies as a partition different from the active file system. The following example shows a volume with these characteristics:

- The total volume capacity (`kbytes` column) is 4,000,000 KB (4 GB): 3,000,000 KB (75 percent) for the active file system, and 1,000,000 KB (25 percent) for Snapshot copies.
- The active file system is using 2,000,000 KB of its 3,000,000 KB capacity (66 percent, rounded to 65 percent in the `capacity` column), leaving 1,000,000 KB (34 percent) available.
- Snapshot copies are using 500,000 KB of their 1,000,000 KB capacity (50 percent in the `capacity` column), leaving 500,000 KB (50 percent of the space allotted for Snapshot copies, not 50 percent of disk space) available.

**Note:** It is important to understand that the `/vol/vol0/.snapshot` line counts data that exists only in a Snapshot copy. The Snapshot copy calculation does not include Snapshot copy data that is shared with the active file system.

```
cluster1:> df
Filesystem          kbytes    used    avail    capacity  Mounted on    Vserver
/vol/vol0/          3000000  2000000  1000000    65%    ---          vs1
/vol/vol0/.snapshot 1000000   500000   500000    50%    ---          vs1
```

## Managing data protection using SnapMirror policies

---

To manage a data protection mirror or SnapVault relationship, you must assign a policy to the relationship. You use the policy to maximize the efficiency of the transfers to the backup secondaries and manage the update operations for SnapVault backups.

FlexVol volumes support data protection mirror and SnapVault relationships and policies. Infinite Volumes support only data protection mirror relationships and policies.

### Commands for managing mirror and SnapVault policies

Cluster administrators can use the `snapmirror policy` commands to create and manage all data protection mirror and SnapVault policies. Storage Virtual Machine (SVM) administrators can use the same commands to create and manage all data protection mirror and SnapVault policies within SVMs.

- All policy-management commands (except for the `snapmirror policy show` command) must be run on the SVM that contains the destination volume.
- Commands for SnapVault policies are supported only by FlexVol volumes.

If you want to...	Use this command...
Add a new rule to a SnapVault policy	<code>snapmirror policy add-rule</code>
Create a new mirror or SnapVault policy	<code>snapmirror policy create</code>
Delete a mirror or SnapVault policy	<code>snapmirror policy delete</code>
Modify a mirror or SnapVault policy	<code>snapmirror policy modify</code>
Modify an existing rule in a SnapVault policy	<code>snapmirror policy modify-rule</code>
Remove a rule in a SnapVault policy	<code>snapmirror policy remove-rule</code>
Display a list of mirror and SnapVault policies	<code>snapmirror policy show</code>

See the man page for each command for more information.

## How SnapMirror policies work with clusters and Storage Virtual Machines (SVMs)

A SnapMirror policy in which the `vserver` parameter contains the cluster name is a cluster-wide policy. You can assign cluster-wide SnapMirror policies to relationships in a cluster. Cluster-wide policies can be configured only by a cluster administrator.

A SnapMirror policy in which the `vserver` parameter contains an SVM name is an SVM-wide policy. You can assign SVM-wide SnapMirror policies to relationships within the SVM in which the policy was created. SVM policies can be configured by either a cluster administrator or an SVM administrator.

## Comparison of what cluster administrators and Storage Virtual Machine (SVM) administrators can manage

Cluster administrators and SVM administrators have different privileges for creating, managing, and assigning policies to mirror and SnapVault relationships.

Cluster administrators can do the following:

- Create and manage any policy in a cluster or SVM  
For cluster-wide policies, the `vserver` parameter contains the cluster name. For SVM-wide policies, the `vserver` parameter contains the SVM name.
- View, modify, or delete policies in a cluster or SVM
- Assign a cluster-wide or SVM-wide policy to a relationship

SVM administrators can do the following:

- Create and manage policies within an SVM  
Policies created by an SVM administrator are automatically configured with the SVM name in the `vserver` parameter.
- View cluster-wide policies and SVM-wide policies created in a specified SVM  
Although SVM administrators can view cluster-wide policies, they cannot modify or delete them. SVM administrators can view only the SVM-wide policies that were created within the SVM on which the `snapmirror policy show` command is executed.
- Assign a cluster-wide or SVM-wide policy to a relationship

SVM administrators cannot access SVM-wide policies of another SVM.

## Guidelines for naming Snapmirror and SnapVault policies

Before you create a SnapMirror or SnapVault policy, you should ensure that the policy name is unique.

Cluster-wide policy names must be unique within the cluster and must not conflict with any Storage Virtual Machine (SVM)-wide policy names.

SVM-wide policy names must be unique within the SVM in which the policy is created. However, an SVM policy name can be the same as a policy name created in a different SVM, as long as the name does not conflict with any cluster-wide policy name.

## Preserving Snapshot copies for SnapVault relationships after reaching retention limit

After the Snapshot copy retention limit defined by a SnapMirror policy for a SnapVault relationship is reached, the oldest Snapshot copy is automatically deleted to create space before transferring a new Snapshot copy. You can configure or modify the policy rule to preserve all Snapshot copies.

### About this task

You can configure or modify the policy rule to preserve all Snapshot copies when you create the SnapMirror policy rule, or you can modify a previously created SnapMirror policy. Configuring or modifying the policy rule to preserve Snapshot copies causes incremental updates to the SnapVault secondary to fail after Snapshot copies reach the retention count.

### Steps

1. Configure or modify the policy rule to preserve Snapshot copies by using the `snapmirror policy add-rule` command or the `snapmirror policy modify-rule` command with the `-preserve` parameter.

### Example

The following example configures the XDPDefault policy rule to preserve the 40 `sm_created` Snapshot copies that you want to retain.

```
cluster1::> snapmirror policy add-rule -vserver vs1 -policy
XDPDefault
-snapmirror-label sm_created -keep 40 -preserve true
```

**Example**

The following example modifies the XDPDefault policy rule to preserve the 40 sm\_created Snapshot copies that you want to retain.

```
cluster1::> snapmirror policy modify-rule -vserver vs1 -policy
XDPDefault
-snapmirror-label sm_created -preserve true
```

2. Optionally, check the policy rules to ensure that you enabled the `-preserve` parameter by using the `snapmirror policy show` command with the `-instance` parameter:

**Example**

```
cluster1::> snapmirror policy show -instance

          Vserver: vs1
SnapMirror Policy Name: XDPDefault
          Policy Owner: cluster-admin
            Tries Limit: 8
      Transfer Priority: normal
Ignore accesstime Enabled: false
  Transfer Restartability: always
          Comment: Default policy for XDP relationship with daily and
weekly rules.
    Total Number of Rules: 3
          Total Keep: 139
          Rules: Snapmirror-label
                  Keep Preserve Warn
                  -----
                  daily           7 false 0
                  weekly          52 false 0
                  sm_created       40 true 0
```

## Example of creating a tiered backup policy

Data ONTAP uses the `snapmirror-label` attribute to identify Snapshot copies between primary and secondary FlexVol volumes in a SnapVault relationship. When you configure rules in a SnapVault policy, you enter the `snapmirror-label` name that you want to use to identify the Snapshot copies to which the rule applies.

In a tiered backup strategy, a SnapVault policy might have several rules, each one identifying a different set of Snapshot copies. In this example, you have a volume to which you have assigned a Snapshot policy that specifies the following schedule:

- An hourly Snapshot copy  
Every two hours, a Snapshot copy is created and is assigned the attribute `-snapmirror-label hourly`.
- A daily Snapshot copy

Every day at 5:00 p.m., a Snapshot copy is created and is assigned the attribute `-snapmirror-label daily`.

- A weekly Snapshot copy

Every Friday at 6:00 p.m., a Snapshot copy is created and is assigned the attribute `weekly`.

In addition, the volume is part of an Oracle database. Using the online management tool for Host Services Agent for Oracle, you set up a schedule that creates a Snapshot copy every day at 5:00 p.m. These Snapshot copies are assigned the attribute `-snapmirror-label Oracle-consistent`.

To set up tiered, disk-to-disk data protection for this volume, in which only the Snapshot copies labeled `daily`, `weekly`, and `Oracle-consistent` are replicated to the SnapVault backup, you do the following:

1. Create a separate rule for each of the three types of Snapshot copies that you want replicated to the SnapVault secondary volume.  
You should have three rules. Each rule must specify the retention count. For this example, you configure a retention count of 20 for the daily Snapshot copies, 24 for the weekly Snapshot copies, and 100 for the Oracle-consistent Snapshot copies.
2. Create a new “TieredOracle” SnapVault policy by using the `snapmirror policy create` command, and add the rules you created in Step 1.
3. Assign the new SnapVault policy to the SnapVault relationship that exists between the primary and secondary volumes.

The new SnapVault policy configuration is as follows:

Vserver Name	Policy Name	Number Of Rules	Tries	Transfer Priority	Restart	Comment
vs1	TieredOracle	3	8	normal	default	Example of a tiered backup policy
	SnapMirror-label:	daily		Keep:	20	
		weekly			24	
		Oracle-consistent			100	

### Related concepts

[Commands for managing mirror and SnapVault policies](#) on page 51

### Related references

[Commands for managing mirror and SnapVault policies](#) on page 51

# Providing disaster recovery using mirroring technology

---

Stored data is susceptible to disaster, either through hardware failure or environmental catastrophe. You can use mirroring technology to create an identical second set of data to replace the primary set of data, should something happen to the primary set of data.

## Managing data protection mirror copies for FlexVol volumes

Data protection mirror management for FlexVol volumes consists of activities such as creating data protection mirror copies for source volumes, modifying data protection mirror copies, and monitoring the status of data protection mirror copies.

For information about creating and managing data protection mirrors for Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

## Creating a data protection mirror copy for FlexVol volumes

You can protect data by replicating it to data protection mirror copies. You can use data protection mirror copies to recover data when a disaster occurs.

### Before you begin

- You must have installed a SnapMirror license on both the source and destination cluster.
- You must have created the cluster and Storage Virtual Machine (SVM) peering relationship. To learn about creating cluster and SVM peering, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

### About this task

You can create data protection mirror copies in a cluster using FlexVol volumes only.

### Steps

1. Create a destination volume on the destination SVM that will become the data protection mirror copy by using the `volume create` command.

### Example

The following command creates a data protection mirror volume named `dept_eng_dr_mirror1` on SVM `vs1.example.com`. The destination volume is located on an aggregate named `aggr3`. The destination volume is also on SVM `vs1.example.com`.



```
vs1::> vol create -volume dept_eng_dr_mirror1 -aggregate aggr3 -size
20MB -type DP
```

If you are creating a data protection mirror copy on an SVM peer, the destination volume is created on the SVM peer:

```
vs2::> volume create -volume dept_eng_dr_mirror1
-aggregate aggr3 -size 20MB -type DP
```

2. Create a data protection mirror relationship by using the `snapmirror create` command.

### Example

The following command creates a data protection relation with the destination volume named `dept_eng_dp_mirror2` of the source volume named `dept_eng`. The SVM is named `vs1`.

```
vs1::> snapmirror create -destination-path
vs1.example.com:dept_eng_dp_mirror2
-source-path vs1.example.com:dept_eng -type DP -schedule 5min
```

If you are creating the data protection mirror relationship with the destination volume on an SVM peer, you create the data protection mirror relationship from the SVM that contains the destination volume. For example, if the destination volume were on the SVM peer named `vs2`, the command to create the data protection mirror relationship is as follows:

```
vs2::> snapmirror create -destination-path
vs2.example.com:dept_eng_dp_mirror2
-source-path vs1.example.com:dept_eng -type DP -schedule 5min
```

Data ONTAP creates the data protection mirror relationship, but the relationship is left in an uninitialized state.

3. Initialize the data protection mirror copy by using the `snapmirror initialize` command.

### Example

The following command initializes a data protection mirror copy named `dept_eng_dp_mirror2` of a source volume named `dept_eng`. The source volume and the data protection mirror copy are on the same SVM named `vs1.example.com`.

```
vs1::> snapmirror initialize -destination-path
vs1.example.com:dept_eng_dp_mirror2
```

If you are initializing the data protection mirror relationship with the destination volume on an SVM peer, you must initialize the data protection mirror relationship from the SVM that contains the destination volume. For example, if the destination volume of the single SVM example were

on an SVM peer named `vs2.example.com`, the command to create the data protection mirror relationship is as follows:

```
vs2::> snapmirror initialize -destination-path  
vs2.example.com:dept_eng_dp_mirror2
```

## Correcting a SnapMirror initialization failure

A SnapMirror initialization can fail with the error message `Volume volume_name is restricted` if a previous initialization attempt failed. The initialization fails because the destination volume was restricted in the first failed attempt.

### About this task

You correct a SnapMirror initialization failure by changing the state of the destination volume from restricted to online and then running another initialization attempt.

### Steps

1. Change the state of the destination volume by using the `volume modify` command with the `-state` parameter.

### Example

```
vs2::> volume modify -vserver vs2.example.com -volume vol3_dst  
-state online
```

2. Initialize the SnapMirror relationship by using the `snapmirror initialize` command.

### Example

```
vs2::> snapmirror initialize -destination-path  
vs2.example.com:vol3_dst
```

## Managing mirror relationships

You manage mirror relationships to optimize the performance of those relationships.

### Commands for managing SnapMirror relationships

Data ONTAP includes many commands for managing SnapMirror relationships of FlexVol volumes and Infinite Volumes.

