



Connecting to SAN Snapshot™ Copies on HP-UX™

Description of Connecting Network Appliance Snapshot™ Copies to an HP-UX System in a SAN Environment

by Richard Jooss, Network Appliance, Inc.

May, 2004 | TR 3318

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Table of Contents

- 1. Scope**
- 2. Intended Audience**
- 3. Requirements and Assumptions**
- 4. Mapping from File System to NetApp Filer Disks**
- 5. Overview of the Process**
- 6. Steps to Connect to LUNs Created with Snapshot Copies**
- 7. Example of Connecting to a LUN from a Snapshot Copy**
- 8. Graphical Representation of Example**
- 9. Steps to Return to the Original State**
- 10. Example of Returning to the Original State**
- 11. Automation/Scripting**

Abstract

This document describes the steps necessary to connect LUNs created with NetApp filer Snapshot copy technology to an HP-UX system using HP's Logical Volume Manager (LVM) in a SAN environment.

1. Scope

This document covers the steps needed for mounting NetApp filer Snapshot copies from LUNs from HP-UX systems. In other words, it describes how to gain read/write access to Snapshot copies created for file systems on HP-UX systems. It does not expand on the multitude of reasons for wanting to gain access to these Snapshot copies, for example, restore, backup, consistency checking, data mining, etc.

2. Intended Audience

This paper is intended for system and storage administrators.

3. Requirements and Assumptions

For the methods and procedures in this document to be useful to the reader, several assumptions are made:

- The reader has at least basic HP-UX administration skills and has access to the administrative login for the server.
- The reader has at least basic Network Appliance administration skills and has administrative access to the filer via the command-line interface.
- The filer and host have the necessary licenses to perform the activities outlined in this document.
- The target system has the required block-level and network protocol interconnects to perform the activities outlined in this document.
- Snapshot copies of the file system or, more correctly stated, Snapshot copies of the NetApp filer volumes containing the LUNs for the file system have already been created, and they have been created in such a way that the content of the LUNs is consistent. The recommend method is using NetApp SnapDrive™ for UNIX®.
- The Snapshot copies will be mounted on the same system where the original LUNs are mounted. It is generally considered a simpler case to mount the LUNs to a different (often referred to as nonoriginating) host because the problem of "duplicate volume group IDs" generally does not exist. This procedure is considered a safe approach because the HP-UX LVM volume group (VG) ID is changed before mounting and will work equally well on originating or nonoriginating hosts.

- In the examples in this report, all administrative commands are performed at the server or filer console for clarity. Web-based management tools can also be used.

4. Mapping from File System to NetApp Filer Disks

It is important to understand the logical stack that exists between the file system and the NetApp filer disks. There are many possible combinations, and the diagram below shows the most common configurations. On the host the stack consists of:

- A file system created on a logical volume.
- The logical volume consists of a portion of a single or multiple physical volumes.
- Physical volumes are formatted raw devices or LUNs from a host perspective. There is a one-to-one relationship between physical volumes and LUNs on the NetApp filer.

On the NetApp filer the stack consists of:

- Each LUN is mapped to one and only one NetApp filer volume.
- Each volume is mapped to one or more RAID groups.
- Each RAID group consists of a number of disks.

For the sake of simplicity, the logical diagram shown below does not show the multiple paths to each LUN.

