Dell Storage Center with HP UX 11i v3 Best Practices

Daniel Tan, UNIX/Linux Product Specialist
Dell Storage Engineering
December 2015

Dell Best Practices

Revisions

Date	Rev.	Description	Author
July 2014	1	Initial draft	Daniel Tan
May 2015	2	Introduce connectivity to Dell Storage SCv2000	Daniel Tan
December 2015	3	Refresh paper	Daniel Tan

Table of contents

Executive summary

HP-UX from Hewlett Packard Enterprise is a robust and scalable enterprise-class operating system. Correctly configured, using the recommendations presented in this paper, the HP-UX operating system provides an optimized experience for use with the Dell™ Storage SC Series arrays. These recommendations include guidelines for configuring volume discovery, multipath, file system, SCSI configuration and queue depth management.

The scope of this paper discusses features of HP-UX 11i v3 with Storage Center OS (SCOS) version 6.5.x thru 6.7.x. Because there are often various methods for accomplishing these tasks, this paper provides a reference starting point for end users and system administrators.

This guide focuses almost exclusively on the command line interface (CLI) because it is often the most universally applicable across UNIX and Linux distributions.

1 Overview

Enterprise environments often require the minimization or elimination of single points of failure in server installations including all required resources such as power, storage and networking. In this paper, assumption is made that both networking and power are redundant and configured correctly. This paper focuses on the implementation of SC Series arrays in an enterprise environment with HP-UX 11i v3 hosts.

The optimum configuration for HP-UX servers accessing an SC Series array consists of at least one pair of controllers, two physically separate Fibre Channel (FC) fabrics and at least two host bus adapters (HBA) in each HP-UX host.

This paper was written to discuss the configuration shown in Figure 1 and address using the mass storage I/O stack introduced with HP-UX 11i v3.

The diagram below provides an example of a multi-fabric, redundant enterprise environment.

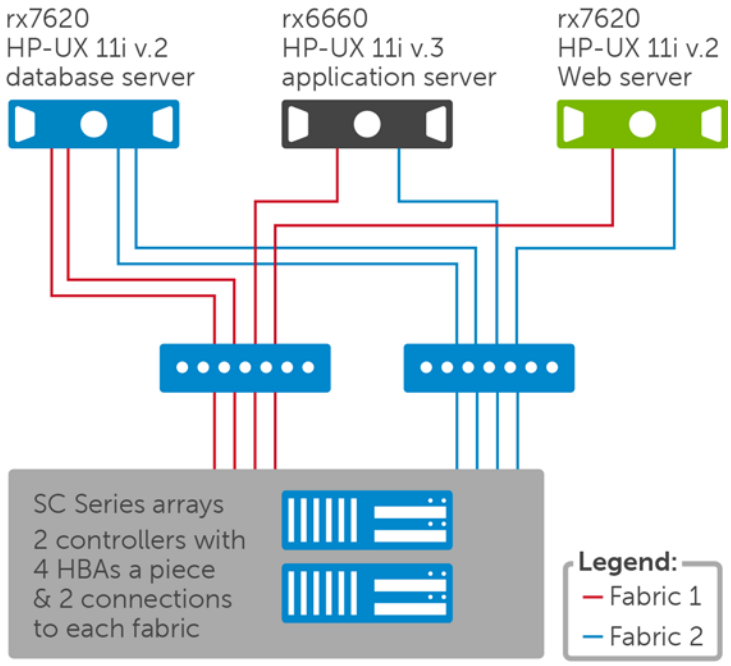

Figure 1 HP-UX best practice configuration

2 Configuration

HP-UX 11i v3 supports both legacy device special files (DSFs) and agile persistent DSFs. SC Series arrays must be configured in legacy port mode when used with HP-UX and legacy DSFs. Either legacy or virtual port mode is acceptable when used with HP-UX and agile persistent DSFs. In mixed-use environments, the SC Series arrays should be configured in Legacy port mode.

When using HP-UX with Legacy DSF, the fibre channel switches connecting the HP-UX hosts to the SC Series array must be capable of WWPN binding and have this feature enabled. Cisco® switches refer to this feature as persistent FCID, while Brocade® switches call it persistent PID. This feature is not required when using HP-UX with only Agile DSF.

2.1 HBA discovery

HBA devices and their respective worldwide port names (WWPN) are discovered using this sample script shown below.

The WWPNs are used to configure the necessary server object(s) onto the SC Series array.

```
# for i in `ls -1 /dev/*d* | grep -e td -e fcd`
> do
> echo "=== $i"
> /opt/fcms/bin/fcmsutil $i | grep Port
> echo
> done
=== /dev/fcd0
                    Local N_Port_id is = 0x270900
                 Previous N_Port_id is = None
            N_Port Node World Wide Name = 0x50014380017af449
            N_Port Port World Wide Name = 0x50014380017af448
            Switch Port World Wide Name = 0x200900053394 08e8
              N_Port Symbolic Port Name = rx3600_fcd0
              N_Port Symbolic Node Name = rx3600_HP-UX_B.11.31

=== /dev/fcd1
                    Local N_Port_id is = 0x280900
                 Previous N_Port_id is = None
            N_Port Node World Wide Name = 0x50014380017af44b
            N_Port Port World Wide Name = 0x50014380017af44a
            Switch Port World Wide Name = 0x20090005339ed94b
              N_Port Symbolic Port Name = rx3600_fcd1
              N_Port Symbolic Node Name = rx3600_HP-UX_B.11.31

=== /dev/fcd2
                    Local N_Port_id is = None
                 Previous N_Port_id is = None
            N_Port Node World Wide Name = 0x500110a0003db439
```

```
N_Port Port World Wide Name = 0x500110a0003db438
Switch Port World Wide Name = 0000000000000000
  N_Port Symbolic Port Name = rx3600_fcd2
  N_Port Symbolic Node Name = rx3600_HP-UX_B.11.31
```

2.2 Volume discovery

A 100G volume is created on the SC Series array and mapped to the HP-UX host. Scanning and discovering new volumes on HP-UX is accomplished with the following commands.

```
# insf –e
# ioscan –fn
```