If you want to...	Use this command...
<p>Abort an active transfer. You can use the <code>snapmirror show</code> command to determine the status of the abort operation.</p> <p><b>Note:</b> If you are using Infinite Volumes, management tasks must be performed on the Infinite Volume and not its individual constituents.</p>	<pre>snapmirror abort</pre>
<p>Make a data protection mirror copy destination writeable. This command must be used from the destination Storage Virtual Machine (SVM). You must not have I/O traffic running on the volume when you use this command.</p> <p>When you use this command, the common Snapshot copy between your source and destination is not protected on your source and can be deleted. If you use this command, you should create your own Snapshot copy on the source that will not get deleted automatically, and replicate it to the destination volume before breaking the relationship.</p>	<pre>snapmirror break</pre>
<p>Create a new data protection mirror relationship. This command must be used from the destination SVM.</p> <p><b>Note:</b> If you are using Infinite Volumes, you can create data protection mirror relationships between clusters only, not within a cluster.</p>	<pre>snapmirror create</pre>

If you want to...	Use this command...
<p>Delete a data protection mirror relationship. This command must be used from the destination SVM.</p> <p><b>Note:</b> If you are using Infinite Volumes, management tasks must be performed on the Infinite Volume and not its individual constituents.</p>	<pre>snapmirror delete</pre>
<p>Start a baseline transfer. This command must be used from the destination SVM.</p>	<pre>snapmirror initialize</pre>
<p>Display a list of data protection mirror relationships whose source endpoints are in the current SVM. This command must be used from the source SVM.</p>	<pre>snapmirror list-destinations</pre>
<p>Modify a data protection mirror relationship. This command must be used from the destination SVM.</p>	<pre>snapmirror modify</pre>
<p>Display a list of data protection and load-sharing mirror relationships or display the state of a scheduled transfer for a SnapMirror relationship. The information that this command shows is updated periodically; therefore, any changes to a relationship might not be displayed immediately. This command must be used from the destination SVM.</p>	<pre>snapmirror show</pre>
<p>Disable future transfers for a mirror relationship. This command must be used from the destination SVM.</p>	<pre>snapmirror quiesce</pre>
<p>Enable future transfers for a mirror relationship. This command must be used from the destination SVM.</p>	<pre>snapmirror resume</pre>

If you want to...	Use this command...
<p>Start a resynchronize operation. This command must be used from the destination SVM.</p> <p>You must not have I/O traffic running on the volume when you use this command.</p> <p>Quotas are turned off on the volume you resynchronize. After resynchronizing, you must activate quotas on the volume, if you had quotas on the volume before the resynchronization.</p> <p><b>Attention:</b> A resynchronize operation can cause data loss on the destination volume because the command can remove the exported Snapshot copy on the destination volume.</p>	<pre>snapmirror resync</pre>
<p>Add an owner to prevent premature deletion of a user-created Snapshot copy for a SnapMirror mirror-to-SnapVault cascade configuration.</p> <p>A typical use case is to preserve an application-consistent Snapshot copy.</p> <p><b>Note:</b> This task is not supported for Infinite Volumes.</p>	<pre>snapmirror snapshot-owner create</pre>
<p>Delete an owner used to preserve a user-created Snapshot copy for a SnapMirror mirror-to-SnapVault cascade configuration.</p> <p><b>Note:</b> This task is not supported for Infinite Volumes.</p>	<pre>snapmirror snapshot-owner delete</pre>
<p>Show all the Snapshot copies with owners that were added using the <code>snapmirror snapshot-owner create</code> command.</p> <p><b>Note:</b> This task is not supported for Infinite Volumes.</p>	<pre>snapmirror snapshot-owner show</pre>

If you want to...	Use this command...
<p>Start an incremental transfer. This command must be used from the destination cluster.</p> <p><b>Note:</b> If you are using Infinite Volumes, aggregate requirements must be met before performing the incremental transfer. Management tasks must be performed on the Infinite Volume and not its individual constituents.</p> <p>You can disregard error messages that result from updating a SnapMirror relationship from a Snapshot copy that exists on the destination volume. Any such messages are for support use.</p>	<pre>snapmirror update</pre>
<p>Create a new policy for a data protection mirror relationship.</p>	<pre>snapmirror policy create</pre>
<p>Delete a policy of a data protection mirror relationship.</p>	<pre>snapmirror policy delete</pre>
<p>Add a new rule to a SnapVault relationship.</p>	<pre>snapmirror policy add-rule</pre>
<p>Modify an existing rule in the policy of a SnapVault relationship.</p>	<pre>snapmirror policy modify-rule</pre>
<p>Modify a policy of a data protection mirror relationship.</p>	<pre>snapmirror policy modify</pre>
<p>Remove a rule from the policy of a data protection mirror relationship.</p>	<pre>snapmirror policy remove-rule</pre>
<p>Show the policy of a data protection mirror relationship.</p>	<pre>snapmirror policy show</pre>
<p>Copy data to a volume.</p> <p>Quotas are turned off on the volume you restore. After the restore, you must activate quotas on the volume, if you had quotas on the volume before the restore.</p> <p><b>Note:</b> This task is not supported for Infinite Volumes.</p>	<pre>snapmirror restore</pre>
<p>Remove the SnapMirror relationship information from the source SVM. This command must be used from the source SVM.</p>	<pre>snapmirror release</pre>

See the man page for each command for more information.

## Using extended queries to operate on many SnapMirror relationships

You can use extended queries to perform SnapMirror operations on many SnapMirror relationships at one time. For example, you might have many uninitialized SnapMirror relationships that you want to initialize using one command.

### About this task

You can apply extended queries to the following SnapMirror operations:

- Initializing many `Uninitialized` SnapMirror relationships
- Resuming many `Quiesced` SnapMirror relationships
- Resynchronizing many `Broken` SnapMirror relationships
- Updating many `Idle` SnapMirror relationships
- Aborting many currently `Transferring` SnapMirror relationships

### Step

1. You perform a SnapMirror operation on many SnapMirror relationships by using the following syntax: `snapmirror command {-state state } *`

### Example

The following command initializes only SnapMirror relationships that are in an `Uninitialized` state:

```
vs1::> snapmirror initialize {-state Uninitialized} *
```

## What tape seeding is

Tape seeding is an SMTape functionality that helps you initialize a destination FlexVol volume in a data protection mirror relationship.

Tape seeding enables you to establish a data protection mirror relationship between a source system and a destination system over a low-bandwidth connection. Incremental mirroring of Snapshot copies from the source to the destination is feasible over a low bandwidth connection. However, an initial mirroring of the base Snapshot copy would take a long time over a low-bandwidth connection. In such a case, you can perform an SMTape backup of the source volume to a tape and use the tape to transfer the initial base Snapshot copy to the destination. You can then set up incremental SnapMirror updates to the destination system using the low-bandwidth connection.

## Performing tape seeding using SMTape

Using SMTape, cluster administrators can perform tape seeding to initialize a destination FlexVol volume in a data protection mirror relationship. The time taken to initialize this destination volume

over a low bandwidth connection using SMTape is faster when compared to using the `snapmirror initialize` command.

### Before you begin

Volume and tape must be located on the same node.

### About this task

To perform tape seeding successfully, you must ensure that all nodes in a cluster are running clustered Data ONTAP 8.2 and later. You can view the event log files at any point in time. For more information about accessing and viewing event log files, see *Clustered Data ONTAP Data Protection Tape Backup and Recovery Guide*. For more information about SMTape commands or other commands, see the man pages.

**Note:** Both source and destination volumes must be located on storage systems running clustered Data ONTAP.

### Steps

1. Determine which Snapshot copy you want to use for tape seeding by using the `volume snapshot show` command.

### Example

The following example lists the Snapshot copies:

```
clus1::> vol snapshot show -vserver vs1 -volume voll
(volume snapshot show)
```

Vserver	Volume	Snapshot	State	Size	Total%	Used%
vs1	voll	hourly.2013-01-25_0005	valid	224KB	0%	0%
		daily.2013-01-25_0010	valid	92KB	0%	0%
		hourly.2013-01-25_0105	valid	228KB	0%	0%
		hourly.2013-01-25_0205	valid	236KB	0%	0%
		hourly.2013-01-25_0305	valid	244KB	0%	0%
		hourly.2013-01-25_0405	valid	244KB	0%	0%
		hourly.2013-01-25_0505	valid	244KB	0%	0%

7 entries were displayed.

2. If you do not have an existing Snapshot copy, manually make a Snapshot copy of a source volume by using the `volume snapshot create` command.

You must specify Storage Virtual Machine (SVM, formerly known as Vserver) name, source volume name, and Snapshot copy name.

**Attention:** You must not delete this Snapshot copy until tape seeding is over.



**Example**

The following example shows how to create a Snapshot copy named mysnap of a source volume named src1 on an SVM named vs1. You can view the details of the Snapshot copy mysnap by using the `volume snapshot show` command:

```
clus1::> volume snapshot create -vserver vs1 -volume src1 -snapshot mysnap
clus1::> volume snapshot show -vserver vs1 -volume src1 -snapshot mysnap

                Vserver: vs1
                Volume: src1
                Snapshot: mysnap
                Creation Time: Thu Aug 09 12:03:46 2012
                Snapshot Busy: false
                List of Owners: -
                Snapshot Size: 52KB
                Percentage of Total Blocks: 0%
                Percentage of Used Blocks: 1%
                Comment:
                7-Mode Snapshot: false
                Label for SnapMirror Operations: -
                Snapshot State: valid
                Constituent Snapshot: false
```

3. Move and position the tape correctly by using the `run -node <node_name> mt` command at the nodeshell.

**Note:** For more information about the `mt` command, see the man pages.

**Example**

The following example moves and positions a no rewind tape device st01:

```
clus1::> run -node clus1-01 "mt -t nrst01 rewind"
clus1::> run -node clus1-01 "mt -t nrst01 status"

Tape drive: Hewlett-Packard LTO-4
Status: ready, write enabled
Format: LTO-2(ro)/3 2/400GB
fileno = 0 blockno = 0 resid = 0
```

4. Use the `smtape backup` command to copy all volume Snapshot copies including the base Snapshot copy to tape.

**Note:** If you back up 32-bit volumes, then you can restore these volumes only in the Data ONTAP 8.2 release family.

**Example**

The following example backs up Snapshot copy mysnap to the tape device st01:

```
clus1:> system smtape backup -vserver vs1 -volume src1 -backup-snapshot
mysnap -tape /clus1-01/nrst01
```

```
Session 35 created successfully
```

5. Use the `smtape status show` command to see the progress of the baseline transfer.

### Example

The following example shows the progress and status of the SMTape backup operation triggered in the previous step:

```
clus1:> system smtape status show -session 35 -instance
```

```
Session Identifier: 35
  Node Name: clus1-01
  Operation Type: backup
  Session Status: ACTIVE
  Path Name: /vs1/src1
  Device Name: /clus1-01/nrst01
  Bytes Transferred: 0B
  Start Time: 8/9/2012 12:03:55
  End Time: -
  Snapshot Name: mysnap
  Tape Block Size: 240
  Error Description: None
```

```
clus1:> smtape status show
(system smtape status show)
```

Session Type	Status	Progress	Path	Device	Node
35 backup	COMPLETED	6.01MB	/vs1/src1	/clus1-01/nrst01	clus1-01

6. Depending upon the status of the SMTape backup operation, you can perform one of the following actions:

---

#### If the Status shows... Then...

---

COMPLETED	Baseline transfer is complete; go to step 7.
WAITING	<ol style="list-style-type: none"> <li>a. Load and position the new tape by using the <code>run -node &lt;node_name&gt; mt</code> command.</li> <li>b. Continue the SMTape backup operation by using the <code>smtape continue</code> command.</li> </ol>

---

7. Physically transport the tapes to the destination node.
8. Optional: View the data backed up on a tape by using the `smtape showheader` command.

9. Create a destination volume of type DP and appropriate size (same or greater than the source volume size) on the destination cluster that will become the data protection mirror by using the `volume create` command and restrict the volume.

### Example

The following example creates a data protection mirror volume named `dst1` on Storage Virtual Machine (SVM, formerly known as Vserver) named `vs1`. The destination volume is located on an aggregate named `aggr5`. Destination volume `dst1` is in the restricted state:

```
clus1::> volume create -vserver vs1 -volume dst1 -aggregate aggr5 -size 400m -
type DP -state restricted

[Job 83] Job is queued: Create dst1.
[Job 83] Initializing
[Job 83] Job succeeded: Successful
```

10. Move and position the tape correctly by using the `run -node <node_name> mt` command at the nodeshell.

**Note:** For more information about the `mt` command, see the man pages.

### Example

The following example moves and positions a no rewind tape device `st01` at the destination volume:

```
clus1::> run -node clus1-01 "mt -t nrst01 rewind"

clus1::> run -node clus1-01 "mt -t nrst01 status"

Tape drive: Hewlett-Packard LTO-4
Status: ready, write enabled
Format: LTO-2(ro)/3 2/400GB
fileno = 0 blockno = 0 resid = 0
```

11. Use the `smtape restore` command to restore all the volume Snapshot copies including the base Snapshot copy from tape to the destination volume.

### Example

The following example restores all the data from tape to the destination volume `dst1` on Storage Virtual Machine (SVM, formerly known as Vserver) `vs1`:

```
clus1::> system smtape restore -vserver vs1 -volume dst1 -tape /clus1-01/nrst01
Session 36 created successfully
```

12. Use the `smtape status show` command to see the progress of the baseline transfer.

**Example**

The following example shows the progress and status of the SMTape restore operation triggered in the previous step:

```
clus1::> system smtape status show -session 36 -instance

Session Identifier: 36
  Node Name: clus1-01
  Operation Type: restore
  Session Status: ACTIVE
  Path Name: /vs1/dst1
  Device Name: /clus1-01/nrst01
Bytes Transferred: 0B
  Start Time: 8/9/2012 12:04:15
  End Time: -
  Snapshot Name: None
  Tape Block Size: 240
  Error Description: None

clus1::> system smtape status show

Session Type      Status      Progress  Path          Device          Node
-----
   36 restore ACTIVE          0B /vs1/dst1 /clus1-01/nrst01 clus1-01
   35 backup COMPLETED 6.01MB /vs1/src1 /clus1-01/nrst01 clus1-01
2 entries were displayed.
```

13. Depending upon the status of the SMTape restore operation, you can perform one of the following actions:

---

**If the Status shows... Then...**


---

COMPLETED	Baseline transfer is complete; go to step 14.
WAITING	<ol style="list-style-type: none"> <li>a. Load and position the new tape by using the run <code>-node &lt;node_name&gt; mt</code> command.</li> <li>b. Continue the SMTape restore operation by using the <code>smtape continue</code> command.</li> </ol>

---

14. Use the `smtape break` command to break the volume and tape relationship.

**Note:** This command is available for the SVM administrators also.

**Example**

The following example breaks the SMTape relationship between the tape and the volume `dst1`:

```
clus1::> system smtape break -vserver vs1 -volume dst1

[Job 84] Job is queued: snapmirror break for destination vs1:dst1.
[Job 84] Job succeeded: SnapMirror Break Succeeded
```

This makes the destination volume writeable and a mirror relationship can be reestablished.

15. Establish the SnapMirror or SnapVault relationship by using the `snapmirror resync` command.

A SnapMirror relationship creates a data protection copy of the source volume; a SnapVault relationship creates a backup copy. The `-type` parameter determines the type of relationship established. The value of the `-type` parameter for a SnapMirror relationship is `DP` and the value for a SnapVault relationship is `XDP`.

