



NETAPP TECHNICAL REPORT

Using SnapMirror[®] with SnapDrive[®] for UNIX[®]

Anil Degwekar, NetApp
TR-3611

Abstract

This technical note explains how to use SnapDrive for UNIX along with SnapMirror. This information is intended for users of NetApp[™] storage systems that utilize these products.

SnapDrive for UNIX does not include dedicated commands to control SnapMirror operations. But users can operate these two products together by writing scripts or by following a few manual steps. These steps are explained in this technical note. This technical note deals with volume SnapMirror only.

You should be familiar with the workings of Data ONTAP[®], SnapDrive for UNIX, and SnapMirror before reading this document.

Table of Contents

1	Introduction	3
1.1	About SnapDrive for UNIX.....	3
1.2	About the Data Protection Features of Data ONTAP.....	3
1.3	Interoperability between SDU and SnapMirror	3
2	SDU usage with Volume SnapMirror.....	3
2.1	Example Disaster Recovery deployment scenario—SDU with Volume SnapMirror	3
2.2	Software installation and setup	4
2.3	Provisioning considerations	6
2.4	Making Snapshot™ copies of data	6
2.5	Restoring data from primary Snapshot copies.....	6
2.6	Accessing data from the secondary storage system.....	6
2.7	Using SnapMirror for test and development.....	8
2.8	Using SnapMirror with SDU in an NFS setup.....	9
2.9	Situations where a host disk group spans across multiple storage systems	9
3	Summary	11

1 Introduction

1.1 About SnapDrive for UNIX

SnapDrive for UNIX (SDU) is a tool for storage provisioning and Snapshot management on UNIX systems connecting to NetApp FAS series storage systems. It allows a UNIX system administrator easy access to storage provisioning functionality such as creating LUNs, connecting them to the UNIX system, creating a volume group and filesystem on these LUNs, and expanding the volume groups by using simple commands on the UNIX console. It also allows for the creation of copies of user data which are host filesystem-consistent, and for the restoration of data from these Snapshot copies. Other NetApp software products, such as SnapManager® for Oracle® and SnapManager for SAP®, use the underlying functionality provided by SnapDrive for UNIX to allow Snapshot management which is integrated into applications that use this storage.

1.2 About the Data Protection Features of Data ONTAP

Data ONTAP has several unique features for Data Protection, including SnapMirror and SnapVault®. SnapMirror allows you to schedule regular, automatic replication of filesystem Snapshot copies of a volume or qtree onto another volume or qtree (usually on a different storage system). There are two variations of SnapMirror—Volume SnapMirror and Qtree SnapMirror. SnapVault protects the data in one or more qtrees in a series of Snapshot copies stored on a separate storage system. SnapVault maintains an online, asynchronous, permanently read-only replica of the qtree data.

This technical report mainly deals with Volume SnapMirror.

1.3 Interoperability between SDU and SnapMirror

It is a natural need of users of storage to use both of these products together—i.e., firstly to use SDU for creating host-consistent Snapshot copies, and secondly to use SnapMirror for protecting the Snapshot copies created by SDU.

Although there is no direct integration between these two product groups (i.e., SDU does not provide any specific commands or CLIs to control the way SnapMirror behaves), customers can use these products together. This technical report explains the various ways in which customers can use SDU with SnapMirror. It gives sample deployment scenarios, in addition to the steps users need to take to ensure correct usage of these products. This information should be used in conjunction with the various User Guides published for each of these products. This technical report does not reproduce all of the information in the product User Guides.

2 SDU usage with Volume SnapMirror

2.1 Example Disaster Recovery deployment scenario—SDU with Volume SnapMirror

In the example configuration shown below, a UNIX host 'alpha' is connected to a primary NetApp® storage system 'toaster', and the data on the primary storage system 'toaster' is mirrored by a secondary storage system 'oven' for disaster recovery. The secondary storage system may be at a nearby location or at a far off location. The UNIX host uses LUNs created on the primary storage system. The UNIX host 'alpha' is connected to 'toaster' via an FCP fabric. The connection between 'alpha' and 'oven' is over IP WAN. In normal usage, the UNIX system 'alpha' does not use any storage from 'oven'. However, it has connectivity to 'oven' over IP SAN