After running the above commands, all SCSI-based disk devices are displayed with the following command.

```
# ioscan -fnC disk
Class     I  H/W Path     Driver S/W State   H/W Type     Description
==========================================================================
disk      6  0/2/1/0/4/0.41.8.2.0.0.1  sdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                         /dev/dsk/c15t0d1   /dev/rdsk/c15t0d1
disk      7  0/2/1/0/4/0.41.9.2.0.0.1  sdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                         /dev/dsk/c17t0d1   /dev/rdsk/c17t0d1
disk      9  0/2/1/0/4/1.42.8.2.0.0.1  sdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                         /dev/dsk/c31t0d1   /dev/rdsk/c31t0d1
disk     10  0/2/1/0/4/1.42.9.2.0.0.1  sdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                         /dev/dsk/c33t0d1   /dev/rdsk/c33t0d1
disk      0  0/4/1/0.0.0.0.0           sdisk   CLAIMED     DEVICE       HP
DG146A4960
                         /dev/dsk/c0t0d0     /dev/dsk/c0t0d0s1
/dev/dsk/c0t0d0s2   /dev/dsk/c0t0d0s3   /dev/rdsk/c0t0d0    /dev/rdsk/c0t0d0s1
/dev/rdsk/c0t0d0s2  /dev/rdsk/c0t0d0s3
disk      1  0/4/1/0.0.0.1.0           sdisk   CLAIMED     DEVICE       HP
DG146BAAJB
                         /dev/dsk/c0t1d0     /dev/dsk/c0t1d0s1
/dev/dsk/c0t1d0s2   /dev/dsk/c0t1d0s3   /dev/rdsk/c0t1d0    /dev/rdsk/c0t1d0s1
/dev/rdsk/c0t1d0s2  /dev/rdsk/c0t1d0s3
disk      4  255/1/0.0.0  sdisk   CLAIMED     DEVICE       TEAC    DV-28E-V
                         /dev/dsk/c1t0d0   /dev/rdsk/c1t0d0
```

To display volume device correlation, and mapping for Agile DSFs to Legacy DSFs.

```
# ioscan -m dsf
Persistent DSF          Legacy DSF(s)
```

```
========================================
/dev/rdisk/disk2          /dev/rdsk/c0t0d0
/dev/rdisk/disk2_p1       /dev/rdsk/c0t0d0s1
/dev/rdisk/disk2_p2       /dev/rdsk/c0t0d0s2
/dev/rdisk/disk2_p3       /dev/rdsk/c0t0d0s3
/dev/rdisk/disk3          /dev/rdsk/c0t1d0
/dev/rdisk/disk3_p1       /dev/rdsk/c0t1d0s1
/dev/rdisk/disk3_p2       /dev/rdsk/c0t1d0s2
/dev/rdisk/disk3_p3       /dev/rdsk/c0t1d0s3
/dev/rdisk/disk5          /dev/rdsk/c1t0d0
/dev/rdisk/disk8          /dev/rdsk/c15t0d1
                          /dev/rdsk/c17t0d1
                          /dev/rdsk/c31t0d1
                          /dev/rdsk/c33t0d1
```

The multipath details of a specific volume device are displayed with the command as shown.

```
# ioscan -m lun /dev/disk/disk8
Class     I  Lun H/W Path  Driver  S/W State   H/W Type     Health  Description
======================================================================
disk      8  64000/0xfa00/0xb   esdisk  CLAIMED     DEVICE       online
COMPELNTCompellent Vol
             0/2/1/0/4/0.0x5000d31000006505.0x4001000000000000
             0/2/1/0/4/0.0x5000d31000006506.0x4001000000000000
             0/2/1/0/4/1.0x5000d31000006507.0x4001000000000000
             0/2/1/0/4/1.0x5000d31000006508.0x4001000000000000
                      /dev/disk/disk8   /dev/rdisk/disk8
```

HP-UX SCSI-based volume devices are correlated to SC Series volumes using the sample script shown below. This script retrieves the serial number of the volume as represented on the SC Series array.

```
# for i in `ioscan -m dsf | grep disk | awk '{print $1}' | grep -v _p | grep -v
pt`
do
        echo "=== ${i}"
        scsimgr -p get_attr -D ${i} -a serial_number
        echo
done
[snip]
=== /dev/rdisk/disk8
"00000065-00001f86"
```

The new disk devices are prepared for LVM use via the following command.

```
# pvcreate /dev/rdisk/disk8
Physical volume "/dev/rdisk/disk8" has been successfully created.
```

2.3 Volume use

This section demonstrates basic use of the /dev/rdisk/disk8 volume device, creating a volume group, logical volume and mounting this logical volume into the file system.

HP-UX 11i v3 introduces LVM v2 volume groups. All LVM commands used in this paper are specific to LVM v2 command syntax. The command syntax with LVM v2 has been simplified and a single command can be used to create a volume group, its supporting directory structure and group files are shown below. This method should only be used when creating simple volume group structures (for example: not clustered or shared volumes).

```
# vgcreate -V 2.1 -s 64 -S 1p /dev/vgdemo /dev/disk/disk8
Volume group "/dev/vgdemo" has been successfully created.
Volume Group configuration for /dev/vgdemo has been saved in
/etc/lvmconf/vgdemo.conf
```

Where -V represents LVM 2.1, -s represents 64MB PE size and -S is 1PB as the total size limit of the volume group.

The creation of the volume group in this manner respects the native multipath mass storage I/O stack of HP-UX 11i v3 and adds the Agile DSF to the volume group in an active/active path configuration.