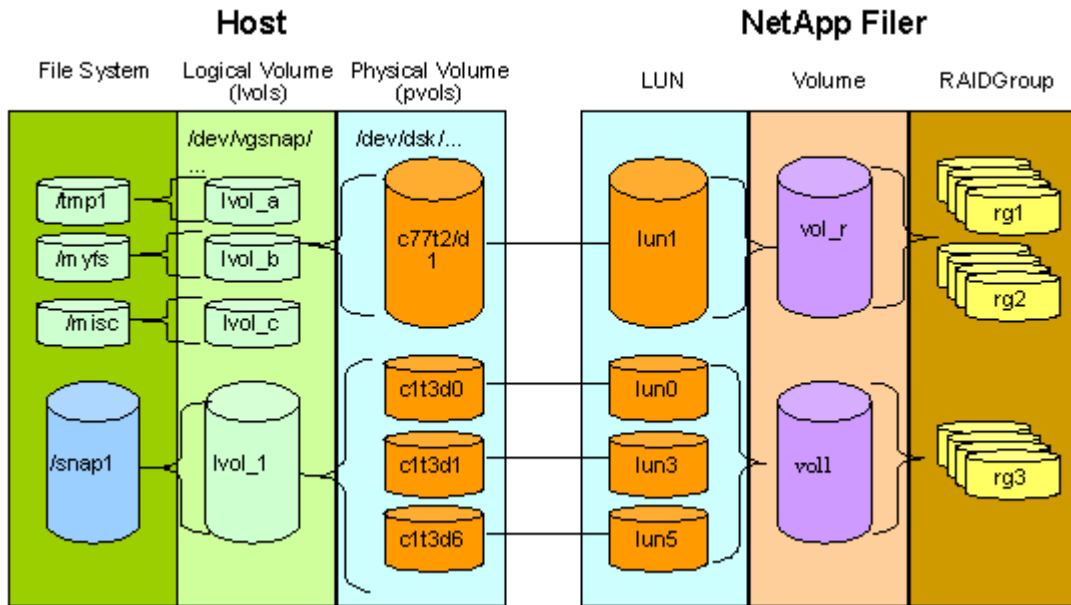


Figure 1. Mapping from File System to NetApp Filer

5. Overview of the Process

Before going into the detailed steps necessary to mount the LUNs from the Snapshot copies, it makes sense to quickly describe the general process. The general principle is that after a Snapshot copy is created (recommended method is using NetApp SnapDrive for UNIX), the LUNs are created on the NetApp filer but not visible to any hosts. Before making the LUNs visible to a host for mounting, a readable/writable copy or Snapshot copy of the Snapshot copy needs to be created, since the LUNs in the Snapshot copy are read only. The only additional space required by this additional copy or Snapshot copy is for the changes that are actually written to the LUN after changes have been made. All operations on this readable/writable version are performed on the same disk mechanisms as the original LUN, so heavy activity can also affect performance. The readable/writable LUNs are then mapped to the NetApp filer interfaces (FC or iSCSI) to make them available to the hosts. After the necessary steps are performed on the host to gain access to the LUNs, the volume ID on the LUNs is modified to avoid having multiple volumes with the same IDs. They are then imported into the volume manager. At that point, it is just necessary to mount the file systems contained on the LUNs to a mount point on the system.

6. Steps to Connect to LUNs Created with Snapshot Copies

The following table describes the steps necessary to connect LUNs created by creating a Snapshot copy of a Network Appliance volume.

Step	Command	Command Location	Description
1	<code>sanlun lun show -p</code>	Host	Shows a list of LUNs visible to the host indicating to which host volume group the LUNs belong and on which filer the LUNs reside. This information can be used to determine on which filer the follow-on steps need to be executed.
2	<code>lun show -m</code>	Filer	Shows a list of LUNs mapped on the filer giving LUN names, filer volume names, and to which host they are mapped. The filer volume names and LUN names will be used in future steps.
3	<code>snap list <filer volume></code>	Filer	Lists the available Snapshot copies for the desired volume. Typically all the LUNs will be contained within one filer volume. If there are multiple filer volumes, this step needs to be repeated for each volume.
4	<code>lun create -b <lun_snapshot_name> <new_lun_name></code>	Filer	Creates a copy or Snapshot copy of the original LUN from the Snapshot copy. Originally the LUN from the Snapshot copy is read-only, and in order to mount/activate the host volume group (VG), a writ able Snapshot copy is needed. The original "snapshotted" LUN is not accessible during the time this LUN created with "lun create -b" exists. The original Snapshot copy state is returned after this LUN is removed. This command has to be done for each needed LUN.
5	<code>lun map <new_lun_name> <igroup> <LUN_ID></code>	Filer	Maps the newly created LUNs so that they are visible for the host. Typically, the same igroup will be used as for the original LUNs, but if for some reason a more restricted group is desired, a new group could be created.
6	<code>ioscan -fnC disk</code>		Scans all I/O devices on the host for newly created disk devices (LUNs in the case of FC/iSCSI devices).
7	<code>insf -eC disk</code>	Host	Creates device files (i.e., /dev/dsk/cXtYdZ) for the newly created and mapped LUNs.
8	<code>sanlun lun show -p</code>	Host	Shows a list of LUNs visible to the host. The newly created and mapped LUNs should appear with no host volume group (VG) assigned. The "-p" option can be left off to get different presentation of the device files.
9	<code>vgchgid /dev/rdisk/<cXtYdZ></code>	Host	The vgchgid (volume group change ID) command changes the host volume group