(iSCSI) and that connectivity needs to be tested before the deployment of this configuration. This connectivity will be used during disaster recovery.¹

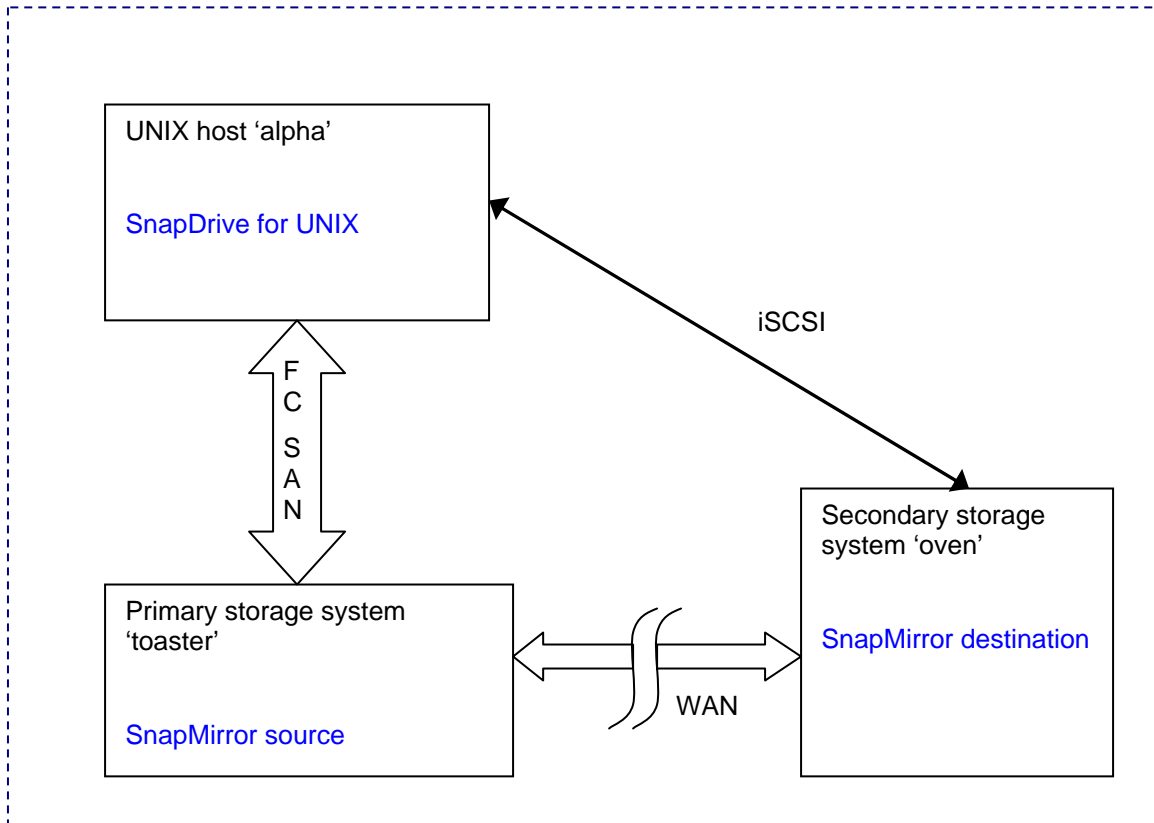


Figure 1) SDU with SnapMirror Configuration 1

This configuration can be further enhanced by using host clustering technology on the host side, which handles failures of the host node.

2.2 Software installation and setup

Installation and setup of this configuration can be done in the normal way. Each storage system is configured with the required licenses (including licenses required for SDU, and those required for SnapMirror). There is no specific order required for installation. Test the connectivity between the systems as follows:

2.2.1 Connectivity between the UNIX host and the primary storage system

Verify that you can create LUNs on the primary storage system using SDU on the UNIX host. To do this, follow the steps described in the SnapDrive for UNIX Installation and Administration Guide. You will have to select “fcp” as the protocol of choice in the SnapDrive Configuration file (snapdrive.conf) on ‘alpha.’² Perform the following steps to verify that all aspects of the connection work properly:

On the storage system:

¹ Note that the IP SAN connectivity may not give the same performance as the FCP SAN connectivity with the primary storage; however, it allows the host ‘alpha’ to continue to operate and access its data until the primary is brought up again. Also note that there may be platform-specific restrictions which do not allow simultaneous use of FCP and iSCSI connectivity on the same host.