The newly created volume group is shown below.

```
# vgdisplay -v vgdemo
--- Volume groups ---
VG Name                     /dev/vgdemo
VG Write Access             read/write
VG Status                   available
Max LV                      2047
Cur LV                      0
Open LV                     0
Max PV                      2048
Cur PV                      1
Act PV                      1
Max PE per PV               262144
VGDA                        2
PE Size (Mbytes)            64
Total PE                    1595
Alloc PE                    0
Free PE                     1595
Total PVG                   0
Total Spare PVs             0
Total Spare PVs in use      0
VG Version                  2.1
VG Max Size                 1p
VG Max Extents              16777216
```

```
   --- Physical volumes ---
   PV Name                     /dev/disk/disk8
   PV Status                   available
   Total PE                    1595
   Free PE                     1595
   Autoswitch                  On
   Proactive Polling           On
```

A 50G logical volume is created and mounted to the directory /vgdemo.

```
# mkdir -p /vgdemo; lvcreate -L 50000 /dev/vgdemo; newfs -F vxfs -o largefiles /
-b 8192 /dev/vgdemo/rlvol1; mount /dev/vgdemo/lvol1 /vgdemo; bdf
Warning: rounding up logical volume size to extent boundary at size "50048" MB.
Logical volume "/dev/vgdemo/lvol1" has been successfully created with
character device "/dev/vgdemo/rlvol1".
Logical volume "/dev/vgdemo/lvol1" has been successfully extended.
Volume Group configuration for /dev/vgdemo has been saved in
/etc/lvmconf/vgdemo.conf
    version 7 layout
    51249152 sectors, 51249152 blocks of size 1024, log size 65536 blocks
    largefiles supported
Filesystem          kbytes    used   avail %used Mounted on
/dev/vg00/lvol3    4292608  230376 4030560    5% /
/dev/vg00/lvol1    2097152  197104 1885288    9% /stand
/dev/vg00/lvol8    8388608 1606832 6737048   19% /var
/dev/vg00/lvol7    8388608 3070056 5277088   37% /usr
/dev/vg00/lvol6    8388608   21032 8302208    0% /tmp
/dev/vg00/lvol5    16777216 4769064 11914368   29% /opt
/dev/vg00/lvol4    8388608   21088 8302152    0% /home
/dev/vgdemo/lvol1  51249152   79205 47971833    0% /vgdemo
```

3 Dynamic LUN Expansion (DLE) and Contraction (DLC)

This section discusses the `vgmodify` command and its use with dynamic volume expansion and contraction tasks.

The use of `vgmodify` requires a minimum of Mar 2008 update release and is only applied towards LVM v2 volume groups.

Identify the current installed and release state with the following command, where `0909` represents the Sep 2009 update.

```
# swlist -l bundle | grep HPUX11i
HPUX11i-DC-OE          B.11.31.0909   HP-UX Data Center Operating Environment
```

3.1 Expanding the LUN

Disk geometry of a LUN is revealed using the following code.

```
# /etc/vx/diag.d/vxscsi -g /dev/rdisk/disk8
geometry of /dev/rdisk/disk8: nhead=8 nsect=525 ncyl=49932 sectsz=512 rpm=15000
cap=209715200
```

The volume on SC Series array is expanded from 100G to 150G. Run the `vxscsi` command again to reveal that the operating system has also detected the disk geometry change.

```
# /etc/vx/diag.d/vxscsi -g /dev/rdisk/disk8
geometry of /dev/rdisk/disk8: nhead=32 nsect=154 ncyl=63833 sectsz=512 rpm=15000
cap=314572800
```

Note: The vxscsi tools are available on HP-UX 11i v3 with the installation of the Base-VxTools-50, B.05.00.01 file set or newer.

3.2 Exposing the new LUN capacity

Though the operating system detected the geometry changes of the LUN, the volume group where it is a member is not yet aware of this change. The number of available physical extents (PE) of the LUN within the volume group is refreshed with the following command.

```
# vgmodify -a -C 100g /dev/vgdemo /dev/rdisk/disk8
Reconfiguration to contract physical volume "/dev/rdisk/disk8"
to requested size 102400 MB succeeded.
Previous number of extents: 2395
Number of extents after reconfiguration: 1595
Physical volume "/dev/rdisk/disk8" was changed.

Volume Group configuration for /dev/vgdemo has been saved.
```

3.3 Contracting the LUN

Volumes used with HP-UX 11i v3 and LVM v2 based volume groups can be dynamically contracted in size as shown where the volume is contracted from 150G to 100G in capacity.

```
# vgmodify -r -a -C 100g /dev/vgdemo /dev/rdisk/disk8
Physical volume "/dev/rdisk/disk8" requires reconfiguration
for contraction to requested size 102400 MB.
Current number of extents: 2395
Number of extents after reconfiguration: 1595

The physical volume "/dev/rdisk/disk8" can be reconfigured for contraction
to a minimum size of 50317 MB by removing free extents from the
end of the physical volume.
Physical volume "/dev/rdisk/disk8" was not changed.
```

Note: Volume size contraction is data-unaware and will blindly reduce the size of the volume as instructed without regard of any data which may reside on it. It is recommended to always capture, validate and maintain qualified backups of data prior to any operation of this nature.

Note: Dell Storage Center does not currently support any volume size contraction on the array itself. It is recommended that any data migration or relocation efforts be performed by presenting a secondary and smaller volume to the host, preparing the volume for LVM use and migrating the data via operating system based tools i.e. pvmove, dd etc. and thereafter retiring and reclaiming the larger and unused volume.

4 Dynamic Root Disk

Among other purposes, the HP Dynamic Root Disk (DRD) toolset is available for production and patch lifecycle management, operating system upgrades and actual or practiced disaster recovery exercises. This toolset is named for its intent to work primarily with your root volume group (vg00 under LVM control). This toolset is available for HP-UX 11i v3 using both LVM v2 or LVM v1 volume groups and is available for download on the HP product site.

The existence of this bundle on the HP-UX host is validated using the code below.

```
# swlist -l bundle | grep Dynamic
  DynRootDisk           B.1131.A.3.10.203 Dynamic Root Disk
  DynamicNPars          B.11.31.0709   Dynamic nPartitions enablement
```

This section discusses the use of SC Series volumes with the DRD toolset and its use in creating Boot from SAN volumes.

4.1 Boot from SAN

The DRD toolset is used to create a point-in-time copy of the vg00 volume group or to clone the vg00 volume group to an alternate volume or disk device.

The point-in-time copy can be used either as a running configuration snapshot or as an alternate copy of the vg00 volume group to test and validate new upgrades or patch sets.

In this paper, the DRD toolset is used to clone the vg00 volume group to an SC Series volume and configure this volume as the primary boot from SAN device.

Note: The SC Series volume presented to the HP-UX host should be of equivalent size to the total capacity of the vg00 volume group.