### Example

The following example reestablishes a SnapMirror relationship between the destination volume `dst1` and the source volume `src1`:

```
clus1:> snapmirror resync -destination-path vs1:dst1 -source-path vs1:src1 -
type DP

[Job 85] Job is queued: initiate snapmirror resync to destination "vs1:dst1".
[Job 85]
[Job 85] Job succeeded: SnapMirror Resync Transfer Queued
```

16. Use the `snapmirror show` command to see the progress of the data protection mirror relationship reestablished between the destination volume and source volume in the previous step.

### Example

The following example shows the data protection mirror relationship established between the source volume `vs1` and destination volume `dst1`. The data protection mirror relationship type established is `DP`:

```
clus1:> snapmirror show -destination-path vs1:dst1

          Source Path: vs1:src1
          Destination Path: vs1:dst1
          Relationship Type: DP
          SnapMirror Schedule: -
          Tries Limit: -
          Throttle (KB/sec): unlimited
          Mirror State: -
          Relationship Status: Transferring
          Transfer Snapshot: snapmirror.58621f01-
e214-11e1-833d-123478563412_2147484708.2012-08-09_120444
          Snapshot Progress: 0B
          Total Progress: 0B
          Snapshot Checkpoint: -
          Newest Snapshot: -
          Newest Snapshot Timestamp: -
          Exported Snapshot: -
          Exported Snapshot Timestamp: -
          Healthy: true
          Constituent Relationship: false
          Relationship ID: 6485d262-e21a-11e1-833d-123478563412
          Transfer Type: resync
          Transfer Error: -
          Current Throttle: 103079214
```

```

Current Transfer Priority: normal
  Last Transfer Type: -
  Last Transfer Error: -
  Last Transfer Size: -
  Last Transfer Duration: -
  Last Transfer From: -
  Progress Last Updated: 08/09 12:04:45
Relationship Capability: 8.2 and above
  Lag Time: -
  Policy: DPDefault

```

When the relationship status shows `idle`, the data protection mirror relationship is established and tape seeding is complete.

### Example

```

clus1::> snapmirror show -destination-path vs1:dst1

      Source Path: vs1:src1
      Destination Path: vs1:dst1
      Relationship Type: DP
      SnapMirror Schedule: -
        Tries Limit: -
      Throttle (KB/sec): unlimited
      Mirror State: Snapmirrored
      Relationship Status: Idle
      Transfer Snapshot: -
      Snapshot Progress: -
      Total Progress: -
      Snapshot Checkpoint: -
        Newest Snapshot: snapmirror.58621f01-
e214-11e1-833d-123478563412_2147484708.2012-08-09_120444
      Newest Snapshot Timestamp: 08/09 12:04:44
        Exported Snapshot: snapmirror.58621f01-
e214-11e1-833d-123478563412_2147484708.2012-08-09_120444
      Exported Snapshot Timestamp: 08/09 12:04:44
      Healthy: true
      Constituent Relationship: false
        Relationship ID: 6485d262-e21a-11e1-833d-123478563412
        Transfer Type: resync
        Transfer Error: -
      Current Throttle: 103079214
      Current Transfer Priority: normal
        Last Transfer Type: resync
        Last Transfer Error: -
        Last Transfer Size: 72KB
        Last Transfer Duration: 0:0:7
        Last Transfer From: vs1:src1
      Progress Last Updated: -
      Relationship Capability: 8.2 and above
        Lag Time: 0:0:7
        Policy: DPDefault

```

### Related information

*[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)*

## Scalability limits for SMTape backup and restore sessions

While performing tape seeding, you must be aware of the maximum number of SMTape backup and restore sessions that can be performed simultaneously on storage systems of different system memory capacities. This maximum number depends on the system memory of a storage system.

System memory of a storage system	Maximum number of NDMP sessions
Less than 16 GB	6
Greater than or equal to 16 GB but less than 24 GB	16
Greater than or equal to 24 GB	32

You can obtain the system memory of your storage system by using the `sysconfig -a` command (available through the nodeshell). For more information about using this command, see the man pages.

## Listing the schedule state of a mirror relationship

You might want to see what state a scheduled transfer for a mirror relationship is in to ensure that jobs are running as they should.

### About this task

The state of a scheduled job might be dormant. The dormant state means that the job is waiting for the scheduled start time to begin the transfer. There is nothing wrong with the job and you do not need to do anything.

### Step

1. To see the state of scheduled jobs, use the `snapmirror show` command.

## Scheduling SnapMirror transfers

If you want scheduled SnapMirror transfers, you can add a schedule to a mirror relationship after you initially create the relationship.

### About this task

Unless you create and implement a schedule for SnapMirror transfers, you are limited to manually updating destination FlexVol volumes or Infinite Volumes with mirror relationships. The following are characteristics of adding a SnapMirror transfer schedule:

- When you add a schedule for a data protection mirror copy of Infinite Volumes, do not schedule updates for less than one-hour intervals.

If you schedule updates for less than one-hour intervals, Data ONTAP tries but cannot meet the schedule for Infinite Volumes, and the data protection mirror relationship is displayed as unhealthy.

- Scheduled SnapMirror transfers (or even manual updates) can disrupt Snapshot copy schedules when a transfer lasts longer than the full retention period of the schedule.

### Steps

1. Create the schedule you want to implement by using the `job schedule cron create` command.

**Note:** You cannot use the `job schedule interval create` command to schedule SnapMirror transfers.

2. Apply the schedule to the mirror relationship by using the `-schedule` option of the `snapmirror modify` command.

See the `snapmirror modify` command man page for more information about the command.

## Changing mirror relationship schedules

You can change a schedule that updates mirror relationships for FlexVol volumes and Infinite Volumes if the schedule impacts other backups or updates.

### Before you begin

- You must have created the cluster and Storage Virtual Machine (SVM) peering relationship. To know about creating cluster and SVM peering, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

### About this task

Changing a schedule affects load-sharing mirror copies differently than it does for data-protection mirror copies. If you change a schedule to a load-sharing mirror relationship, Data ONTAP makes the change to the relationships of all the load-sharing mirror copies in the group. Data ONTAP determines the load-sharing mirror group by the SVM and source volume specified by the command. See the *Clustered Data ONTAP Logical Storage Management Guide* for more information about load-sharing mirror copies.

### Steps

1. Create the new schedule by using the `job schedule cron create` command.

Creating a schedule is described in the cron job creation section of the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*. See the man page for details about the `job schedule cron create` command.

2. Change the schedule for a mirror relationship by using the `snapmirror modify -schedule` command.



This command must be used from the destination SVM.

### Example

The following command changes the update schedule used by a data protection mirror relationship for destination volume named “dept\_eng\_ls1” to a schedule named “dept\_eng\_mirror\_sched”:

```
vs2::> snapmirror modify -source-path vs1:dept_eng
-destination-path vs2:dept_eng_ls1 -schedule dept_eng_mirror_sched
```

## Manually updating data protection mirror copies on destination volumes

You can schedule updates to data protection mirror copies for destination FlexVol volumes or Infinite Volumes, or you can manually update a data protection mirror copy to transfer Snapshot copies between the source and destination volume. However, for Infinite Volumes, you cannot choose which Snapshot copies to transfer.

### Before you begin

- A base Snapshot copy must exist on the source volume and the destination volume.
- The destination volume must be the same size or bigger than the source volume.

### About this task

You can have cluster administrator or Storage Virtual Machine (SVM) administrator privileges to perform this task.

When you update a destination volume, all of the Snapshot copies from the source volume are transferred to the destination volume. In addition, any Snapshot copies deleted from the source volume are deleted from the destination volume during the update. Similarly, any new Snapshot copies on the source volume are transferred to the destination volume.

### Step

1. On the destination cluster, manually update a destination volume by using the `snapmirror update` command.

### Example

The following command updates the data protection mirror relationship for a destination volume named `repo_vol_dest` on an SVM named `vs0_dest`:

```
vs2::> snapmirror update -destination-path vs0_dest:repo_vol_dest
```

## Deleting a mirror copy

You can delete a mirror relationship and the destination FlexVol volume or Infinite Volume if you no longer want the mirror copy.

### About this task

When you delete a mirror copy, you must delete the mirror relationship and the destination volume. Deleting the mirror relationship does not delete SnapMirror created Snapshot copies on either the source or destination volumes. Deleting the mirror relationship attempts to delete Snapshot copy owners for the SnapMirror created Snapshot copies on both source and destination volumes.

When you delete a load-sharing mirror copy from a set of load-sharing mirror copies, the destination volume of the deleted load-sharing mirror relationship cannot be used again as a destination volume of a load-sharing relationship if it contains any data or Snapshot copies.

### Steps

1. Optional: On the source Storage Virtual Machine (SVM), use the `snapmirror list-destination` command to view the list of destination volumes for that source volume.

### Example

```
vs1::> snapmirror list-destinations
```

Source Path	Type	Destination Path	Status	Transfer Progress	Progress Last Updated	Relationship Id
vs1:src_ui	DP	vs2:vsrsrc_ui_ls_mir2	Idle	-	-	3672728c-ad06-11e2-981e-123478563412

2. Use the `snapmirror delete` command to delete a mirror relationship.

### Example

The following command deletes a mirror relationship between a destination volume named `src_ui_ls_mir2` and a source volume named `src_ui`:

```
vs2::> snapmirror delete -source-path vs1:src_ui
-destination-path vs2:src_ui_ls_mir2
```

The command deletes the mirror relationship, but does not delete the destination volume. In the case of a load-sharing mirror copy, the destination volume will be in the restricted state. If you want to use it as a destination volume of a data protection relationship, you must wait at least 10 minutes. This is the amount of time required to refresh internal caches and place the volume back online.

3. Use the `snapmirror release` command from the source SVM to remove the configuration information and Data ONTAP created Snapshot copies on the source volume.

### Example

The following command removes the DP or XDP relationship from the source SVM named `vs1`:

```
vs1::> snapmirror release -source-path vs1:src_ui -destination-path
vs2:src_ui_ls_mir2
```

This command removes the DP or XDP relationship information from the source SVM and does not delete any volumes. This command deletes the base Snapshot copies for the destination volume named `src_ui_ls_mir2` from the source volume named `src_ui`.

4. Optional: Use the `volume delete` command to delete the destination volume.

Delete the destination volume if you no longer need the volume.

## Reversing the data protection mirror relationship when disaster occurs

When disaster disables the source FlexVol volume of a data protection mirror relationship, you can use the destination FlexVol volume to serve data while you repair or replace the source, update the source, and reestablish the original configuration of the systems.

### About this task

The following procedure describes a data protection mirror relationship that has the source volume on one Storage Virtual Machine (SVM) and the destination volume on another SVM. The source and the destination clusters and source and destination SVMs are in peer relationships. The original source (the one disabled by the disaster) is `vs1:volA` and the original destination is `vs2:volB`.

All data from the last scheduled SnapMirror Snapshot copy before the source was disabled and all the data written to `vs2:volB` after it was made writeable is preserved. Any data written to `vs1:volA` between the last SnapMirror Snapshot copy and the time that `vs1:volA` was disabled is not preserved.

For information about retrieving data from Infinite Volumes during disaster recovery, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

### Steps

1. Temporarily make the original source volume a read-only destination volume and reverse the data protection mirror relationship to continue to serve data.

If the source `vs1:volA` is recoverable and its data is intact, complete the following steps:

- a) After the source volume (in this case, `vs1:volA`) is disabled, use the `snapmirror break` command on the destination volume, `vs2:volB`, to make the destination volume, `vs2:volB`, writeable.

### Example

```
vs2::> snapmirror break vs2:volB
```

- b) Redirect the clients of the source volume `vs1:volA` to the new source volume `vs2:volB`.  
The former clients of `vs1:volA` access and write to `vs2:volB`.
- c) On the destination volume, `vs2:volB`, use the `snapmirror delete` command to remove the data protection mirror relationship between the source `vs1:volA` and the destination `vs2:volB`.

### Example

```
vs2::> snapmirror delete vs2:volB
```

- d) On the source volume, `vs1:volA`, use the `snapmirror release` command to remove relationship information from the source.

### Example

```
vs1::> snapmirror release vs2:volB
```

- e) On the new destination volume, `vs1:volA`, use the `snapmirror create` command to create the mirror relationship, but with `vs2:volB` as the new source and `vs1:volA` as the new destination.

### Example

```
vs1::> snapmirror create vs2:volB vs1:volA -type DP
```

- f) If there are LUNs on the original source volume, `vs1:volA`, remove the mapping by using the `lun unmap` command.
- g) On the new destination volume, `vs1:volA`, use the `snapmirror resync` command to resynchronize `vs1:volA` with `vs2:volB`.

### Example

```
vs1::> snapmirror resync vs1:volA
```

- h) If there were LUNs, map the LUNs on the new source `vs2:volB` by using the `lun map` command.

If the source `vs1:volA` is unrecoverable, complete the following steps:

- a) After the source volume (in this case, `vs1:volA`) is disabled, use the `snapmirror break` command on the destination volume, `vs2:volB`, to make the destination volume, `vs2:volB`, writeable.

### Example

```
vs2::> snapmirror break vs2:volB
```

- b) Redirect the clients of the source volume `vs1:volA` to the new source volume `vs2:volB`.  
The former clients of `vs1:volA` access and write to `vs2:volB`.
- c) On the destination volume, `vs2:volB`, use the `snapmirror delete` command to remove the data protection mirror relationship between the source `vs1:volA` and the destination `vs2:volB`.

### Example

```
vs2::> snapmirror delete vs2:volB
```

- d) On the source SVM, `vs1`, use the `snapmirror release` command to remove relationship information from the source.

Even though the source volume is unrecoverable, the data protection mirror relationship still exists and must be removed.

### Example

```
vs1::> snapmirror release vs2:volB
```

- e) Delete the old volume `volA` and use the `volume create` command to create a new data protection destination volume called `vs1:volA`.

**Note:** Remember to use the `-type DP` parameter when creating the destination volume.

### Example

```
vs1::> volume delete -volume vs1:volA
```

```
vs1::> volume create -volume volA -aggr aggr1  
-type DP -vserver vs1
```

- f) On the new destination volume, `vs1:volA`, use the `snapmirror create` command to create the data protection mirror relationship with `vs2:volB` as the new source volume and `vs1:volA` as the new destination volume.