² It is also possible to use iSCSI connectivity between ‘alpha’ and ‘toaster.’

- i. Create a volume 'vol1' on 'toaster'

On the UNIX host:

- ii. Create a test LUN and mount it to the host by using this command:
`snapdrive storage create -fs /mnttest -lun toaster:/vol/voll/testlun -lunsize 1g`
- iii. Copy some data to the LUN using a command such as:
`cp -r /usr/bin/* /mnttest`
- iv. Create a Snapshot copy of this data using a command such as:
`snapdrive snap create -fs /mnttest -snapname trial`
- v. Delete the data that you just copied: `rm -r /mnttest/*`
- vi. Restore the data using a command such as:
`snapdrive snap restore -fs /mnttest -snapname trial`
- vii. Verify that the data was properly restored using a command such as:
`ls -l /mnttest`

Once connectivity has been established, you can delete this Snapshot copy using a command such as:
`snapdrive snap delete -snapname toaster:/vol/voll:trial`. You can then delete the test LUN using a command such as: `snapdrive storage delete -fs /mnttest`.

2.2.2 Connectivity between SDU and the secondary storage system

Verify that you can create LUNs on the secondary storage system using SDU on the UNIX host. To do this, follow the steps described in the SnapDrive for UNIX Installation and Administration Guide. You will have to select "iscsi" as the protocol of choice in the SnapDrive Configuration file (`snapdrive.conf`) on 'alpha.'³ Perform the same steps that you followed in the above example, substituting the storage system name 'oven' in place of 'toaster.' Note that even though the connectivity is different here (iSCSI instead of FCP), the command syntax remains the same. At the end of the test, change the preferred connectivity in the 'snapdrive.conf' file back to FCP.

2.2.3 Connectivity between the two storage systems

Set up a SnapMirror relationship between the two storage systems. You can use any of the following methods to do this:

- Data ONTAP CLI commands
- FilerView®
- Operations Manager

Important SnapMirror configuration parameters:

- The secondary volume should be of a size equal to or greater than the primary system volume.
- The versions of Data ONTAP used between the two systems should be compatible for using SnapMirror.
- Set up the mirroring schedule based on your availability requirements. Note that you may lose data up to one mirroring interval in the event of a disaster at the primary site. The schedule also depends on the bandwidth available on the WAN link, and the frequency of data change that you expect on the storage system.
- SnapMirror can operate either synchronously or asynchronously. This technical report will focus on asynchronous mirroring.
- You can use Volume SnapMirror, Qtree SnapMirror, or Synchronous SnapMirror. However, in case of WAN connectivity between 'toaster' and 'oven,' Synchronous SnapMirror will overload the network and so is not recommended.

³ In our sample configuration, the connectivity between 'alpha' and 'oven' is iSCSI, but in theory an FCP connectivity will also work in the same way. Most installations will use iSCSI if the secondary system is at a remote location.

2.3 Provisioning considerations

Once the setup has been configured as discussed above, further usage of SDU on the UNIX host is straightforward. While provisioning, use the volume which is configured for mirroring. SnapMirror will ensure that the data is replicated to the secondary storage system volume at regular intervals. The initial (baseline) transfer might take more time, depending upon how much data is already present in the volume⁴. After that, SnapMirror works independent of the activity on the host.

Example command for storage provisioning:

```
snapdrive storage create -fs /mnt/data -hostvol SduDg/SduHv -lun  
toaster:/vol/vol1/lun1 -lunsize 100g
```

2.4 Making Snapshot copies of data

Using commands like “snapdrive snap create,” you can create Snapshot copies of the data. These Snapshot copies will be host filesystem-consistent, because SDU ensures that host filesystem and volume manager buffers are flushed to disk before making a Snapshot copy.

Example command for making Snapshot copies:

```
snapdrive snap create -fs /mnt/data -snapname 21march11pm
```

2.5 Restoring data from primary Snapshot copies

This works the same regardless of whether SnapMirror is used or not. The primary storage system allows all storage features like single-file SnapRestore®, LUN cloning, LUN resize and so on.