The total capacity of the vg00 volume group is obtained with the following code.

```
# vgdisplay -v vg00 | grep PE | head -3
Max PE per PV               4357
PE Size (Mbytes)            32
Total PE                    4347
```

Where the total capacity of vg00 is calculated by multiplying **PE Size** by **Total PE** (32 x 4347 = 139104 / 1024 = 136G volume).

Present the SC Series volume of similar capacity to the HP-UX host. In this scenario, the /dev/rdisk/disk15 volume of 150G capacity was used for this exercise.

With the SC Series volume in place and discovered on the HP-UX host, issue the command below.

```
# /opt/drd/bin/drd clone -v -x overwrite=true -t /dev/disk/disk15
```

```
=======  11/04/15 14:04:38 CST  BEGIN Clone System Image (user=root)
(jobid=rx3600)

       * Reading Current System Information
       * Selecting System Image To Clone
       * Selecting Target Disk
       * The disk "/dev/disk/disk15" contains data which will be overwritten.
       * Selecting Volume Manager For New System Image
       * Analyzing For System Image Cloning
       * Creating New File Systems
       * Copying File Systems To New System Image
WARNING: The following files could not be copied to the clone.
WARNING: This may be caused by updating files during the copy.
WARNING: Uncopied file: /var/opt/OV/datafiles/coda00000
       * Copying File Systems To New System Image succeeded with 3 warnings.
       * Making New System Image Bootable
       * Unmounting New System Image Clone
       * System image: "sysimage_001" on disk "/dev/disk/disk15"

=======  11/04/15 14:17:05 CST  END Clone System Image succeeded with 3
warnings. (user=root)  (jobid=rx3600)
```

Note: With Itanium-based operating system installations, the **`drd clone`** command only clones the s1 and s2 partitions to the target volume. The s3 partition (HP Service Partition) will not be cloned.

The newly created drd clone volume is verified as shown and unmounted accordingly with the **`drd umount`** command.

```
# /opt/drd/bin/drd mount; bdf

=======  11/04/15 14:18:12 CST  BEGIN Mount Inactive System Image (user=root)
(jobid=rx3600)

       * Checking for Valid Inactive System Image
       * Locating Inactive System Image
       * Preparing To Mount Inactive System Image
       * Selected inactive system image "sysimage_001" on disk
"/dev/disk/disk15".
       * Mounting Inactive System Image

=======  11/04/15 14:18:32 CST  END Mount Inactive System Image succeeded.
(user=root)  (jobid=rx3600)

Filesystem          kbytes    used   avail %used Mounted on
/dev/vg00/lvol3    4292608  402600 3859704    9% /
/dev/vg00/lvol1    2097152  384896 1698944   18% /stand
/dev/vg00/lvol8    8388608 2074624 6272704   25% /var
```

```
/dev/vg00/lvol7    8388608 3270800 5077856   39% /usr
/dev/vg00/lvol6    8388608   21048 8302192    0% /tmp
/dev/vg00/lvol5    16777216 5568792 11120912   33% /opt
/dev/vg00/lvol4    8388608   21088 8302152    0% /home
/dev/drd00/lvol3   4292608  400192 3862096    9% /var/opt/drd/mnts/sysimage_001
/dev/drd00/lvol4   8388608   21088 8302152    0%
/var/opt/drd/mnts/sysimage_001/home
/dev/drd00/lvol5   16777216 5554944 11134656   33%
/var/opt/drd/mnts/sysimage_001/opt
/dev/drd00/lvol1   2097152  384904 1698936   18%
/var/opt/drd/mnts/sysimage_001/stand
/dev/drd00/lvol6   8388608   21048 8302192    0%
/var/opt/drd/mnts/sysimage_001/tmp
/dev/drd00/lvol7   8388608 3198816 5149336   38%
/var/opt/drd/mnts/sysimage_001/usr
/dev/drd00/lvol8   8388608 2001344 6337456   24%
/var/opt/drd/mnts/sysimage_001/var
```

Finally, the drd clone volume is configured as the primary boot from SAN device. This command also initiates a system reboot and attempts to boot from the SAN volume. The drd clone volume should be unmounted from the file system before attempting this command.

```
# /opt/drd/bin/drd activate –x reboot=true
=======  11/04/15 14:26:14 CST  BEGIN Activate Inactive System Image (user=root)
(jobid=rx3600)
[snip]
* Rebooting System
```

Note: The HP-UX 11i v3 IgniteUX C.7.13.259 IINSTALLFS kernel or newer contains the necessary drivers to detect SC Series array presented volumes during an IgniteUX-based installation process. You may either install HP-UX directly to the SC Series volume using the IgniteUX interactive ascii-based GUI or with the drd clone process outlined above.

The boot configuration is verified using the following code.

```
# setboot
Primary bootpath : 0/2/1/0/4/0.0x5000d31000006505.0x4 (/dev/rdisk/disk15)
HA Alternate bootpath :
Alternate bootpath : 0/4/1/0.0x5000cca0000ae311.0x0 (/dev/rdisk/disk2)

Autoboot is ON (enabled)
Hyperthreading : ON
               : ON (next boot)
```

5 Thin provisioning and Replays

All volumes on the SC Series array are thin provisioned by default. Thin provisioning enables interaction between an operating system, volume, file management platform and the SC Series array. In this interaction, the operating system, volume and file management communicate through SCSI requests and calls with the SC Series array and continually monitors the amount of space allocated versus the amount of space actually used. In doing so the volumes can appear to be larger than what is actually allocated on the array. This creates a model that provides just enough storage at just the right time and results in huge savings from both CapEx and OpEx perspectives.

Replays are a feature of SC Series arrays that enable the capture of point-in-time copies of the volume state and data. Replays are metadata snapshots only and for this reason consume minimal additional storage capacity on the array.

5.1 Thin provisioning

HP-UX 11i v3 supports thin provisioning interoperability with the SC Series arrays using the installation of the HP-UX Sep 2011 Support Pack Update release or newer. This release installs the HP-licensed Veritas file system VxFS 5.0.1. The VxFS 5.0.1 release also improves direct I/O in Base-VxFS as well as concurrent I/O in OnlineJFS. More information is available on the HPE Website.