	<code>/dev/rdisk/<cXtYdQ></code>		(VG) ID on the LUNs to avoid the machine having two VGs with the same ID. All the LUNs making up the volume group need to be included in a single invocation of the command.
10	<code>ls -s /dev/*/group</code>	Host	This command gives a list of host volume groups showing the group name (<code>/dev//group</code>) and minor number (<code>0xXX0000</code>). The major number of 64 is also displayed. This list is used to pick a unique VG name and minor number for the next steps.
11	<code>mkdir /dev/<vgsnap_directory></code>	Host	This is to make the necessary directory for the new VG group that will be created for accessing the Snapshot copy.
12	<code>mknod /dev/<vgsnap_directory>/group c 64 0xXX0000</code>	Host	This command creates the special character device for a host volume group (VG). The directory and minor number need to be unique.
13	<code>vgimport /dev/<vgsnap> /dev/dsk/<cXtYdZ>/dev/dsk/ ...</code>	Host	This command will import the new Snapshot copy LUNs into the new host volume group that was created. Note that with multiple paths there will be multiple devices files per LUN and all the possible device files pointing to the LUNs should be entered.
14	<code>vgchange -a y /dev/<vgsnap></code>	Host	This command will activate the host VG.
15	<code>mkdir 777 /<snap></code>	Host	This command creates a directory for mounting the file system.
16	<code>mount /dev/<vgsnap>/<logical volume> /<snap></code>	Host	This command mounts the file system on the directory that was just created.

7. Example of Connecting to a LUN from a Snapshot Copy

The following table shows an example where a file system is built using a single LUN located on a single filer. Typically, a logical volume will consist of multiple LUNs, but for simplicity only a single LUN is used here. If multiple LUNs are used, steps 4, 5, 9, and 13 need to be repeated or expanded to include the additional LUNs. If the Snapshot copies are from multiple filers, steps 2, 3, 4, and 5 will need to be executed on each filer.

Step	Example
1	<pre> root> sanlun lun show -p filer1:/vol/voll/sh-hp2-lun2049 (LUN 2049) VG: /dev/vg_ntap01 700m (734003200) lun state: GOOD ----- path path /dev/dsk host target partner PVlinks state type filename HBA port port priority ----- up primary /dev/dsk/c44t0d1 td0 4a 0 up primary /dev/dsk/c45t0d1 td1 5a 1 up secondary /dev/dsk/c51t0d1 td1 5b 2 up secondary /dev/dsk/c50t0d1 td0 4b 3 root> </pre>
2	<pre> filer1> lun show -m LUN path Mapped to LUN ID ----- /vol/voll/sh-hp1-lun-010 sh-hp1 10 /vol/voll/sh-hp2-lun2049 sh-hp2 2049 filer1> </pre>
3	<pre> filer1> snap list voll Volume voll working... %/used %/total date name ----- 0% (0%) 0% (0%) Feb 29 12:52 test_hpux_snap filer1> </pre>
4	<pre> filer1> lun create -b /vol/voll/.snapshot/test_hpux_snap/sh-hp2-lun2049 /vol/voll/snap-sh-hp2-lun2049 filer1> </pre>
5	<pre> filer1> lun map /vol/voll/snap-sh-hp2-lun2049 sh-hp2 13 filer1> </pre>
6	<pre> root> ioscan -fnC disk Class I H/W Path Driver S/W State H/W Type Description ===== disk 0 0/0/1/1.15.0 sdisk CLAIMED DEVICE HP 18.2GMAN3184MC /dev/dsk/ctl15d0 /dev/rdisk/ctl15d0 disk 9 0/4/0/0.1.4.0.0.1.5 sdisk CLAIMED DEVICE NETAPP LUN disk 1 0/4/0/0.1.4.0.16.0.1 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c44t0d1 /dev/rdisk/c44t0d1 disk 15 0/4/0/0.1.7.0.0.1.5 sdisk CLAIMED DEVICE NETAPP LUN disk 7 0/4/0/0.1.7.0.16.0.1 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c50t0d1 /dev/rdisk/c50t0d1 disk 10 0/6/2/0.1.12.0.0.1.5 sdisk CLAIMED DEVICE NETAPP LUN disk 2 0/6/2/0.1.12.0.16.0.1 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c45t0d1 /dev/rdisk/c45t0d1 disk 16 0/6/2/0.1.15.0.0.1.5 sdisk CLAIMED DEVICE NETAPP LUN disk 8 0/6/2/0.1.15.0.16.0.1 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c51t0d1 /dev/rdisk/c51t0d1 root> </pre>
7	<pre> root> insf -Cdisk insf: Installing special files for sdisk instance 9 address 0/4/0/0.1.4.0.0.1.5 </pre>