Example command for restoring data from Snapshot copies:

```
snapdrive snap restore -fs /mnt/data -snapname 21march11pm
```

2.6 Accessing data from the secondary storage system

In case the primary storage system is not available, the UNIX host can access data from the secondary storage system. This is not considered a normal “restore,” but it is essentially a “connection” to the secondary system’s copy of the data. There is no direct command in SDU to accomplish this, and so you need to take steps to accomplish this access. These steps are explained below.

2.6.1 Special case: secondary storage configured as a cluster failover partner

In case the secondary storage system (‘oven’ in this example) is configured as a clustered failover partner of the primary storage system (‘toaster’ in this example), then a fast and automatic mechanism of Snapshot restore is supported by SDU. It is essential that the volume name of the source and destination of SnapMirror be identical in such a configuration.

Step 1: Make the secondary volume writable

In case SnapMirror detects a failure of the primary storage system, this will happen automatically. Otherwise, the user needs to do this by using storage system management software such as FilerView.

Step 2: List the Snapshot copies on the storage system using SDU

Use an SDU command such as:

⁴ To speed up the baseline transfer, the secondary system is installed next to the primary system. After the baseline transfer, it is shipped to the backup site.

```
snapdrive snap list -filervol oven:/vol/vol1
```

The Snapshot copy of interest ('21march11pm' in this case) should be listed in the output list of Snapshot copies.

Step 3: Connect to the Snapshot copy

Use an SDU command such as:

```
snapdrive snap connect -fs /mnt/data -snapname  
oven:/vol/vol1:21march11pm -autorename
```

This should create a new filesystem at the new mount point ('/mnt/remote') with the contents of the Snapshot copy taken from the original data. SDU will pick up the correct names of the disk group and host volume from the Snapshot copy automatically.

2.6.2 General case: secondary storage is not a clustered failover partner

In this configuration, the previous approach (using the "snapdrive snap connect" command) will not work. This is because the snap connect command works only on the original storage system ('toaster') or on its clustered failover partner.

However, there is a 4-step procedure which users can follow to achieve the same objective. This procedure is explained below.

Step 1: List the Snapshot copies on the storage system using SDU

Use an SDU command such as:

```
snapdrive snap list -filervol oven:/vol/mirrorvol
```

The Snapshot copy of interest ('21march11pm' in this case) should be listed in the output list of Snapshot copies. Notice that the volume name of the SnapMirror destination volume can be selected arbitrarily—it doesn't need to be identical to the SnapMirror source volume.

Step 2: Create a FlexClone® volume based on the Snapshot copy of interest

Use FilerView or the storage system CLI to create a FlexClone volume from the SnapMirror destination volume. This FlexClone volume should be based on the Snapshot copy of interest ('21march11pm' in this case).

In FilerView, you can use the "create FlexClone" wizard to accomplish this. The CLI to create a FlexClone on the storage system console is:

```
vol clone create vol21m -s volume -b mirrorvol 21march11pm
```

Here 'vol21m' is the new FlexVol® clone which is based on the volume 'mirrorvol' and its Snapshot copy '21march11pm.'

If the UNIX host is configured for remote shell access to the storage system, then this command can be executed from the host using "rsh/ssh" syntax. Otherwise, it needs to be executed from the storage system console.

Step 3: Unmap the LUN connection to the host

The new FlexClone 'vol21m' contains a copy of the LUN 'lun1' which shows a mapping to an Initiator group which was created when the LUN was originally provisioned. It is necessary to remove this stale mapping from the LUN before it can be used.

Use FilerView to remove this mapping. In FilerView, you can use the "LUN Manage" menu to unmap the LUN.

Step 4: Connect to the Snapshot copy

On the UNIX host 'alpha,' use an SDU command such as:

```
snapdrive storage connect -fs /mnt/remote -hostvol SduDg/SduLv -lun  
oven:/vol/vol21m/lun1
```

This should create a new filesystem at the new mount point ('/mnt/remote') with the contents of the Snapshot copy taken from the original data. Note that in this command you need to specify a new mount point ('/mnt/remote' in this case). But the disk group name ('SduDg' in this case) and the host volume name ('SduHv' in this case) need to be identical to the ones you specified when creating the storage. The full LUN path from the FlexClone volume also needs to be provided.