5.1.1 SCSI UNMAP

SCSI UNMAP is achieved on HP-UX via the use of the Veritas File System (VxFS) as well. The VxFS file system should be applied to the logical volume prior to mounting it for use as shown.

```
# newfs -F vxfs /dev/vgdemo/rlvol1
```

The VxFS file system automatically manages free space recovery when inodes become available on the file system and direct the SC Series array accordingly.

Additional information is provided in the *HP-UX VxFS tuning and performance* technical white paper.

5.2 Replays

This section discusses the use of SC Series Replays with HP-UX.

5.2.1 vgchange

Quiesce the volume group prior to capturing a replay. Volume groups can be quiesced without unmounting it from the file system or disabling the volume group.

Note: Consider the business needs before quiencing one or multiple volume groups. Once the volume group is quiesced, the volume group, all its logical volumes and disk devices are placed into a write-locked state and I/O will be interrupted.

```
# vgchange -Q w /dev/vgdemo (use rw parameters to place volume group into read
and write-locked state)
Quiesced writes to volume group "/dev/vgdemo".
Volume group "/dev/vgdemo" has been successfully changed.

# vgdisplay –v /dev/vgdemo
--- Volume groups ---
VG Name                     /dev/vgdemo
VG Write Access             read/write, write-quiesced
VG Status                   available
[snip]
```

Note: Replays on SC Series arrays can also be captured without quiescing the volume groups or interrupting I/O and integrity of data. This procedure should be validated within respective business environments to deem its suitability for production data use cases.

The volume group is resumed to normal operating I/O state with the command:

```
# vgchange -R /dev/vgdemo
```

5.2.2 Capture the Replay

Capturing a Replay of an SC volume is described below.

1. Right click on the target volume in the Dell Enterprise Manager and click **Create Replay**.

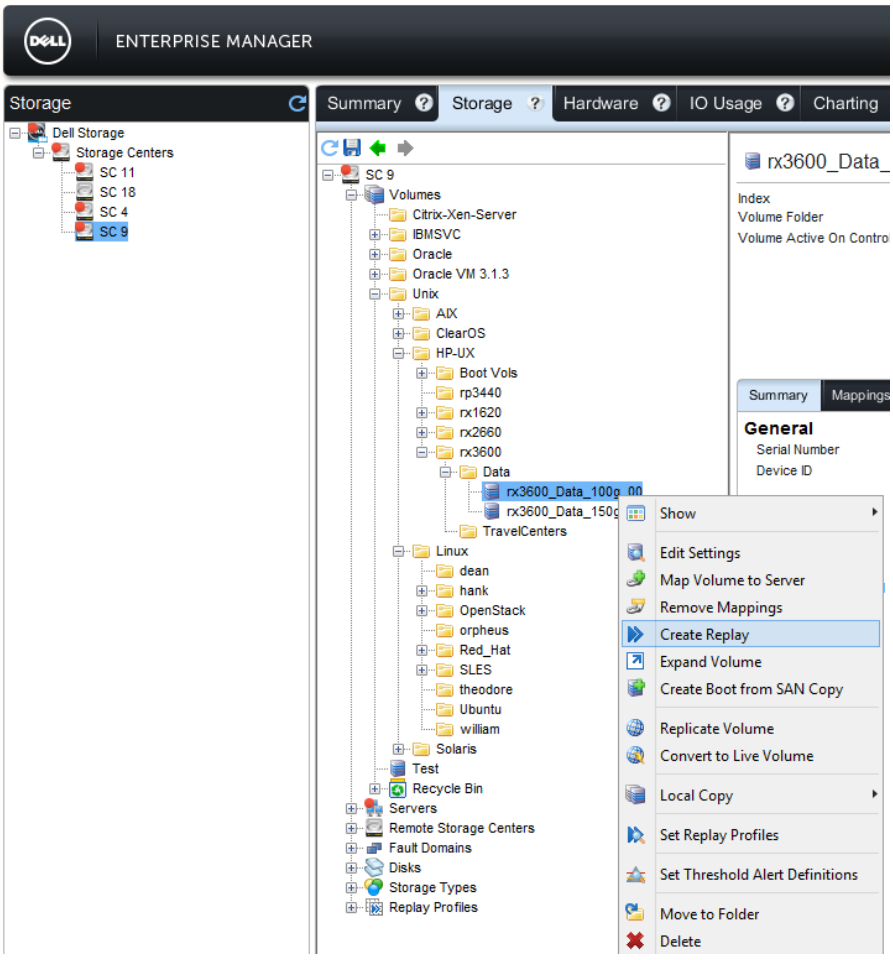

2. Select **Do Not Expire** and enter a description for the volume; Click **OK**.

3. View the replays for this volume in the **Replays** tab.

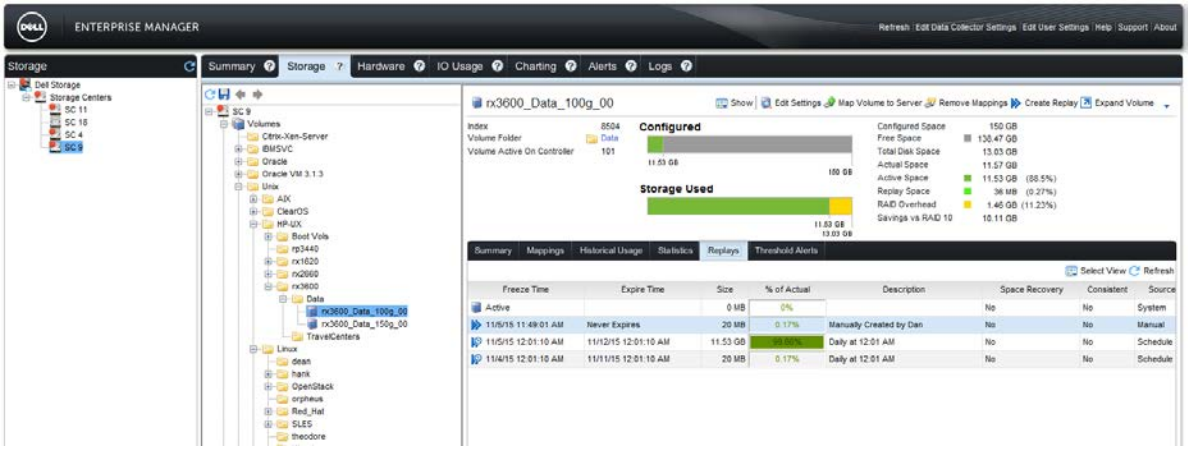

4. Create additional volumes from this replay and present back to the same host or an alternate host as required. These new volumes are referenced as replay view volumes.