	<pre> insf: Installing special files for sdisk instance 15 address 0/4/0/0.1.7.0.0.1.5 insf: Installing special files for sdisk instance 10 address 0/6/2/0.1.12.0.0.1.5 insf: Installing special files for sdisk instance 16 address 0/6/2/0.1.15.0.0.1.5 </pre>
8	<pre> root> sanlun lun show -p filer1:/vol/vol1/sh-hp2-lun2049 (LUN 2049) VG: /dev/vg_ntap01 700m (734003200) lun state: GOOD ----- path path /dev/dsk host target partner PVLlinks state type filename HBA port port priority ----- up primary /dev/dsk/c44t0d1 td0 4a 0 up primary /dev/dsk/c45t0d1 td1 5a 1 up secondary /dev/dsk/c51t0d1 td1 5b 2 up secondary /dev/dsk/c50t0d1 td0 4b 3 filer1:/vol/vol1/snap-sh-hp2-lun2049 (LUN 13) VG: none 700m (734003200) lun state: GOOD ----- path path /dev/dsk host target partner PVLlinks state type filename HBA port port priority ----- up secondary /dev/dsk/c55t1d5 td1 5b up primary /dev/dsk/c53t1d5 td1 5a up secondary /dev/dsk/c54t1d5 td0 4b up primary /dev/dsk/c52t1d5 td0 4a root> </pre>
9	<pre> root> vgchgid /dev/rdisk/c52t1d5 /dev/rdisk/c41t1d6 </pre>
10	<pre> root> ls -l /dev/*/group crw-r----- 1 root sys 64 0x000000 Aug 12 2002 /dev/vg00/group crw-r--r-- 1 root sys 64 0x010000 Oct 30 18:04 /dev/vgntap01/group </pre>
11	<pre> root> mkdir /dev/vgsnap </pre>
12	<pre> root> mknod /dev/vgsnap/group c 64 0x440000 </pre>
13	<pre> root> vgimport /dev/vgsnap /dev/dsk/c52t1d5 /dev/dsk/c53t1d5 /dev/dsk/c55t1d5 /dev/dsk/c54t1d5 vgimport: Warning: Volume Group contains "1" PVs, "4" specified. Continuing. Warning: A backup of this volume group may not exist on this machine. Please remember to take a backup using the vgcfgbackup command after activating the volume group. </pre>
14	<pre> root> vgchange -a y /dev/vgsanp Activated volume group Volume group "/dev/vgsnap" has been successfully changed. root> </pre>
15	<pre> root> mkdir 777 /mnt/snap root> </pre>
16	<pre> root> mount /dev/vgsnap/lvol1 /mnt/snap </pre>

8. Graphical Representation of Example

The diagram below depicts the result of the steps described in the preceding table. The first row, Original, shows how the original file system is mounted. The second row, Snapshot Copy of Original, shows how the Snapshot copy exists on the filer, but the LUNs are not visible to the host. The LUNs in this row are read only and not available during the time the LUNs in the third row exist. The third line, lun create -b of original, displays the readable/writable LUNs that were made visible to the host to allow the mounting of the file system. Any changes made to the LUNs in this third row will be lost as soon as the LUNs are removed or destroyed. The diagram below differs from the example in that it shows two LUNs, which make up the file system, and the names not representative.