In case the original data was created with the 'nolvm' option (to create the filesystem directly on a raw LUN), then the same option needs to be provided at the time of storage connect.

2.7 Using SnapMirror for test and development

A variation of the configuration shown in Figure 1 is when two UNIX hosts are used— one at a primary site, and another at a remote site. The remote site may use the same data which is created at the primary site, for test and development purposes. This configuration can also be useful when it is required to present a read-only copy of the data at a remote location. Figure 2 shows this configuration.

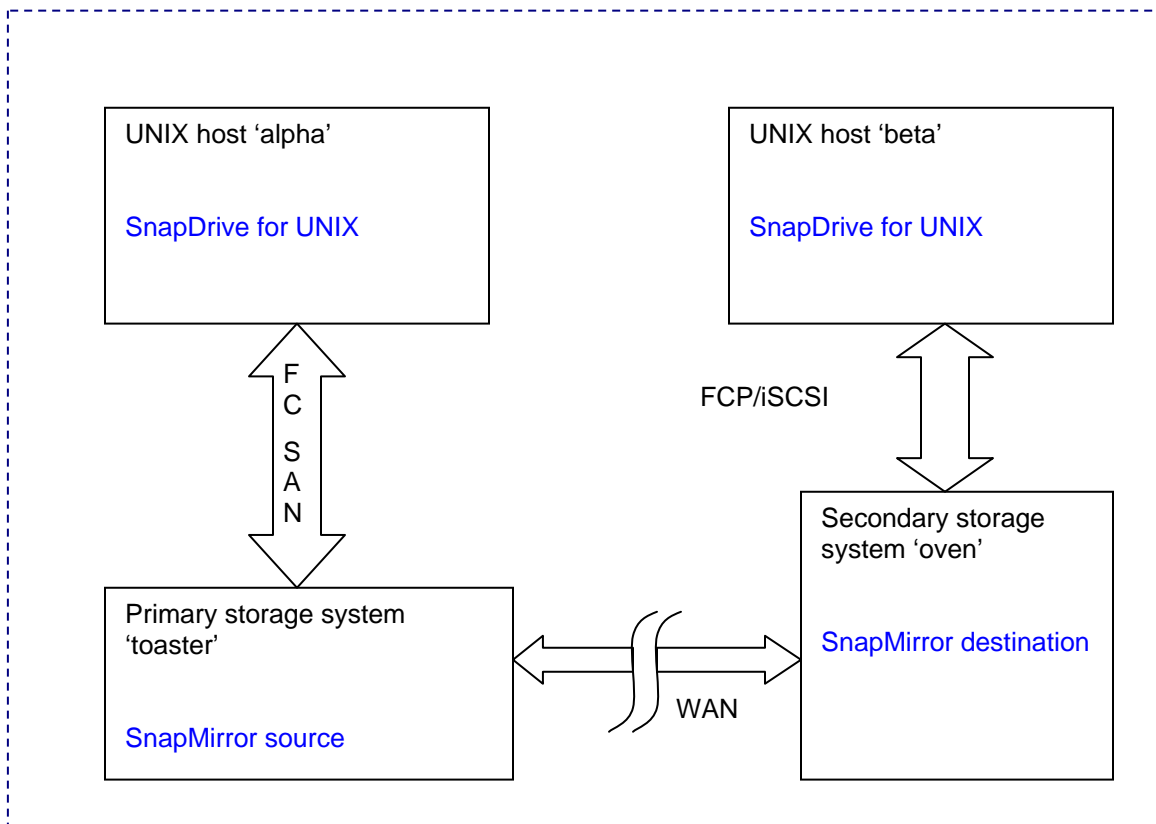


Figure 2) SDU with SnapMirror Configuration 2

In this configuration, a second UNIX host ('beta') is used at the remote site.

Most of the description from the previous section can be applied in this configuration. A few notable changes are as follows:

- While checking the connectivity, check connectivity from the primary host to the primary storage system; and from the secondary host to the secondary storage system. Connectivity from the primary host to the secondary storage system is not needed.
- Connectivity from the primary storage system to the secondary storage system is required.