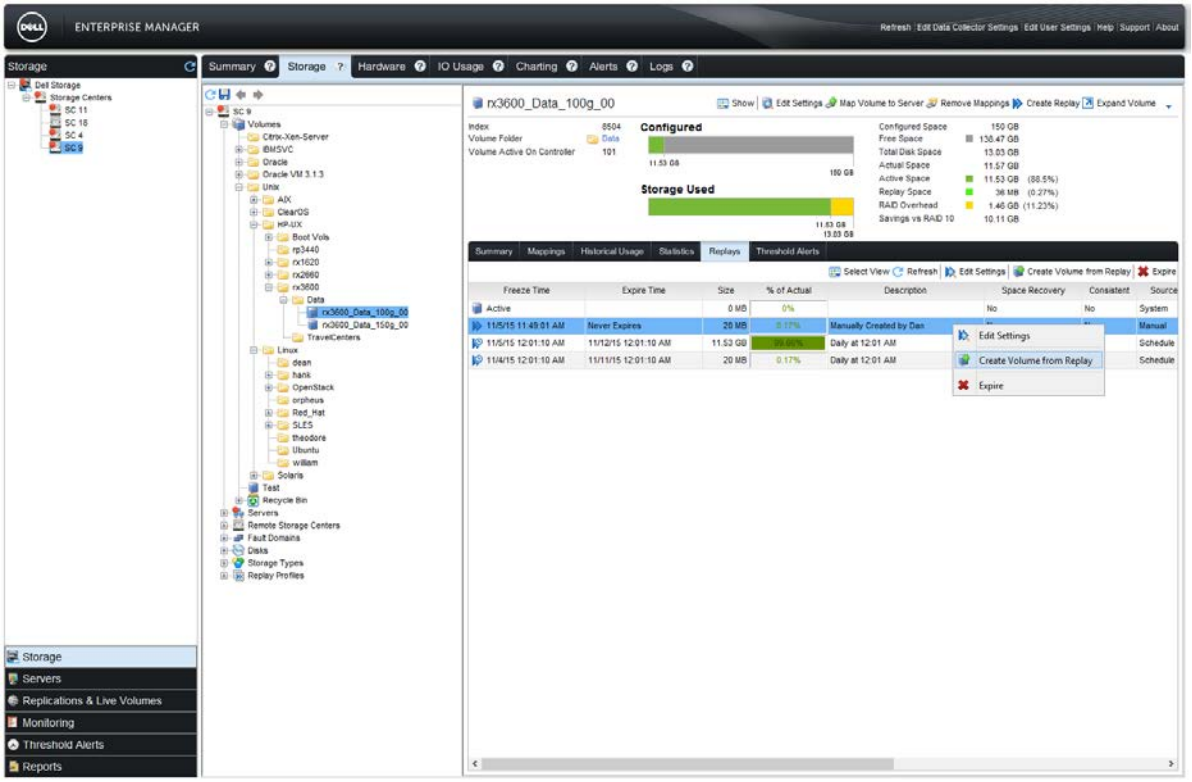

5.2.3 Mounting to the same host

Prior to presenting a replay view volume to the same host, capture the following data using the command line interface.

```
# ioscan -kfnNC disk
Class     I  H/W Path  Driver S/W State   H/W Type     Description
==========================================================================
disk      2  64000/0xfa00/0x0   esdisk   CLAIMED     DEVICE       HP
DG146A4960
                      /dev/disk/disk2      /dev/disk/disk2_p3
/dev/rdisk/disk2_p2
                      /dev/disk/disk2_p1   /dev/rdisk/disk2
/dev/rdisk/disk2_p3
                      /dev/disk/disk2_p2   /dev/rdisk/disk2_p1
[snip]
disk      5  64000/0xfa00/0x2   esdisk   CLAIMED     DEVICE       TEAC    DV-
28E-V
                      /dev/disk/disk5   /dev/rdisk/disk5
disk      8  64000/0xfa00/0xb   esdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                      /dev/disk/disk8      /dev/disk/disk8_p3
/dev/rdisk/disk8_p2
                      /dev/disk/disk8_p1   /dev/rdisk/disk8
/dev/rdisk/disk8_p3
                      /dev/disk/disk8_p2   /dev/rdisk/disk8_p1
disk     15  64000/0xfa00/0x10  esdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                      /dev/disk/disk15   /dev/rdisk/disk15
# strings /etc/lvmtab
/dev/vg00
/dev/disk/disk2_p2
# strings /etc/lvmtab_p
/dev/vgdemo
A0000000000000002Thu Nov  5 17:38:21 20151ab0bf92-18a9-11dd-af0f-3b63380fc243
/dev/disk/disk8
```

Additionally, capture a vgexport map of the /dev/vgdemo volume group. The map file is still created with the -p (preview) flag. However, the vgdemo volume group is not actually exported.

```
# vgexport -p -m /tmp/vgdemo.map vgdemo
vgexport: Volume group "vgdemo" is still active.
vgexport: Preview of vgexport on volume group "vgdemo" succeeded.
```

Present the replay view volume from the SC Series array to the host and run the commands as shown where the new volume is discovered as /dev/disk/disk18.

```
# insf –e; ioscan -fnNC disk
```

```
[snip]
disk      8  64000/0xfa00/0xb   esdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                     /dev/disk/disk8      /dev/disk/disk8_p3
/dev/rdisk/disk8_p2
                     /dev/disk/disk8_p1   /dev/rdisk/disk8
/dev/rdisk/disk8_p3
                     /dev/disk/disk8_p2   /dev/rdisk/disk8_p1
disk     15  64000/0xfa00/0x10  esdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                     /dev/disk/disk15   /dev/rdisk/disk15
disk     18  64000/0xfa00/0x11  esdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                     /dev/disk/disk18   /dev/rdisk/disk18
```