**Example**

```
vs1::> snapmirror create vs2:volB vs1:volA -type DP
```

- g) On the new destination volume, `vs1:volA`, use the `snapmirror initialize` command to create the baseline on the data protection mirror copy.

This command also makes `vs1:volA` a read-only destination.

**Example**

```
vs1::> snapmirror initialize vs1:volA
```

- h) If there were LUNs, map the LUNs on the new source `vs2:volB` by using the `lun map` command.

You can keep this configuration or, after resolving the problem with the original source volume, you can complete the following steps to reestablish the original data protection mirror relationship.

2. On the new destination volume, `vs1:volA`, update the new destination volume `vs1:volA` to transfer the latest data from the new source volume `vs2:volB` by using the `snapmirror update` command.

**Example**

```
vs1::> snapmirror update vs1:volA
```

3. On the new destination volume, `vs1:volA`, use the `snapmirror break` command to make `vs1:volA` writeable.

**Example**

```
vs1::>> snapmirror break -source-path vs2:volB -destination-path vs1:volA
```

4. On the new destination volume, `vs1:volA`, use the `snapmirror delete` command to remove the data protection mirror relationship between the new source `vs2:volB` and the new destination `vs1:volA`.

**Example**

```
vs1::> snapmirror delete vs1:volA
```

5. On the new source volume, `vs2:volB`, use the `snapmirror release` command to remove the data protection mirror relationship between the new source `vs2:volB` and the new destination `vs1:volA`.

**Example**

```
vs2::> snapmirror release vs1:volA
```

6. On the original destination volume, `vs2:volB`, use the `snapmirror create` command to re-create the original data protection mirror relationship with `vs1:volA` as the source and `vs2:volB` as the destination.

**Example**

```
vs2::> snapmirror create vs1:volA vs2:volB -type DP
```

7. If there are LUNs on the source `vs2:volB`, remove the mapping by using the `lun unmap` command.
8. On the original destination volume, `vs2:volB`, use the `snapmirror resync` command to resynchronize the original source and original destination volumes.

**Example**

```
vs2::> snapmirror resync vs2:volB
```

9. Redirect the clients from `vs2:volB` back to their original source volume `vs1:volA`.
10. If there were LUNs, map them back to the original source `vs1:volA`.

## Converting a data protection mirror destination to a writable volume

You can convert the destination volume of a data protection mirror relationship to a writable volume if you want to use that volume to serve data. For example, you might do this if you want to migrate a volume.

**Steps**

1. On the destination Storage Virtual Machine (SVM), make the destination volume writable by using the `snapmirror break` command.
2. Remove the data protection mirror relationship that the destination volume has with the source volume by using the `snapmirror delete` command.
3. On the source SVM, remove the configuration information and Data ONTAP created Snapshot copies by using the `snapmirror release` command.

## Testing database applications

You can use data protection mirror relationships to create a copy of database data on which to test software applications that run on a database, to avoid the possibility of changing or corrupting the database.

### Before you begin

The volume that contains the database must be in a data protection mirror relationship.

### About this task

Although you can use this method to test database applications, the preferred method is to create a FlexCache volume. See the *Clustered Data ONTAP Logical Storage Management Guide* for information about creating FlexCache volumes.

### Steps

1. On the destination Storage Virtual Machine (SVM), use the `snapmirror break` command to make the destination volume writeable.

### Example

```
vs2::>> snapmirror break -destination-path vs2:Test_vol
```

2. Run the application on the data in the former destination volume (Test\_vol).
3. Check the data in the former destination volume (Test\_vol).
4. If testing results in alterations to the data that you do not want, use the `snapmirror resync` command to reestablish the mirror relationship.
5. Repeat Steps 2, 3, and 4, until you are satisfied with the testing.

### After you finish

After completing the test, you can resynchronize the source and the destination volumes.



# Protecting data on FlexVol volumes by using SnapVault backups

---

You can create a SnapVault relationship between FlexVol volumes and assign a SnapVault policy to it to create a SnapVault backup. A SnapVault backup contains a set of read-only backup copies, located on a secondary volume.

**Note:** SnapVault relationships are supported on clusters running Data ONTAP 8.2 or later. SnapVault relationships are not supported on Infinite Volumes.

A SnapVault backup differs from a set of Snapshot copies or a set of mirror copies on a destination volume. In a SnapVault backup, the data in the secondary volume is periodically updated to keep the data in the secondary volume up to date with changes made in the primary data.

## Creating SnapVault backups on FlexVol volumes

You configure a SnapVault relationship and assign a SnapVault policy to the relationship to establish a SnapVault backup.

### About this task

The commands you use to create SnapVault backups are the same commands you use to create data protection mirrors. For a list of commands, see [Commands for managing data protection mirror copies](#) on page 59.

### Related concepts

[Supported data protection deployment configurations](#) on page 34

## Guidelines for creating SnapVault relationships on FlexVol volumes

You must follow certain guidelines when creating SnapVault relationships.

### General guidelines for creating a SnapVault relationship

The following guidelines apply to all SnapVault relationships:

- A volume can be in multiple relationships, either as the secondary or the primary. A volume can be the primary for multiple secondaries and also the secondary for another primary.
- A volume can be the secondary for only one SnapVault relationship.
- You cannot configure SnapVault relationships from multiple primary volumes to a single SnapVault secondary volume.

For example, if you want to back up an entire Storage Virtual Machine (SVM) to a SnapVault backup, then you must create a separate secondary volume for each volume in the SVM, and create a separate SnapVault relationship for each primary volume.

- You can configure SnapVault relationships to be used simultaneously with data protection mirror relationships.
- Primary or secondary volumes cannot be 32-bit volumes.
- The primary of a SnapVault backup should not be a FlexClone volume. The relationship will work, but the efficiency provided by FlexClone volumes is not preserved.
- A SnapVault secondary volume cannot be the primary volume of FlexCache volumes.
- Primary and secondary volumes must have the same `vol lang` settings.
- After you establish a SnapVault relationship, you cannot change the language assigned to the secondary volume.
- A SnapVault relationship can be only one leg of a cascade chain.
- After you establish a SnapVault relationship, you can rename primary or secondary volumes. If you rename a primary volume, it can take a few minutes for the relationship to recover from the name change.

### Guidelines for creating a SnapVault relationship to a prepopulated secondary

Typically, you create a prepopulated secondary volume when you copy a primary volume to a secondary volume using tape. This process is known as *tape seeding*.

If the SnapVault secondary volume already contains data, you can create a SnapVault relationship by using the `snapmirror resync` command with the `-type XDP` option.

Before creating a SnapVault relationship to a prepopulated secondary, you must use the following guidelines:

- The primary and secondary volumes must have a common Snapshot copy.
- Snapshot copies on the secondary volume that are newer than the common Snapshot copy are deleted.

When a SnapVault relationship is created, all Snapshot copies on the secondary volume that are more recent than the common Snapshot copy and that are not present on the primary volume are deleted. Newer Snapshot copies on the primary volume that match the configured SnapVault policy are transferred to the secondary volume according to the SnapVault policy.

You can use the `-preserve` option to keep any Snapshot copies that are more recent than the common Snapshot copy on the SnapVault secondary volume and that are not present on the primary volume.

When you use the `-preserve` option, data on the secondary volume is logically made the same as the common Snapshot copy. All newer Snapshot copies on the primary volume that match the SnapVault policy are transferred to the secondary volume.

This option is useful when the latest common Snapshot copy is deleted from the primary volume but another, older common Snapshot copy between the primary and secondary volumes still exists.

## SnapVault updates fail if destination aggregate is full

If the aggregate that contains the secondary volume of the SnapVault backup is out of space, SnapVault updates fail, even if the secondary volume has space.

Ensure that there is free space in the aggregate and the volume for transfers to succeed.

## Prepopulated SnapVault secondary scenarios

There are several ways in which a secondary FlexVol volume for a SnapVault relationship might be prepopulated with data.

The following are some scenarios in which a SnapVault secondary might be populated before a SnapVault relationship is created:

- You used tape backups to provide a baseline transfer to a secondary volume.
  - Note:** Disk seeding to establish a baseline is not supported for SnapVault backups.
- A SnapVault primary volume in a cascade becomes unavailable.
 

You have a data protection mirror relationship between a source and a destination volume (a mirror relationship from A to B) and a SnapVault relationship between the secondary destination volume and a tertiary destination volume (a SnapVault relationship from B to C). The backup cascade chain is A mirror to B and B SnapVault backup to C. If the volume on B becomes unavailable, you can configure a SnapVault relationship directly from A to C. The cascade chain is now A SnapVault backup to C, where C was prepopulated with data.
- You created a SnapVault relationship between two flexible clones.
 

You create a SnapVault relationship between two flexible clones for which their respective parent volumes are already in a SnapVault relationship.
- You extended the SnapVault backup protection beyond 251 Snapshot copies.
 

To extend the SnapVault backup protection beyond the volume limit of 251 Snapshot copies, you can clone the secondary volume. The original SnapVault secondary volume is the parent volume for the new flexible clone.
- You restored data from a SnapVault secondary to a new primary volume.
 

You have a SnapVault relationship from A to B. A becomes inaccessible, so the SnapVault secondary volume (B) is used for a baseline restore operation to a new SnapVault secondary volume (C).

After the restore operation finishes, you establish a new SnapVault relationship from the new secondary volume (C), which now becomes the primary volume, and the original SnapVault secondary volume (in other words, C to B). The disk to disk backup relationship is now C to B, where B was prepopulated with data.
- You deleted the base Snapshot copy from the primary volume.
 

You deleted the base Snapshot copy from the primary volume that was used for a SnapVault transfer, but another, older Snapshot copy exists that is common between the primary and secondary volumes.

## Creating a SnapVault backup in an empty FlexVol volume

You can protect data that has long-term storage requirements on a FlexVol volume by replicating selected Snapshot copies to a SnapVault backup on another Storage Virtual Machine (SVM) or cluster.

### Before you begin

- You must have cluster administrator privileges to perform this task for a cluster, and SVM administrator privileges to perform this task for an SVM.
- If the primary and secondary volumes are in different SVMs, the SVMs must be in a peer relationship.  
If the primary and secondary volumes are in different clusters, the clusters must be in a peer relationship.  
For information about creating peer relationships, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.
- A SnapVault policy must exist.  
You must either create one or accept the default SnapVault policy (named `XDPDefault`) that is automatically assigned.  
Only Snapshot copies with the labels configured in the SnapVault policy rules are replicated in SnapVault operations.
- The Snapshot policy assigned to the primary volume must include the `snapmirror-label` attribute.  
You can create a new Snapshot policy by using the `volume snapshot policy add-schedule` command, or you can modify an existing policy by using the `volume snapshot policy modify-schedule` command to set the `snapmirror-label` attribute for the set of Snapshot copies that you want backed up to the SnapVault secondary volume. Other Snapshot copies on the primary volume are ignored by the SnapVault relationship.
- Your work environment must be able to accommodate the time it might take to transfer a baseline Snapshot copy with a large amount of data.

### Steps

1. On the destination SVM, create a SnapVault secondary volume with a volume type `DP`.  
For information about creating a FlexVol volume, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.
2. Create a schedule that Data ONTAP uses to update the SnapVault relationship by using the `job schedule cron create` command.

For more information, see [Scheduling SnapMirror transfers](#) on page 71.

### Example

The following command creates a schedule that runs on the weekend at 3 a.m.:

```
vserverB::> job schedule cron create -name weekendcron -dayofweek
"Saturday, Sunday" -hour 3 -minute 0
```

3. On the source SVM, create a Snapshot copy policy that contains the schedule of when Snapshot copies with `snapmirror-label` attributes occur by using the `volume snapshot policy create` command with the `snapmirror-label` parameter, or use the default Snapshot copy policy called `default`.

### Example

The following command creates a Snapshot copy policy called “keep-more-snapshot”:

```
vserverB::> snapshot policy create -vserver vs1 -policy keep-more-
snapshot
-enabled true -schedule1 weekly -count1 2 -prefix1 weekly -schedule2
daily
-count2 6 -prefix2 daily -schedule3 hourly -count3 8 -prefix3 hourly
```

The name specified in the `snapmirror-label` attribute for the new Snapshot policy must match the `snapmirror-label` attribute that is specified in the SnapVault policy. This ensures that all subsequent Snapshot copies created on the primary volume have labels that are recognized by the SnapVault policy.

The default Snapshot copy policy has two `snapmirror-label` attributes associated with it, `daily` and `weekly`.

4. Create a SnapVault policy by using the `snapmirror policy create` command, or use the default SnapVault policy called `XDPDefault`.

### Example

The following command creates a SnapVault policy called “vserverB-vault-policy”:

```
vserverB::> snapmirror policy create -vserver vserverB -policy
vserverB-vault-policy
```

5. Add the `snapmirror-label` attribute to the SnapVault policy you created by using the `snapmirror policy add-rule` command.

If you used the `XDPDefault SnapMirror` policy, you do not need to perform this step. The `XDPDefault SnapVault` policy uses the `daily` and `weekly` `snapmirror-label` attributes specified by the default Snapshot copy policy.

### Example

The following command adds a rule to the `vserverB-vault-policy` to transfer Snapshot copies with the “weekly” `snapmirror-label` attribute and to keep 40 Snapshot copies:

```
vserverB::> snapmirror policy add-rule -vserver vserverB -policy
vserverB-vault-policy -snapmirror-label weekly -keep 40
```

6. On the destination SVM, create a SnapVault relationship and assign a SnapVault policy by using the `snapmirror create` command with the `type XDP` parameter and the `policy` parameter.

In the path specification, a single name is interpreted as a volume name in the SVM from which the command is executed. To specify a volume in a different SVM or in a different cluster, you must specify the full path name.

### Example

The following command creates a SnapVault relationship between the primary volume “srcvolA” on SVM “vserverA” and the empty secondary volume “dstvolB” on SVM “vserverB”. It assigns the SnapVault policy named “vserverB-vault-policy” and uses the “weekendcron” schedule:

```
vserverB::> snapmirror create -source-path vserverA:srcvolA
-destination-path vserverB:dstvolB -type XDP -policy
vserverB-vault-policy -schedule weekendcron
```

7. On the destination SVM, initialize the SnapVault relationship by using the `snapmirror initialize` command to start a baseline transfer.