- There is no change needed in the mechanism for setting up the SnapMirror relationship, provisioning, creating Snapshot copies, and restoring Snapshot copies from the primary host.
- Steps explained in section 2.6.1 are not applicable here, since this procedure is not meant for accessing data in a disaster scenario.
- The SDU commands explained in section 2.6.2 should be executed on the secondary host ('beta' in this case) when accessing data from the secondary storage system.
- It is important not to delete the Snapshot copy of interest from the primary storage system while the secondary has a FlexClone dependent on it. Doing so will interfere with SnapMirror operations because the Snapshot copy is locked on the secondary system.

2.8 Using SnapMirror with SDU in an NFS setup

SDU supports configurations where the host is connected to the storage system on NFS. In such configurations, SnapMirror can be used for disaster recovery or for giving a secondary host access to data for test and development purposes. The steps involved are similar to the previous section.

- Provisioning needs to happen on the storage system. Use commands like 'vol create' to create a flexible volume, and enable NFS export permissions for the host (alpha). Then issue a command like:
`mount toaster:/vol/vol1 /mnt/data`
- Example command for making Snapshot copies remains the same:
`snapdrive snap create -fs /mnt/data -snapname 21march11pm`
- To access storage from a cluster storage partner, you can mount it on the same host (alpha) or on another host (beta), in read-only mode:
`mount -r oven:/vol/vol1 /mnt/mirroredata`
- To create a clone volume based on the Snapshot copy, use this command on the storage system (same as shown in the previous section):
`vol clone create vol21m -s volume -b mirrorvol 21march11pm`
- To access this cloned volume storage for read-write, use:
`mount oven:/vol/vol21m /mnt/mirroredata`
- Note: in order to allow root access, you may need to configure the export permissions on the storage system accordingly.

2.9 Situations where a host disk group spans across multiple storage systems

SDU supports disk groups containing multiple LUNs—these LUNs can be spread over multiple storage systems. Users can use such a configuration to obtain a high bandwidth to the storage subsystem. These disk groups are typically hand crafted (i.e., not created using SDU commands). So, the steps required to create them can be summarized briefly as follows:

1. User creates LUNs and maps them to the host using SDU commands (`snapdrive storage create -lun ...`).
2. User then creates a disk group from these LUNs using LVM commands.
3. User creates logical volumes on the disk groups, creates a filesystem on each of them, and mounts them on the UNIX machine.

If the user wants to use SnapMirror on such configurations, it is possible—provided the following conditions are met:

- Each storage system volume used on the primary storage must be a SnapMirror source volume.
- Use the Consistency Groups feature of Data ONTAP and SDU to make a consistent Snapshot copy across these volumes.

- It is not possible to use the special case of storage systems configured as a clustered failover partner (section 2.6.1).
- Note the names of disk groups and logical volumes carefully—the same should be used subsequently.
- For each SnapMirror destination volume, create a FlexClone volume based on the Snapshot copy of interest.
- Connect the corresponding LUNs to the UNIX machine (alpha or beta) using `snapdrive storage connect -lun` commands.
- Note down the device names of these LUNs (as they get mapped to the host), and then use these to create disk groups.
- Mount the filesystem to the desired mount point on the host.
- The disk group configuration (number of LUNs, order of the LUNs, name of the disk group, name of the logical volume) needs to be identical to the original disk group.
- Some volume managers (e.g., Veritas® VxVM, Linux® LVM) do not allow an identical disk group to be connected to the system while the original disk group is still active. So this procedure can be used only when the original storage is lost, or on a different UNIX host.

Note that this section does not provide example commands. You can create your own commands based on the information provided in the previous sections. Note that this can be a difficult task because of the complexities associated with using volume manager commands in the correct manner.

3 Summary

SnapDrive for UNIX (SDU) is a tool for storage provisioning and Snapshot management on UNIX systems connecting to NetApp FAS series storage systems. Data ONTAP has several unique features for Data Protection, including SnapMirror and SnapVault. This technical note details methods of using these tools together to achieve Data Protection and Disaster Recovery.