The replay view volume /dev/disk/disk18 contains the same UUID as its parent volume /dev/disk/disk8. This UUID needs to be changed with the following command.

vgchgid /dev/rdisk/disk18

The replay view volume is then imported and activated.

```
# vgimport -v vgdemo_import /dev/disk/disk18
Beginning the import process on Volume Group "vgdemo_import".
Logical volume "/dev/vgdemo_import/lvol1" has been successfully created
with lv number 1.
vgimport: Volume group "/dev/vgdemo_import" has been successfully created.
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.
# vgchange -a y /dev/vgdemo_import
Activated volume group.
Volume group "/dev/vgdemo_import" has been successfully changed.
```

The volume group, folder and group files are automatically created and incremented as required.

```
# ls -la /dev/*/group
crw-r-----   1 root       sys         64 0x000000 Oct 28 13:38 /dev/vg00/group
crw-r-----   1 root       sys        128 0x000000 Nov  3 15:11 /dev/vgdemo/group
crw-r-----   1 root       sys        128 0x001000 Nov  5 13:19
/dev/vgdemo_import/group
```

The directory structure is created, the logical volume(s) are verified for integrity and mounted to the file system.

```
# mkdir -p /vgdemo_import; fsck -F vxfs /dev/vgdemo_import/lvol1; mount
/dev/vgdemo_import/lvol1 /vgdemo_import; bdf
log replay in progress
replay complete - marking super-block as CLEAN
```

```
Filesystem          kbytes    used   avail %used Mounted on
/dev/vg00/lvol3    4292608  403224 3859080    9% /
/dev/vg00/lvol1    2097152  384896 1698944   18% /stand
/dev/vg00/lvol8    8388608 2077872 6269512   25% /var
/dev/vg00/lvol7    8388608 3270800 5077856   39% /usr
/dev/vg00/lvol6    8388608   21056 8302184    0% /tmp
/dev/vg00/lvol5    16777216 5568792 11120912   33% /opt
/dev/vg00/lvol4    8388608   21088 8302152    0% /home
/dev/vgdemo/lvol1  51249152   79205 47971833    0% /vgdemo
/dev/vgdemo_import/lvol1
                   51249152   79205 47971833    0% /vgdemo_import
```

5.2.4 Mounting to an alternate host

The methods used to mount a replay view volume to an alternate host mirrors the section above with certain additional caveats.

The /tmp/vgdemo.map file previously created needs to be relocated to the target host where the replay view volume will be presented. Since this volume is new to this target host, there is no UUID conflict and the `vgchgid` command is not used.

The `vgimport` command is issued as shown.

```
# vgimport –m /tmp/vgdemo.map –v vgdemo_import /dev/disk/diskX
# vgchange –a y /dev/vgdemo_import
```

Where X represents the diskX value from an `ioscan` command.

The directory structure is created, the logical volume(s) are verified for integrity and mounted to the file system.

Note: HP-UX 11i v2 and older platforms use LVM v1 structures while HP-UX 11i v3 uses LVM v1, v2 or v2.1 structures. LVM structures are **not** backward compatible. A volume group and logical volume created on HP-UX 11i v3 is accessible when mounted on HP-UX 11i v1. In addition, VxFS versions are **not** backward compatible either. Even if you create an LVM v1 volume group with HP-UX 11i v3, the file system that you place on any logical volumes with HP-UX 11i v3 will not be suitable for use by any previous version of HP-UX. For all practical purposes, replay(s) are operating system forward compatible but not backward compatible. This is a limitation of the HP-UX, LVM structures and VxFS and not a part of Storage Center Replay functionality.

5.2.5 Consistency groups

SC Series consistency groups is a feature that allows multiple volumes on the array to be flagged as logical members of a single entity or container. The replay capture of this entity or container ensures that all members within it have a consistent snapshot at the same point in time.

This feature is especially useful and applicable when working with multi-device based volume groups (concatenated, striped, mirrored or otherwise).

Consistency groups are managed using Storage Center Replay profiles applied to volumes that are intended to be members of this entity.

Consistency groups are setup and configured as shown below.