The command creates a new Snapshot copy that is transferred to the secondary volume and used as a baseline for subsequent incremental Snapshot copies. The command does not use any Snapshot copies that currently exist on the primary volume.

**Note:** Creating a baseline for a large amount of data might take a while.

### Example

The following command begins the relationship initialization by creating and transferring a baseline Snapshot copy to the destination volume “dstvolB” on SVM “vserverB”:

```
vserverB::> snapmirror initialize -destination-path
vserverB:dstvolB
```

## Related concepts

[Guidelines for creating SnapVault relationships on FlexVol volumes](#) on page 81

## Related references

[Commands for managing mirror and SnapVault policies](#) on page 51

## Creating the SnapVault relationship of a mirror-SnapVault cascade

The SnapVault relationship of a mirror-SnapVault cascade requires a different configuration from a SnapVault relationship that is not a part of a mirror-SnapVault cascade.

### Before you begin

- You must have cluster administrator privileges to perform this task for a cluster, and Storage Virtual Machine (SVM) administrator privileges to perform this task for an SVM.
- If the primary and secondary volumes are in different SVMs, the SVMs must be in a peer relationship.

If the primary and secondary volumes are in different clusters, the clusters must be in a peer relationship.

For information about creating peer relationships, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

### About this task

The Snapshot copies that are exported to the mirror destination are ones that are created by Data ONTAP. These Snapshot copies are called “sm\_created” Snapshot copies. Only these Snapshot copies are replicated from the mirror to the SnapVault backup. If the default SnapVault policy is used, the SnapVault secondary accumulates up to 251 “sm\_created” Snapshot copies. The next Snapshot copy transferred after this limit is reached will be added and the oldest “sm\_created” Snapshot copy will be rotated out. You can manage this retention and rotation behavior by adding a rule specifying the “sm\_created” SnapMirror label to the default SnapVault policy.

### Steps

1. On the destination SVM, create a SnapVault secondary volume with a volume type DP.

For information about creating a FlexVol volume, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

2. Create a SnapVault policy by using the `snapmirror policy create` command, or use the default SnapVault policy called `XDPDefault`.

### Example

This procedure uses the `XDPDefault` policy in the examples.

3. Add the `sm_created` `snapmirror-label` to the SnapVault policy by using the `snapmirror policy add-rule` command.

Only the `sm_created` rule is observed. Any other rules associated with the SnapVault policy, such as the `daily` or `weekly` rule, are disregarded.

**Example**

The following command adds a rule to the `XDPDefault` policy to retain 40 `sm_created` Snapshot copies on the SnapVault secondary:

```
vserverB::> snapmirror policy add-rule -vserver vserverC -policy
XDPDefault -snapmirror-label sm_created -keep 40
```

4. On the destination SVM, create a SnapVault relationship and assign a SnapVault policy by using the `snapmirror create` command with the `type XDP` parameter and the `policy` parameter.

**Example**

The following command creates a SnapVault relationship between the primary volume `srcvolB` on SVM `vserverB` and the empty secondary volume `dstvolC` on SVM `vserverC`. It assigns the SnapVault policy named `XDPDefault`:

```
vserverC::> snapmirror create -source-path vserverB:srcvolB
-destination-path vserverC:dstvolC -type XDP -policy XDPDefault
```

5. On the destination SVM, initialize the SnapVault relationship by using the `snapmirror initialize` command to start a baseline transfer.

**Note:** Creating a baseline for a large amount of data can take many hours.

**Example**

The following command begins the relationship initialization by creating and transferring a baseline Snapshot copy to the secondary volume `dstvolC` on SVM `vserverC`:

```
vserverC::> snapmirror initialize -destination-path
vserverC:dstvolC
```

**Related tasks**

*[How a mirror-SnapVault cascade works](#)* on page 38

**Preserving a Snapshot copy on the primary source volume**

In a mirror-SnapVault cascade, you must preserve a Snapshot copy on the primary source volume until it transfers to the secondary volume of the SnapVault backup. For example, you want to ensure that application-consistent Snapshot copies are backed up.

**Before you begin**

You must have created the mirror-SnapVault cascade.



**Steps**

1. Ensure that the Snapshot copy you want to preserve has a `snapmirror-label` by using the `volume snapshot show` command.
2. If the Snapshot copy does not have a `snapmirror-label` associated with it, add one by using the `volume snapshot modify` command.

**Example**

The following command adds a `snapmirror-label` called “`expl`” to the Snapshot copy called “`snapappa`”:

```
clust1::> volume snapshot modify -volume vol1 -snapshot snapappa
-snapmirror-label expl
```

3. Preserve the Snapshot copy on the source volume by using the `snapmirror snapshot-owner create` command to add an owner name to the Snapshot copy.

**Example**

The following command adds `ApplicationA` as the owner name to the `snap1` Snapshot copy in the `testvol` volume on the `vs1` Storage Virtual Machine (SVM):

```
clust1::> snapmirror snapshot-owner create -vserver vs1 -volume vol1
-snapshot snapappa -owner ApplicationA
```

4. Update the destination volume of the data protection mirror relationship by using the `snapmirror update` command.  
  
Alternatively, you can wait for the scheduled update of the data protection mirror relationship to occur.
5. Update the secondary volume of the SnapVault relationship to transfer the specific Snapshot copy from the SnapMirror destination volume to the SnapVault secondary volume by using the `snapmirror update` command with the `-source-snapshot` parameter.
6. Remove the owner name from the primary source volume by using the `snapmirror snapshot-owner delete` command.

**Example**

The following command removes `ApplicationA` as the owner name to the `snap1` Snapshot copy in the `testvol` volume on the `vs1` SVM:

```
clust1::> snapmirror snapshot-owner delete -vserver vs1 -volume vol1
-snapshot snapappa -owner ApplicationA
```

## Creating a SnapVault backup in a prepopulated FlexVol volume

You can protect data that has long-term storage requirements on a FlexVol Volume by replicating selected Snapshot copies to a SnapVault backup on another Storage Virtual Machine (SVM) or cluster. The SnapVault secondary volume might contain data that already exists from a previous data protection mirror or SnapVault relationship or has been loaded from a tape backup.

### Before you begin

- You must have cluster administrator privileges to perform this task for a cluster, and you must have SVM administrator privileges to perform this task for an SVM.
- If the primary and secondary volumes are in different SVMs, the SVMs must be in a peer relationship.

If the primary and secondary volumes are in different clusters, the clusters must be in a peer relationship.

For information about creating peer relationships, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

- The secondary volume must be prepopulated with data.
- A SnapVault policy must exist.

You must either create one or accept the default SnapVault policy (named `XDPDefault`) that is automatically assigned.

The SnapVault policy configuration includes the `snapmirror-label` attribute that is used to select Snapshot copies on the primary volume and match Snapshot copies between the primary and secondary volumes. Only Snapshot copies with the labels configured in the SnapVault policy rules are replicated in SnapVault operations.

- The Snapshot policy assigned to the primary volume must include the `snapmirror-label` attribute.

The name specified in the `snapmirror-label` attribute for the new Snapshot policy must match the `snapmirror-label` attribute that is specified in the SnapVault policy. This ensures that all subsequent Snapshot copies created on the primary volume have labels that are recognized by the SnapVault policy.

You can create a new Snapshot policy by using the `volume snapshot policy add-schedule`, or you can modify an existing Snapshot policy by using the `volume snapshot policy modify-schedule` command to set the `snapmirror-label` attribute for the set of Snapshot copies that you want replicated to the SnapVault secondary volume. Other Snapshot copies on the primary volume are ignored by the SnapVault relationship.

- Your work environment must be able to accommodate the time it might take to transfer a baseline Snapshot copy with a large amount of data.

### Step

1. On the destination SVM, establish the relationship by using the `snapmirror resync` command and the `-type XDP` parameter.

If the most recent common Snapshot copy between the primary and the secondary is deleted from the primary but there exists another, older common Snapshot copy, you can also use the `-preserve` option. This option performs a logical local rollback to make the data in the primary and the secondary the same, and then it replicates all newer Snapshot copies from the source that match the SnapVault policy.

### Example

The following command creates a SnapVault relationship between the primary volume `srcvolA` on SVM `vserverA` and the prepopulated secondary volume `dstvolB` on SVM `vserverB`:

```
vserverB::> snapmirror resync -source-path vserverA:srcvolA -  
destination-path vserverB:dstvolB -type XDP
```

### Related concepts

[Guidelines for creating SnapVault relationships on FlexVol volumes](#) on page 81

[Prepopulated SnapVault secondary scenarios](#) on page 83

### Related tasks

[Creating a destination baseline using a tape backup](#) on page 91

### Related references

[Commands for managing mirror and SnapVault policies](#) on page 51

## Creating a destination baseline using a tape backup

You can perform a baseline transfer from local tape copies to a SnapVault secondary volume to manage your bandwidth or timing constraints over a network.

### Before you begin

- You must have cluster administrator privileges to perform this task for a cluster.
- You must have Storage Virtual Machine (SVM) administrator privileges to perform this task for an SVM.
- The destination volume must not contain data.

### About this task

This operation physically copies data from tape to one or more secondary volumes. When the operation finishes, the secondary volume contains all the Snapshot copies that existed on the primary volume at the time the tape copy was created.

### Steps

1. Create a copy of the primary volume on the tape by using the `system smtape backup` command.

For information about backing up and restoring from tape, see [Performing tape seeding using SMTape](#) on page 63.

2. Restore the data to the empty secondary volume from the tape copy.

For information about backing up and restoring from tape, see [Performing tape seeding using SMTape](#) on page 63.

3. Initialize the SnapVault relationship by using the `snapmirror resync` command with the `-typeXDP` parameter on the secondary volume, and enable incremental updates.

## Converting a data protection destination to a SnapVault secondary

You convert a data protection destination volume to a SnapVault secondary volume after a tape seeding operation or after you lose a SnapVault secondary volume in a backup to disaster protection mirror cascade.

### Before you begin

- You must have cluster administrator privileges to perform this task for a cluster.
- You must have Storage Virtual Machine (SVM) administrator privileges to perform this task for an SVM.

### About this task

In the case of tape seeding, after you transfer the data from the tape to the volume, the volume is a data protection destination volume.

In the case of a SnapVault secondary volume to disaster protection volume cascade, if the SnapVault secondary volume is lost, you can resume SnapVault protection by creating a direct relationship between the SnapVault primary volume and the disaster protection destination volume. You must make the disaster protection destination volume a SnapVault secondary volume to do this.

### Steps

1. Break the data protection mirror relationship by using the `snapmirror break` command.  
The relationship is broken and the disaster protection volume becomes a read-write volume.
2. Delete the existing data protection mirror relationship, if one exists, by using the `snapmirror delete` command.
3. Remove the relationship information from the source SVM by using the `snapmirror release` command.  
This also deletes the Data ONTAP created Snapshot copies from the source volume.
4. Create a SnapVault relationship between the primary volume and the read-write volume by using the `snapmirror create` command with the `-type XDP` parameter.
5. Convert the destination volume from a read-write volume to a SnapVault volume and establish the SnapVault relationship by using the `snapmirror resync` command.

## Managing backup and restore operations for SnapVault backups

You configure SnapVault relationships on FlexVol volumes to establish SnapVault backups. You manage SnapVault relationships to optimize the performance of the relationships.

### Backing up from a Snapshot copy that is older than the base Snapshot copy

You might want to replicate a special, manually initiated Snapshot copy to the SnapVault backup. The Snapshot copy is one that is not in the sequence scheduled by the SnapVault policy assigned to the SnapVault relationship.

#### Before you begin

You must have cluster administrator privileges to perform this task for a cluster. You must have Storage Virtual Machine (SVM) administrator privileges to perform this task for an SVM.

#### Step

1. Begin the backup transfer of the older Snapshot copy by using the `snapmirror update` command.

#### Example

The following command starts an out-of-order transfer of Snapshot copy SC3 from the source volume `srcvolA` on SVM `vserverA` and the secondary volume `dstvolB` on SVM `vserverB`:

```
vserverA::> snapmirror update -source-path vserverA:srcvolA -  
destination-path vserverB:dstvolB -snapshot SC3
```

#### Result

After the backup finishes, the transferred Snapshot copy becomes the base.

### How an out-of-order Snapshot copy transfer works

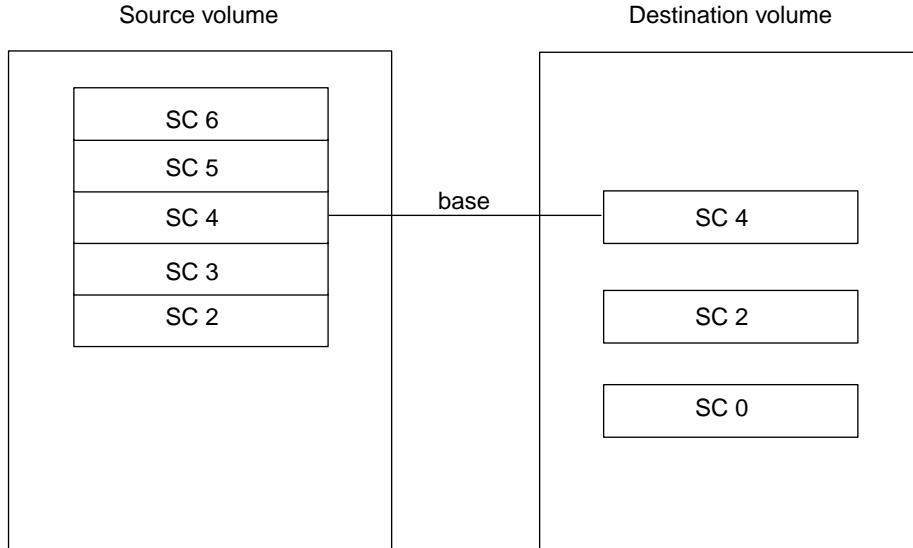
The transfer of a Snapshot copy that does not conform to the usual sequence scheduled by a SnapVault policy is an out-of-order Snapshot copy transfer.

In SnapVault relationships, Snapshot copies are selected and transferred from the primary volume to the secondary volume, according to the configured SnapVault policy. Only Snapshot copies that are newer than the common Snapshot copy between the primary and secondary volume are transferred. However, you can use the `snapmirror update` command to initiate the transfer of a Snapshot copy that was not originally selected and transferred.