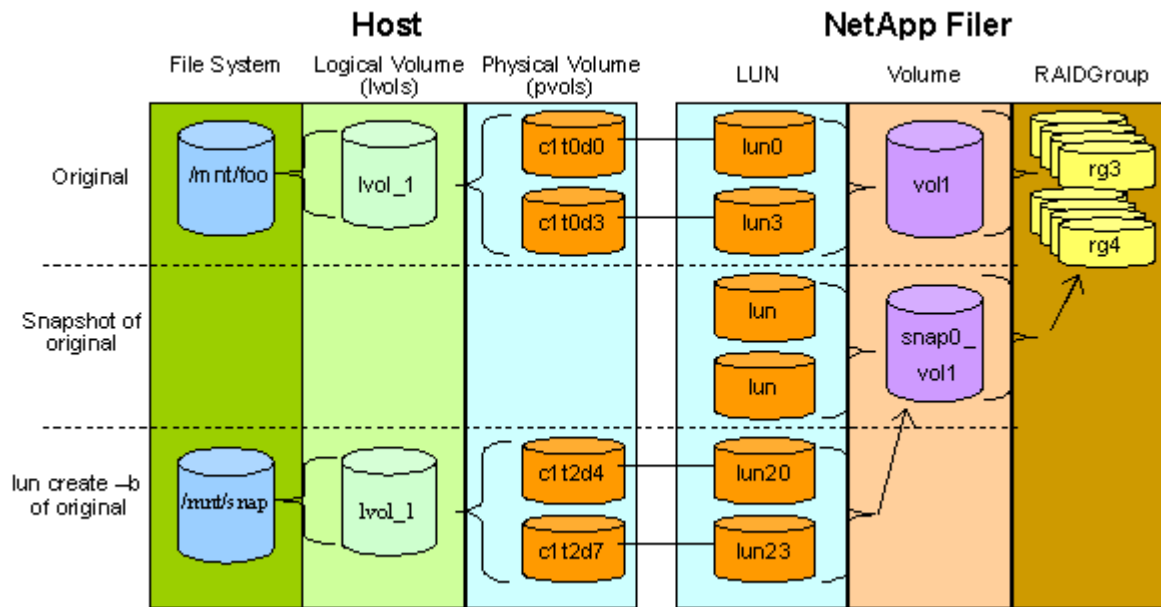


Figure 2. Mapping of Snapshot Copy and Readable/Writable Copy of Snapshot Copy

9. Steps to Return to the Original State

The following table describes the steps necessary to undo what was done to connect to the Snapshot copy and return to the original state. To correlate the steps, these steps are numbered r<number> for reverse step.

Step	Command	Command Location	Description
r16	<code>umount / <snap></code>	Host	This command unmounts the file system.
r15	<code>rmdir /<snap></code>	Host	This command removes the directory where the file system was mounted.
r14	<code>vgchange -a n /dev/<vgsnap></code>	Host	This command will deactivate the host VG.

r13-r10	<code>vgexport /dev/<vgsnap></code>	Host	This command will remove the host VG, including the <code>/dev/<vgsnap></code> directory and group file from the host.
r9	No action necessary		This change, which was written to the disk, will be lost when the LUNs are destroyed in a later step.
r8	No action necessary		This was for information only.
r7-r6	<code>rmsf -k /dev/dsk/<cXtYdZ> /dev/dsk/<cXtYdZ> /dev/dsk/<cXtYdQ> ...</code> <code>rmsf /dev/dsk/<cXtYdZ> /dev/dsk/<cXtYdZ> /dev/dsk/<cXtYdQ> ...</code>	Host	These commands remove the device files from the system kernel (-k) and from the file system. Both commands need to be executed and should be executed with the same list of device files. The list of device files should be the same as what was given in step 13 during the create process.
r5	<code>lun unmap <new_lun_name><igroup></code>	Filer	This command removes the mapping of the LUNs from the igroup.
r4-r1	<code>lun destroy <new_lun_name></code>	Filer	This command removes the created LUN. After this command is executed, the LUN within the Snapshot copy is returned to its original state.

10. Example of Returning to the Original State

The following table shows the steps necessary to return to the original state from the example steps show in the preceding table.

Step	Example
r16	<code>root> umount /mnt/snap</code>
r15	<code>root> rmdir /mnt/snap</code>
r14	<code>root> vgchange -a n /dev/vgsnap</code> Volume group "/dev/vgsnap" has been successfully changed. <code>root></code>
r13-r10	<code>root> vgexport /dev/vgsnaproot></code>
r7-r6	<code>root> rmsf -k /dev/dsk/c52t1d5 /dev/dsk/c53t1d5 /dev/dsk/c55t1d5 /dev/dsk/c54t1d5</code> <code>root></code> <code>root> rmsf /dev/dsk/c52t1d5 /dev/dsk/c53t1d5 /dev/dsk/c55t1d5 /dev/dsk/c54t1d5</code> <code>root></code>
r5	<code>filer1> lun unmap /vol/vol1/snap-sh-hp2-lun2049 sh-hp2filer1></code>
r4-r1	<code>filer1> lun destroy /vol/vol1/snap-sh-hp2-lun2049</code>

11. Automation/Scripting

This is a procedure that typically is done repeatedly and is automated through scripts where the data is then backed up to tape, analyzed, checked for consistency, etc. Typically, the Snapshot copies are created and renamed using the same names each time to make the scripting process easier. Also, by using the same mappings each time, the steps of finding the LUNs and creating the device files can be eliminated.

In a scripted solution, normally only steps 4, 5, 9, 11 through 14, and 16 will need to be done each time to gain access to the Snapshot copies, and steps r16, r14, r13, r5, and r4 are necessary to return to the starting state.