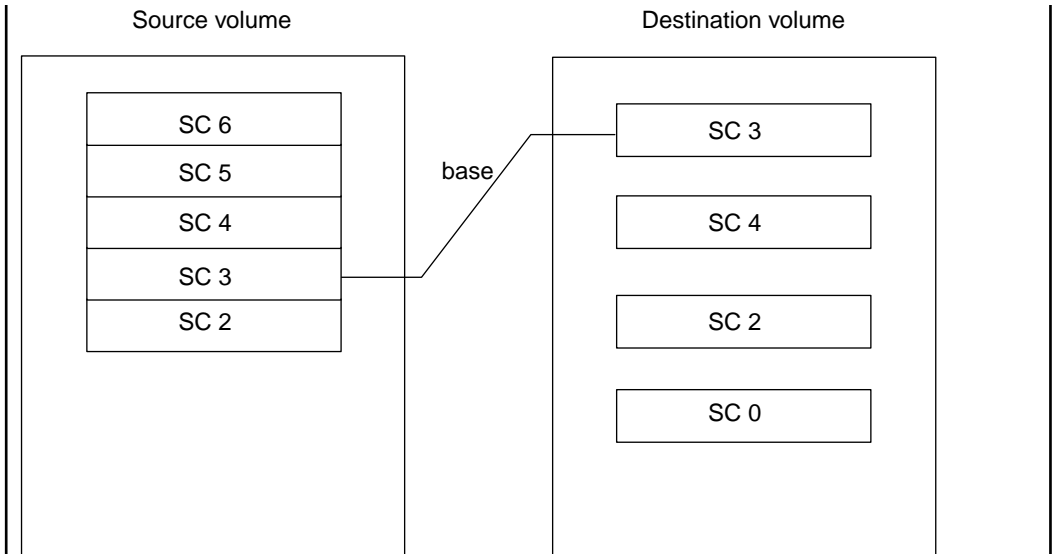
When you initiate an out-of-order transfer, an older Snapshot copy is used to establish the base. To avoid subsequent transfers of Snapshot copies that already exist on the SnapVault secondary volume, the list of Snapshot copies that are selected for transfer in this update cycle are reconciled against the Snapshot copies that are already present on the secondary volume. Snapshot copies that are already present on the secondary volume are discarded from the transfer list.

### Example of a new base that is established from an out-of-order Snapshot copy transfer

In this example, the SnapVault policy has a schedule in which only the even-numbered Snapshot copies on the primary volume are transferred to the secondary volume. Before the out-of-order transfer begins, the primary volume contains Snapshot copies 2 through 6; the secondary volume contains only the even-numbered Snapshot copies (noted as “SC” in the figures). Snapshot copy 4 is the common Snapshot copy that is used to establish the base, as shown in the following figure:

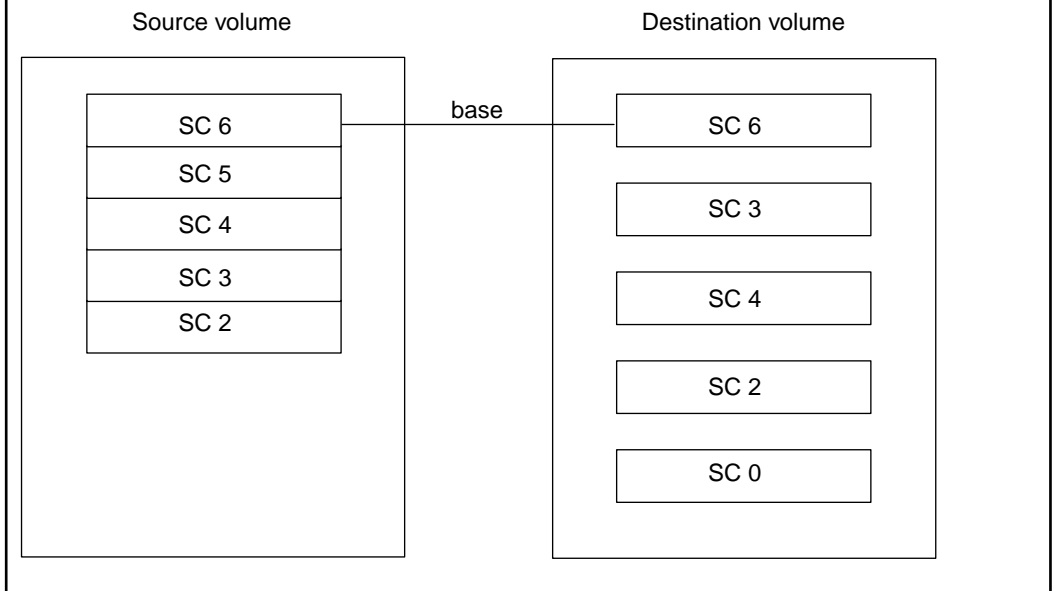


After Snapshot copy 3 is transferred to the secondary volume, out of order, it becomes the new common Snapshot copy that is used to establish the base, as shown in the following figure:



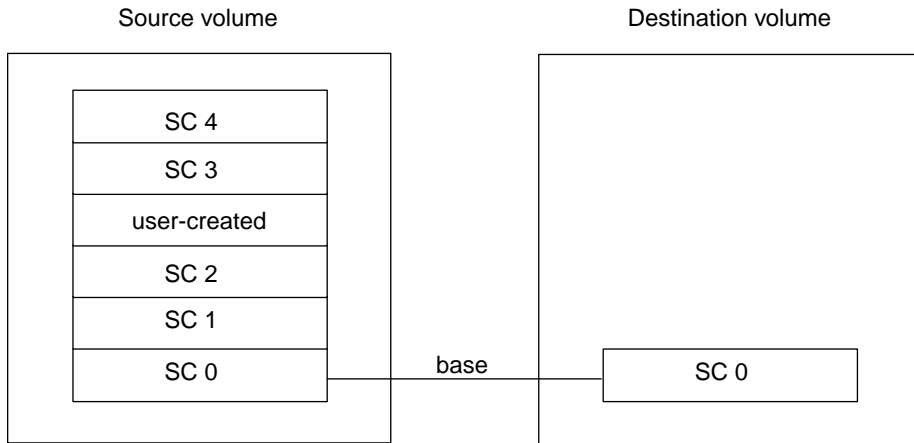
**Note:** Although Snapshot copy 3 is now the base, the exported Snapshot copy is still Snapshot copy 4.

When Snapshot copies are selected for subsequent updates according to the SnapVault policy, the policy selects Snapshot copy 4 and Snapshot copy 6 for transfer to the secondary volume. When the transfer list is reconciled, Snapshot copy 4 is removed from the transfer list because it already exists on the secondary volume. Only Snapshot copy 6 is transferred, which becomes the new common Snapshot copy that is used to establish the base, as shown in the following figure:

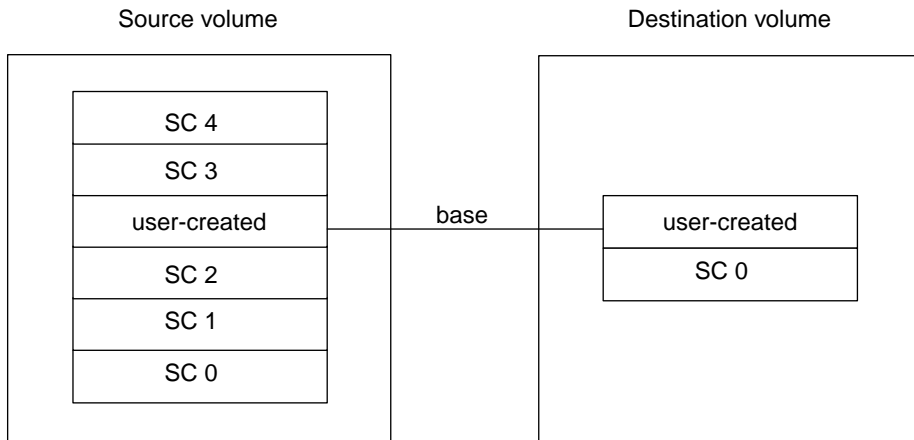


### Example of SnapVault transfer behavior with an out-of-order Snapshot copy transfer

In this example, Data ONTAP created two SnapVault-labeled Snapshot copies, you made a user-created Snapshot copy, and then more SnapVault-labeled Snapshot copies were created. The order of Snapshot copies made would appear as shown in the following figure:

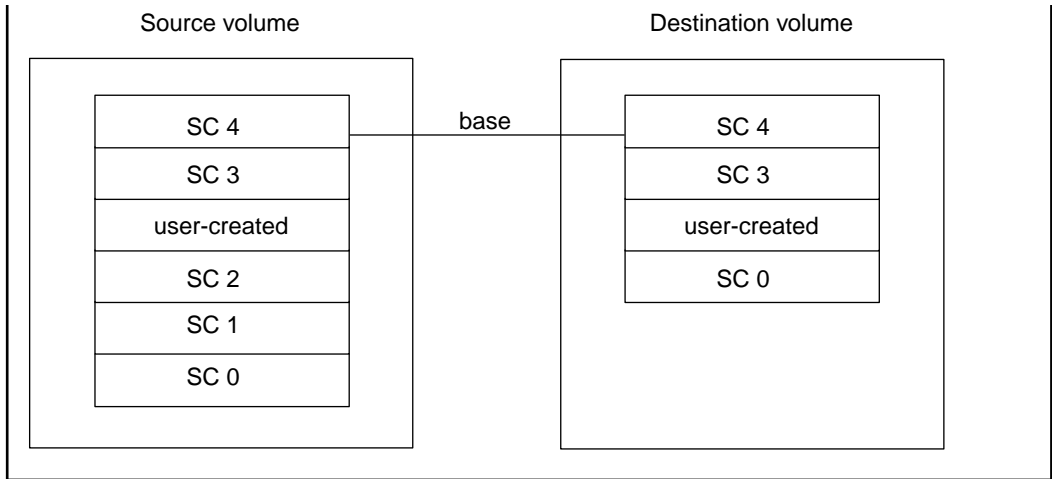


You perform an out-of-order Snapshot transfer using the user-created Snapshot copy, which establishes that Snapshot copy as the new base Snapshot copy, as shown in the following figure:



When the next SnapVault scheduled transfer occurs, only the SnapVault labeled Snapshot copies made after the user-created Snapshot copy are transferred. This occurs because the Snapshot copies created between the previous base Snapshot copy and the current base Snapshot copy are not transferred.





## Backing up FlexVol volumes that contain the maximum limit of Snapshot copies

To work around the limit of 251 Snapshot copies per volume, you can create a new destination volume clone, then establish a SnapVault relationship with the new clone.

### Before you begin

You must have cluster administrator privileges to perform this task for a cluster. You must have Storage Virtual Machine (SVM) administrator privileges to perform this task for an SVM.

### About this task

Creating a new SnapVault relationship to a new volume clone enables you to continue SnapVault protection with minimum disruption on the clone volume and without starting a new baseline transfer. Because the source clone and the volume clone share the latest common Snapshot copy, subsequent updates are performed as usual, according to the policy assigned to the SnapVault relationship.

### Steps

1. Quiesce the SnapVault relationship between the primary volume and the secondary volume by using the `snapmirror quiesce` command.

This step prevents updates from starting until after the task is complete.

2. Verify that there are no active transfers on the relationship by using the `snapmirror show` command.

The Relationships field should be `Idle`.

3. Create a volume clone based on the most recent common Snapshot copy between the SnapVault primary volume and the SnapVault secondary volume by using the `volume clone create` command with the `-type DP` parameter.
4. Establish the SnapVault relationship between the primary volume and the newly created secondary volume clone by using the `snapmirror resync` command and the `-type XDP` parameter.
5. Delete the SnapVault relationship between the primary volume and the original SnapVault secondary volume by using the `snapmirror delete` command.

## Managing the backup of a copied source volume

If you use the `volume copy` command to copy the primary volume of a SnapVault relationship to a different volume, Data ONTAP does not copy SnapMirror labels for Snapshot copies, and you lose the capability to back up from the primary volume copy.

### About this task

You must add the SnapMirror labels back before you can back up the volume copy.

### Step

1. Add the SnapMirror labels to the copied volume by using the `volume snapshot modify` command or by using the `snapmirror update -s` command.

## Guidelines for restoring the active file system

The restore operation from a SnapVault backup copies a single, specified Snapshot copy from a SnapVault secondary volume to a specified volume. Restoring a volume from a SnapVault secondary volume changes the view of the active file system but preserves all earlier Snapshot copies in the SnapVault backup.

Before restoring a volume, you must shut down any application that accesses data in a volume to which a restore is writing data. Therefore, you must dismount the file system, shut down any database, and deactivate and quiesce the Logical Volume Manager (LVM) if you are using an LVM.

The restore operation is disruptive. When the restore operation finishes, the cluster administrator or Storage Virtual Machine (SVM) administrator must remount the volume and restart all applications that use the volume.

The restore destination volume must not be the destination of another mirror or the secondary of another SnapVault relationship.

You can restore to the following volumes:

- Original source volume  
You can restore from a SnapVault secondary volume back to the original SnapVault primary volume.
- New, empty secondary volume

You can restore from a SnapVault secondary volume to a new, empty secondary volume. You must first create the volume as a data protection (DP) volume.

- New secondary that already contains data

You can restore from a SnapVault secondary volume to a volume that is prepopulated with data. The volume must have a Snapshot copy in common with the restore primary volume and must not be a DP volume.

## Guidelines for restoring LUNs in SAN environments

The restore operation from a SnapVault backup copies a single, specified LUN from a SnapVault secondary volume to a specified volume. Restoring a LUN from a SnapVault secondary volume changes the view of the active system on the volume to which data is being restored, preserving all earlier Snapshot copies.

The following guidelines apply only to SAN environments:

- You can restore a single file or single LUN from a SnapVault secondary volume by using the NetApp OnCommand management software online management tools.
- When LUNs are restored to existing LUNs, new access controls do not need to be configured. You must configure new access controls for the restored LUNs only when restoring LUNs as newly created LUNs on the volume.
- If LUNs on the SnapVault secondary volume are online and mapped before the restore operation begins, they remain so for the duration of the restore operation and after the operation finishes.
- The host system can discover the LUNs and issue non-media access commands for the LUNs, such as inquiries or commands to set persistent reservations, while the restore operation is in progress.
- You cannot create new LUNs in a volume during a restore operation with the `lun create` command.
- Restore operations from tape and from a SnapVault backup are identical.
- You cannot restore a single LUN from a SnapVault secondary volume that is located on a system that is running in 7-Mode.

For more information about backing up and restoring data in a SAN environment, see the *Clustered Data ONTAP SAN Administration Guide*.

## How restore operations work from a SnapVault backup

A restore operation from a SnapVault backup consists of a series of actions performed on a temporary restore relationship and on the secondary volume.

During a restore operation, the following actions occur:

1. A new temporary relationship is created from the restore source (which is the original SnapVault relationship secondary volume) to the restore destination.  
The temporary relationship is a restore type (RST). The `snapmirror show` command displays the RST type while the restore operation is in progress.

The restore destination might be the original SnapVault primary or might be a new SnapVault secondary.

2. During the restore process, the restore destination volume is changed to read-only.
3. When the restore operation finishes, the temporary relationship is removed and the restore destination volume is changed to read-write.

## Restoring a volume from a SnapVault backup

If the data on a volume becomes unavailable, you can restore the volume to a specific time by copying a Snapshot copy in the SnapVault backup. You can restore data to the same primary volume or to a new location. This is a disruptive operation.

### Before you begin

- You must have cluster administrator privileges to perform this task for a cluster.
- You must have Storage Virtual Machine (SVM) administrator privileges to perform this task for an SVM.
- CIFS traffic must not be running on the SnapVault primary volume when a restore operation is running.

### About this task

This task describes how to restore a whole volume from a SnapVault backup. To restore a single file or LUN, you can restore the whole volume to a different, non-primary volume, and then select the file or LUN, or you can use the NetApp OnCommand management software online management tools.

### Steps

1. If the volume to which you are restoring has compression enabled and the secondary volume from which you are restoring does not have compression enabled, disable compression.  
You disable compression to retain storage efficiency during the restore.
2. Restore a volume by using the `snapmirror restore` command.

### Example

```
vs1:~> snapmirror restore -destination-path
vs1:vol1
  -source-path vs2:vol1_dp_mirror2 -source-snapshot
snap3
Warning: All data newer than Snapshot copy snap6 on volume
vs1:vol1
      will be
deleted.

Do you want to continue? {y|n}:
y
```

```
[Job 34] Job is queued: snapmirror restore from source
vs2:vol1_dp_mirror2 for the snapshot snap3.
```

For more information about the `snapmirror restore` command, see the man pages.

3. If the volume had quotas before the restore operation, activate the quotas on the restored volume by using the `volume quota modify` command with the `-state` parameter.

Quotas are not turned on when you restore a volume.

4. Remount the restored volume and restart all applications that use the volume.
5. If you previously disabled compression, reenable compression on the volume.

### Related concepts

[Guidelines for restoring the active file system](#) on page 98

[Guidelines for restoring LUNs in SAN environments](#) on page 99

[How restore operations work from a SnapVault backup](#) on page 99

## Managing a SnapVault-mirror cascade when the SnapVault backup is unavailable

You can manipulate relationships in a SnapVault-mirror cascade to maintain data backup relationships if the secondary of the SnapVault relationship becomes unavailable.

### Before you begin

You must have a SnapVault-mirror cascade already configured.

### About this task

The destination of the SnapVault relationship is the middle of the SnapVault-mirror cascade. If it becomes unavailable, you might have the following issues:

- You cannot update the SnapVault backup.
- You cannot update the mirror copy of the SnapVault secondary.

To manage this issue, you can temporarily remove the SnapVault secondary volume from the cascade and establish a SnapVault relationship to the mirror copy of the SnapVault secondary volume. When the unavailable secondary volume becomes available, you can reestablish the original cascade configuration.

In the following steps, the primary volume of the cascade is called “A”, the secondary volume of the SnapVault relationship is called “B”, and the destination volume of the data protection mirror relationship is called “C”.

**Steps**

1. Identify the current exported Snapshot copy on C by using the `volume snapshot show` command with the `-fields busy` parameter.

The busy field is set to `true` for the exported Snapshot copy.

**Example**

```
volume snapshot show C -fields busy
```

2. Break the data protection mirror relationship by using the `snapmirror break` command on C.

**Example**

```
snapmirror break C
```

3. Create a dummy `snapmirror-label` on the exported Snapshot copy you previously identified by using the `volume snapshot modify` command with the `-snapmirror-label` parameter.

If a `snapmirror-label` already exists for the exported Snapshot copy, you do not need to perform this step.

**Example**

```
volume snapshot modify -volume C -snapshot name -snapmirror-label expl
```

4. Create a Snapshot owner on the exported Snapshot copy of C by using the `snapmirror snapshot-owner create` command.

This prevents clustered Data ONTAP from deleting the Snapshot copy.

**Example**

```
snapmirror snapshot-owner create -volume C -snapshot exported -owner admin1
```

5. Delete the data protection mirror relationship between B and C by using the `snapmirror delete` command.

**Example**

```
snapmirror delete C
```

6. Create the SnapVault relationship between A and C by using the `snapmirror resync` command and the `-type XDP` parameter.

**Example**

```
snapmirror resync -source-path A -destination-path C -type XDP
```

You can maintain this SnapVault relationship until you recover the original SnapVault secondary volume. At that time, you can reestablish the original cascade relationship by using the steps that follow this step.

7. Delete the data protection mirror relationship between A and B by using the `snapmirror delete` command.
8. Perform a disaster recovery resynchronization from C to B by using the `snapmirror resync` command.

#### Example

```
snapmirror resync -source-path C -destination-path B
```

This step copies from C to B, all of the Snapshot copies made after B became unavailable.

9. Identify the current exported Snapshot copy on B by using the `volume snapshot show` command with the `-fields busy` parameter.

#### Example

```
volume snapshot show B -fields busy
```

The `busy` field is set to `true` for the exported Snapshot copy.

10. Break the data protection mirror relationship by using the `snapmirror break` command on B.

#### Example

```
snapmirror break B
```

11. Create a dummy `snapmirror-label` on the exported Snapshot copy you previously identified by using the `volume snapshot modify` command with the `-snapmirror-label` parameter.

#### Example

```
volume snapshot modify -volume B -snapshot name -snapmirror-label exp2
```

If a `snapmirror-label` already exists for the exported Snapshot copy, you do not need to perform this step.

12. Create a Snapshot owner on the exported Snapshot copy of B by using the `snapmirror snapshot-owner create` command.

This prevents clustered Data ONTAP from deleting the Snapshot copy.

#### Example

```
snapmirror snapshot-owner create -volume B -snapshot exported -owner admin1
```

13. Delete the data protection mirror relationship between C and B by using the `snapmirror delete` command.

14. Perform a SnapVault resynchronization from A to B by using the `snapmirror resync` command and the `-type XDP` parameter.

**Example**

```
snapmirror resync -source-path A -destination-path B -type XDP
```

New Snapshot copies that meet the Snapshot policy of the SnapVault relationship are transferred from A to B.

15. Delete the data protection mirror relationship between A and C by using the `snapmirror delete` command.
16. Perform a disaster recovery resynchronization from B to C by using the `snapmirror resync` command.

This step copies from B to C, all of the Snapshot copies made after reestablishing the A to B relationship without deleting any Snapshot copies on C.

**Example**

```
snapmirror resync -source-path B -destination-path C
```

17. Remove the Snapshot copy owner from volumes B and C by using the `snapmirror snapshot-owner delete` command.

**Example**

```
snapmirror snapshot-owner delete -volume B -snapshot exported_snap
```

18. Remove SnapMirror labels that you created from volumes B and C by using the `snapshot modify` command.

**Example**

```
snapshot modify -volume B -snapshot exported_snap -snapmirror-label text
```

**Example**

```
snapshot modify -volume C -snapshot exported_snap -snapmirror-label text
```

## Managing storage efficiency for SnapVault secondaries

SnapVault relationships preserve storage efficiency when backing up data from the primary volume to the secondary volume, with one exception: if post-process and optionally inline compression are



enabled on the secondary volume, storage efficiency is not preserved for data transfers between the primary and secondary volumes.

## Guidelines for managing storage efficiency for SnapVault backups

If both the primary and secondary volumes in a SnapVault relationship have storage efficiency enabled, then data transfers to the SnapVault secondary volume preserve storage efficiency. If the primary volume does not have storage efficiency enabled, you might want to enable storage efficiency only on the secondary volume.

Because SnapVault secondary volumes typically contain a large amount of data, storage efficiency on SnapVault secondary volumes can be very important.

### If storage efficiency is enabled on the primary volumes

If the primary volume in a SnapVault relationship is enabled for storage efficiency, all data backup operations preserve the storage efficiency.

### If storage efficiency is enabled only on the secondary volume

If the primary volume in a SnapVault relationship does not have storage efficiency enabled, you might want to enable storage efficiency for the secondary volume because it is likely to contain a large amount of data over time.

You can use the `volume efficiency` command to start a scan on the volume if there is already data on the volume from transfers. If this is a new relationship with no transfers, then there is no need to run the scan manually.

Changes to the volume's efficiency schedule do not take effect for a SnapVault secondary volume. Instead, when storage efficiency is enabled, the SnapVault relationship manages the schedule. When a data transfer begins, the storage efficiency process automatically pauses until the transfer is finished, and then automatically begins again after the data transfer is complete. Because data transfers to a SnapVault secondary volume might include more than one Snapshot copy, the storage efficiency process is paused for the entire duration of the update operation. After the transfer is finished and the post-transfer storage efficiency process is complete, the last Snapshot copy created in the secondary volume is replaced by a new, storage-efficient Snapshot copy.

If the last Snapshot copy that is created in the secondary volume is locked before it can be replaced by a new, storage-efficient Snapshot copy, then a new, storage-efficient Snapshot copy is still created, but the locked Snapshot copy is not deleted. That Snapshot copy is deleted later during the storage-efficient cleanup process after a subsequent update to the SnapVault secondary volume and after the lock is released. A Snapshot copy in a SnapVault secondary volume might be locked because the volume is the source in another relationship, such as a data protection mirror relationship.

**If the secondary volume has additional compression enabled, storage efficiency is not preserved**

Storage efficiency on all data transfers in SnapVault relationships is not preserved when the secondary volume has additional compression enabled. Because of the loss of storage efficiency, a warning message is displayed when you enable compression on a SnapVault secondary volume. After you enable compression on the secondary volume, you can never have storage-efficient transfers.

**Related tasks**

*Enabling storage efficiency on a SnapVault secondary volume* on page 106

**Enabling storage efficiency on a SnapVault secondary volume**

If the primary volume does not have storage efficiency enabled, you can enable storage efficiency on a SnapVault secondary volume by enabling storage efficiency on the volume.

**Before you begin**

You must have cluster administrator privileges to perform this task for a cluster. You must have Vserver administrator privileges to perform this task for an SVM.

**About this task**

For information about increasing storage efficiency using deduplication and compression, see the *Clustered Data ONTAP Logical Storage Management Guide*.

**Steps**

1. Use the `volume efficiency` command with the `-on` parameter to enable storage efficiency.
2. If the volume already has data which you want to make storage efficient, use the `volume efficiency` command with the `-start` and `-scan-old-data` parameters to start a scan of the volume.

**Related concepts**

*Guidelines for managing storage efficiency for SnapVault backups* on page 105

## 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, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

## 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

- A**
- about
    - snapshot copy reserve [20](#)
  - active file disk space
    - how Data ONTAP uses deleted [21](#)
  - active file systems
    - access on destination volumes [25](#)
    - guidelines for restoring from SnapVault backups [98](#)
    - when available on destination volumes [25](#)
  - automatic Snapshot copies
    - using prefixes to name [46](#)
  - automatically deleting
    - Snapshot copies [17](#)
- B**
- backup policies
    - example of creating tiered [54](#)
  - backup tasks
    - Snapshot copy [13](#)
  - backups
    - guidelines for managing storage efficiency for SnapVault [105](#)
    - guidelines for restoring active file system from SnapVault [98](#)
    - how restore operations work from SnapVault [99](#)
    - managing copied source volume, in SnapVault relationships [98](#)
    - SnapVault, how they work [31](#)
    - source-to-destination-to-tape, defined [36](#)
  - baseline transfers
    - defined [28](#)
- C**
- cascades
    - creating SnapVault backup in mirror-SnapVault [87](#)
    - managing SnapVault-mirror, when SnapVault backup is unavailable [101](#)
    - mirror-SnapVault, how they work [38](#)
    - SnapVault-SnapMirror, how they work [39](#)
  - CIFS clients
    - destination volume not accessible [27](#)
  - CIFS users
    - accessing Snapshot copies [12](#)
  - client access
    - to active file systems on destination volumes [25](#)
  - cluster administrators
    - mirror and SnapVault relationship management capabilities [52](#)
  - clusters
    - guidelines for creating mirror or SnapVault relationships between [25](#)
    - how SnapMirror policies work with [52](#)
  - commands
    - for managing mirror and SnapVault policies [51](#)
    - for managing SnapMirror relationships [59](#)
    - for managing Snapshot copy policies and schedules [44](#)
    - for monitoring Snapshot copy disk consumption) [49](#)
    - for scheduling when SnapMirror transfers occur [71](#)
    - job show [71](#)
    - snapmirror break [79, 80](#)
    - snapmirror delete [79](#)
    - snapmirror delete, when deleting mirror relationships [74](#)
    - volume delete, when deleting mirror relationships [74](#)
    - volume snapshot autodelete show [18](#)
    - volume snapshot partial-restore-file [48](#)
    - volume snapshot partial-restore-file-list-info [48](#)
    - volume snapshot restore [48](#)
    - volume snapshot restore-file [47](#)
    - volume snapshot restore-file-info [47](#)
  - configurations
    - deployment, how mirror-SnapVault cascades work [38](#)
    - source-to-destination-to-tape, defined [36](#)
    - supported data protection deployment [34](#)
  - copied source volumes
    - managing SnapVault backup of [98](#)
- D**
- data compression
    - how SnapVault backups work with [32](#)
  - data loss
    - tools to protect against [9](#)
  - data loss disaster [8](#)
  - data protection
    - in SAN environments [10](#)
    - levels of provided by mirror relationships [23](#)

- overview of methods for [7](#)
- data protection mirror copies
  - updating manually [73](#)
- data protection mirrors
  - active file systems for [25](#)
  - creating [56](#)
  - FlexVol volumes, support of [23](#)
  - introduction to managing [56](#)
  - making a destination volume writable [79](#)
  - reestablishing the relationship after disaster [75](#)
  - reversing the relationship after disaster [75](#)
  - using to test databases [80](#)
- data protection policies
  - types of [10](#)
- data protection strategy
  - planning [12](#)
- databases
  - protection of [8](#)
- deleting
  - Snapshot copies automatically [17](#)
- deployment configurations
  - basic, defined [36](#)
  - how mirror-mirror cascades work [37](#)
  - how mirror-SnapVault cascades work [38](#)
  - how SnapVault-SnapMirror cascades work [39](#)
  - mirror-SnapVault fanout, defined [39](#)
  - multiple-mirrors fanout, how they work [40](#)
  - source-to-destination-to-tape, defined [36](#)
  - supported data protection [34](#)
- destination volumes
  - access to active file systems on [25](#)
  - CIFS clients cannot access [27](#)
  - components of a mirror relationship [23](#)
  - converting to SnapVault secondary [92](#)
  - matching source volume size [24](#)
  - supported number of SnapMirror relationship fanout [27](#)
- disaster recovery
  - reestablishing the original data protection mirror relationship [75](#)
  - reversing the data protection mirror relationship [75](#)
- disasters
  - tools to protect against data-loss [9](#)
- disk consumption
  - monitoring Snapshot copy [49](#)
- disk space
  - recovery of [22](#)
- disks
  - monitoring Snapshot copy consumption [49](#)
- DP mirror relationships

- where to find information about Infinite Volumes and [24](#)

## E

- error messages
  - correcting restricted volume [58](#)
- extended queries
  - using to operate on many SnapMirror relationships [63](#)

## F

- failures
  - correcting SnapMirror initialization [58](#)
- fanout support
  - SnapMirror relationship [27](#)
- fanouts
  - multiple-mirrors, how they work [40](#)
- file system
  - recovery of disk space for use by the [22](#)
- file systems
  - guidelines for restoring active, from SnapVault backups [98](#)
- files
  - restoring from Snapshot copies of FlexVol volumes [47](#)
- firewalls
  - intercluster SnapMirror relationship requirements [36](#)
- FlexVol volumes
  - commands for managing SnapMirror relationships of [59](#)
  - creating SnapVault backup in prepopulated [90](#)
  - creating SnapVault backups in empty [84](#)
  - guidelines for creating SnapVault relationships on [81](#)
  - mirror relationships for [23](#)
  - restoring part of a file [48](#)
  - SnapVault backup limitations [32](#)
  - which data gets backed up and restored from [29](#)
  - which data is not backed up to SnapVault backups [30](#)

## I

- incremental transfers
  - defined [28](#)
- Infinite Volumes
  - commands for managing SnapMirror relationships of [59](#)

- data protection mirror relationships for, where to find information [24](#)
- how Snapshot policies are associated with [43](#)
- maximum number of Snapshot copies for [14](#)
- Snapshot copies of, where to find information [14](#)
- updating data protection mirror copies manually [73](#)
- initialization failures
  - correcting SnapMirror [58](#)
- intercluster networks
  - firewall requirements for SnapMirror relationships [36](#)

## L

- language settings
  - requirement between SVMs [25](#)
- limitations
  - mirror relationship [27](#)
  - SnapVault backup [32](#)
- load-sharing mirrors
  - modifying relationship schedules [72](#)
- LUNs
  - guidelines for restoring in SAN environments [99](#)
  - protecting data [10](#)

## M

- methods
  - overview of data protection [7](#)
- mirror copies
  - manually updating data protection, on volumes [73](#)
- mirror policies
  - commands for managing [51](#)
  - description of [10](#)
- mirror relationships
  - components of [23](#)
  - differences between cluster administrator and SVM administrator policy management privileges [52](#)
  - guidelines for creating mirror or SnapVault, between clusters or SVMs [25](#)
  - limitations [27](#)
- mirror-SnapVault cascades
  - creating SnapVault backup for [87](#)
  - how they work [38](#)
  - preserving Snapshot copies on primary source volume of [88](#)
- mirrors
  - creating data protection [56](#)
  - deleting [74](#)
  - deleting Snapshot copies automatically [27](#)

- listing the state of a scheduled transfer [71](#)
- making a destination volume writable [79](#)
- modifying load-sharing relationship schedules [72](#)
- reestablishing the relationship after disaster [75](#)
- reversing the relationship after disaster [75](#)
- scheduling when SnapMirror transfers occur [71](#)
- Snapshot copy limit [28](#)
- using to test databases [80](#)
- multiple-mirrors fanout deployments
  - how they work [40](#)

## N

- namespaces
  - protecting SVM [32](#)
- naming guidelines
  - SnapMirror and SnapVault policy [53](#)
- naming guidelines for SnapMirror [53](#)
- NFS users
  - accessing Snapshot copies [12](#)
- NVRAM
  - warning of database validity [8](#)

## O

- older-than-base Snapshot copies
  - backing up from [93](#)
- out-of-order Snapshot copies
  - how they work [93](#)

## P

- path names
  - abbreviating [24](#)
  - pattern matching of [24](#)
  - wildcard use for [24](#)
- pattern matching
  - path names [24](#)
- policies
  - commands for managing mirror and SnapVault [51](#)
  - commands for managing Snapshot copy [44](#)
  - differences between cluster administrator and SVM administrator management privileges [52](#)
  - example of creating tiered backup [54](#)
  - introduction to managing data protection using SnapMirror [51](#)
  - introduction to managing Snapshot [43](#)
  - naming guidelines for SnapVault [53](#)
  - SnapMirror, how they work with clusters and SVMs [52](#)



- Snapshot copy schedule and retention planning
  - guidelines for SnapVault backups [33](#)
  - types of data protection [10](#)
- port usage
  - intercluster SnapMirror relationship firewall requirements [36](#)
- prefixes
  - using to name automatic Snapshot copies [46](#)
- primary volumes
  - defined [28](#)
- protection deployments
  - See* deployment configurations
- protection policies
  - types of data [10](#)

**R**

- recovery tasks
  - Snapshot copy [13](#)
- relationships
  - commands for managing SnapMirror [59](#)
  - components of mirror [23](#)
  - creating SnapVault backup, in a prepopulated FlexVol volume [90](#)
  - creating SnapVault backup, in empty FlexVol volumes [84](#)
  - guidelines for creating mirror or SnapVault, between clusters or SVMs [25](#)
  - guidelines for creating SnapVault, on FlexVol volumes [81](#)
  - prepopulated SnapVault secondary scenarios [83](#)
  - SnapVault, creating baseline from tape [91](#)
  - transition (TDP) [7](#)
- requirements
  - SVM language setting [25](#)
- reserve
  - example of what happens when exceeding Snapshot copy [22](#)
- restore operations
  - guidelines for LUNs in SAN environments [99](#)
  - guidelines for restoring active file system from SnapVault backups [98](#)
  - how they work from a SnapVault backup [99](#)
  - restoring volumes from SnapVault backups [100](#)
- restoring Snapshot copies
  - Shadow Copy Client tools [49](#)
- restricted volume errors
  - correcting [58](#)
- root information
  - protecting SVM [32](#)

**S**

- SAN (storage area network)
  - data protection of volumes containing LUNs [10](#)
- SAN environments
  - guidelines for restoring LUNs in [99](#)
  - which LUN data is backed up to SnapVault backups [29](#)
  - which LUN data is not backed up to SnapVault backups [30](#)
- schedules
  - commands for managing Snapshot copy [44](#)
  - creating a Snapshot copy [16](#)
  - for default Snapshot copies [15](#)
  - guidelines for planning Snapshot copy [33](#)
  - strategies for creating Snapshot copy policies [45](#)
- secondary volumes
  - defined [28](#)
  - enabling storage efficiency on SnapVault [106](#)
- SMTape
  - performing tape seeding using [63](#)
- SMTape backup and restore sessions
  - scalability limits for [71](#)
- SnapMirror
  - correcting an initialization failure [58](#)
  - creating a data protection [56](#)
  - deleting a mirror relationship [74](#)
  - firewall requirements for intercluster relationships [36](#)
  - how destination volumes match source volume size [24](#)
  - introduction to managing data protection mirror copies [56](#)
  - listing the state of a scheduled transfer [71](#)
  - modifying relationship schedules [72](#)
  - scheduling when transfers occur [71](#)
  - source and destination port requirements on intercluster relationships [36](#)
  - using extended queries [63](#)
- SnapMirror commands
  - for managing mirror and SnapVault policies [51](#)
- SnapMirror labels
  - defined [28](#)
- SnapMirror policies
  - how they work with clusters and SVMs [52](#)
  - introduction to managing data protection using [51](#)
- SnapMirror policy
  - altering to preserve Snapshot copies after reaching retention limit [53](#)
- SnapMirror relationships

- commands for managing [59](#)
- supported number of fanout volumes [27](#)
- Snapshot copies
  - backing up FlexVol volumes with over 251 [97](#)
  - backing up from older than base [93](#)
  - backup and recovery tasks you can perform with [13](#)
  - commands for managing [42](#)
  - commands for managing policies and schedules [44](#)
  - creating a schedule [16](#)
  - default schedule [15](#)
  - defined [28](#)
  - defining [12](#)
  - deleting automatically [17](#)
  - example of what happens when the reserve is exceeded [22](#)
  - how out-of-order transfers work [93](#)
  - limit on source volume [28](#)
  - maximum number [14](#)
  - monitoring disk consumption of) [49](#)
  - preserving after reaching retention limit [53](#)
  - preserving on primary source volume in mirror-SnapVault cascades [88](#)
  - protecting data, about [42](#)
  - restoring a single file from [47](#)
  - restoring part of a file [48](#)
  - restoring volume contents [48](#)
  - types of user-specified schedules [15](#)
  - user access to [12](#)
  - using prefixes to name automatic [46](#)
  - viewing automatic deletion settings for [18](#)
  - where to find information about Infinite Volumes and [14](#)
- Snapshot copy
  - creating [16](#)
- Snapshot copy reserve
  - how deleted active file disk space consumes [21](#)
- Snapshot policies
  - description of [10](#)
  - how volumes inherit from SVMs [43](#)
  - introduction to managing [43](#)
  - strategies for creating [45](#)
- SnapVault
  - backing up FlexVol volumes with over 255 Snapshot copies [97](#)
  - backing up from older-than-base Snapshot copy [93](#)
  - converting data protection destination to secondary [92](#)
  - creating a baseline from tape [91](#)
  - creating backup, in a prepopulated FlexVol volume [90](#)
  - differences between cluster administrator and SVM administrator policy management privileges [52](#)
  - which data is backed up and restored from FlexVol volume [29](#)
  - which data is not backed up to [30](#)
- SnapVault backups
  - creating in empty FlexVol volumes [84](#)
  - creating in mirror-SnapVault cascades [87](#)
  - guidelines for managing storage efficiency for [105](#)
  - guidelines for planning Snapshot copy schedule and retention for [33](#)
  - guidelines for restoring active file system from [98](#)
  - guidelines for restoring LUNs in SAN environments [99](#)
  - how restore operations work from [99](#)
  - how they work [31](#)
  - how they work with data compression [32](#)
  - introduction to protecting FlexVol volume data using [81](#)
  - limitations for FlexVol volume backup [32](#)
  - managing SnapVault-mirror cascade when unavailable [101](#)
  - preserving Snapshot copies on primary source volume in mirror-SnapVault cascades [88](#)
  - restoring volumes from [100](#)
  - SVM namespace [32](#)
- SnapVault policies
  - commands for managing [51](#)
  - description of [10](#)
  - how out-of-order Snapshot copy transfers work with [93](#)
- SnapVault relationship
  - configuring policy rules to preserve Snapshot copies [53](#)
- SnapVault relationships
  - defined [28](#)
  - example of creating tiered backup policy for [54](#)
  - guidelines for creating mirror or SnapVault, between clusters or SVMs [25](#)
  - guidelines for creating on FlexVol volumes [81](#)
  - managing backup of copied source volumes in [98](#)
  - prepopulated secondary scenarios [83](#)
- SnapVault secondary volumes
  - enabling storage efficiency on [106](#)
  - introduction to managing storage efficiency for [104](#)
- SnapVault updates
  - no space on destination aggregate failure [83](#)
- SnapVault-mirror cascades
  - managing when SnapVault backup is unavailable [101](#)

- SnapVault-SnapMirror cascades
  - how they work [39](#)
- source volumes
  - components of a mirror relationship [23](#)
  - managing SnapVault backup of copied [98](#)
- storage efficiency
  - enabling on SnapVault secondary volume [106](#)
  - guidelines for managing, for SnapVault backups [105](#)
  - how SnapVault backups work with data compression [32](#)
  - introduction to managing, for SnapVault secondary volumes [104](#)
- strategies
  - planning data protection [12](#)
- SVM administrators
  - mirror and SnapVault relationship management capabilities [52](#)
- SVMs
  - default Snapshot policies associated with [43](#)
  - guidelines for creating mirror or SnapVault relationships between [25](#)
  - how SnapMirror policies work with [52](#)
  - mirror language setting requirement [25](#)
  - namespace, protecting [32](#)
  - root information, protecting [32](#)

## T

- tape seeding [63](#)
- tiered backup policies
  - example of creating [54](#)
- tools
  - for data-loss protection [9](#)
- transfers
  - scheduling SnapMirror [71](#)

## V

- volume is restricted errors

- correcting [58](#)
- volume snapshot commands
  - for managing Snapshot copies [42](#)
- volume snapshot policy commands
  - for managing Snapshot copy policies and schedules [44](#)
- volumes
  - components of a mirror relationship [23](#)
  - converting data protection destination to SnapVault secondary [92](#)
  - creating SnapVault backup in prepopulated FlexVol [90](#)
  - creating SnapVault backups in empty FlexVol [84](#)
  - enabling storage efficiency on SnapVault secondary [106](#)
  - firewall requirements for intercluster SnapMirror relationships [36](#)
  - guidelines for creating SnapVault relationships on FlexVol [81](#)
  - how Snapshot policies are associated with [43](#)
  - manually updating data protection mirror copies on [73](#)
  - maximum number of Snapshot copies for [14](#)
  - restoring from SnapVault backup [100](#)
  - SnapVault backup limitations for FlexVol [32](#)
  - which data gets backed up and restored from FlexVol [29](#)

Vservers

*See* SVMs

## W

- wildcards
  - path names using [24](#)