1. Right click **Replay Profiles** and click **Create Replay Profile**.

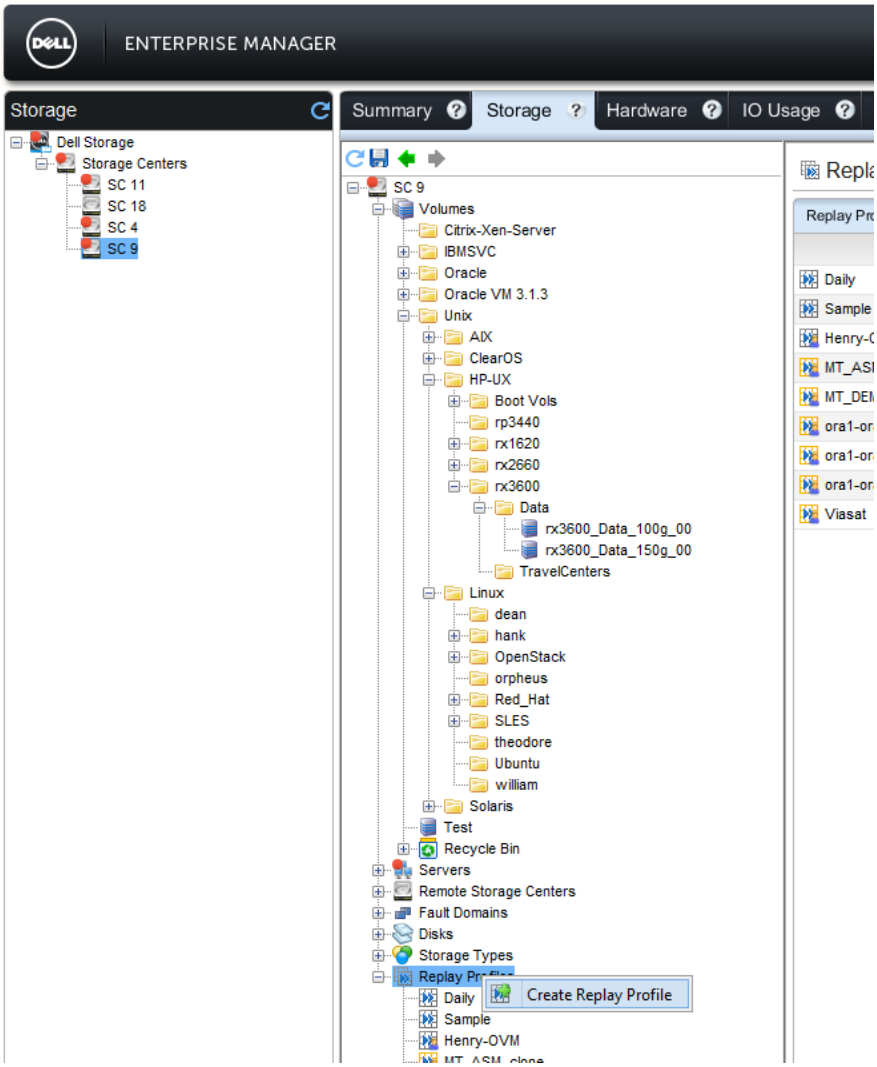

2. Name the replay profile and in the Replay Create Method, select **Consistent**. Click **Add Rule** to open the Schedule dialog.

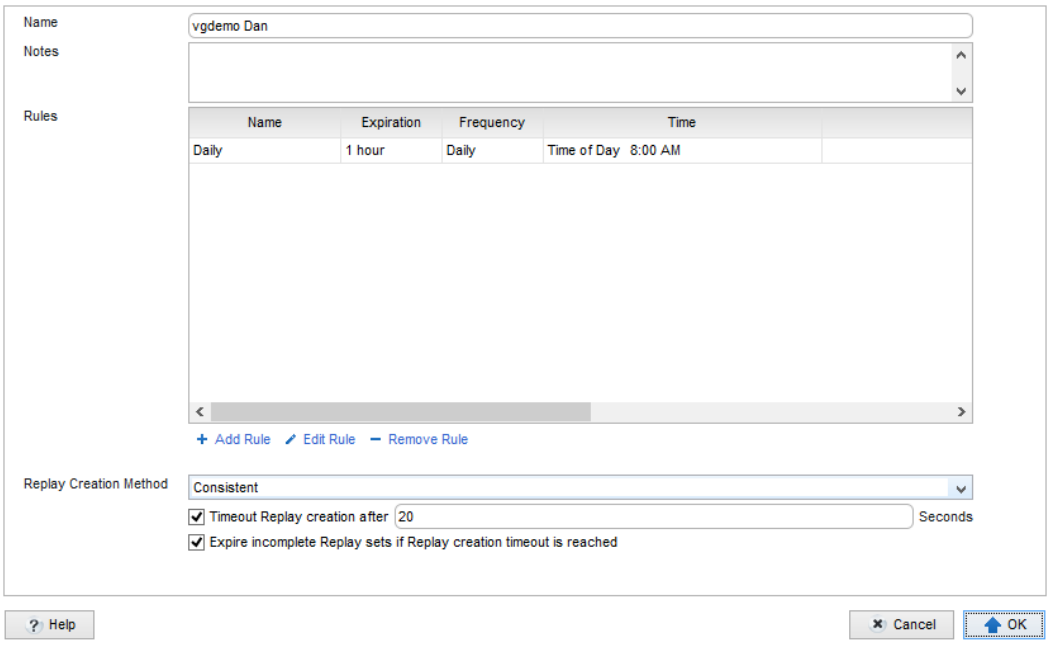

3. Define a basic daily schedule, click **OK** and then click **OK** once more to complete the creation of the replay profile.

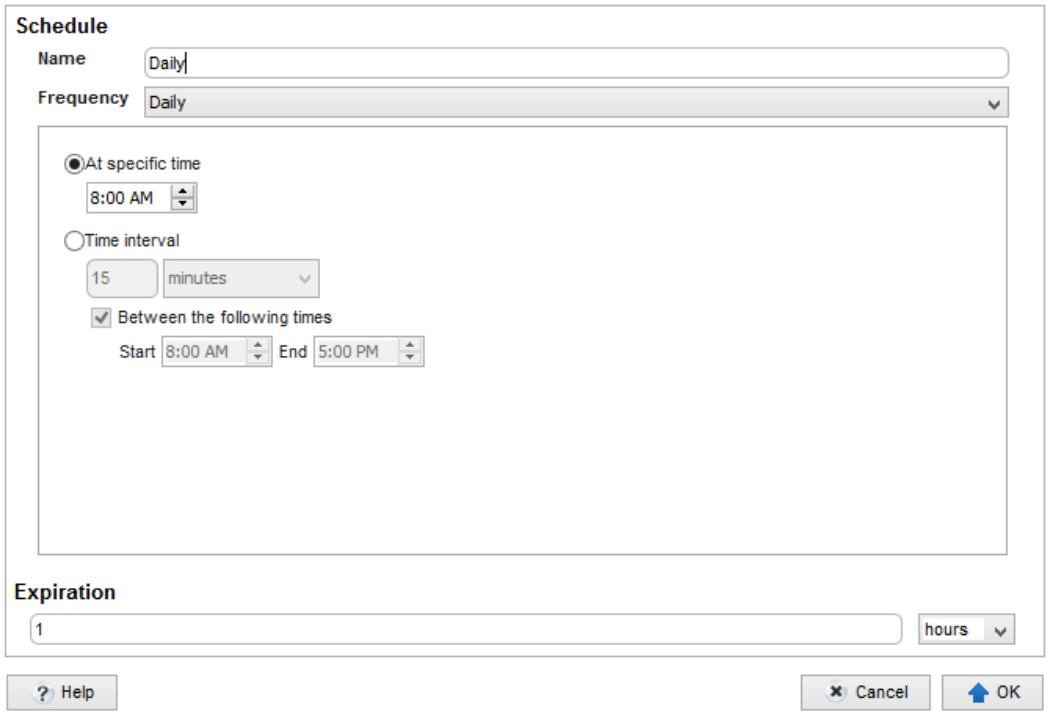

4. Right click on the new replay profile and select **Apply to Volumes**.

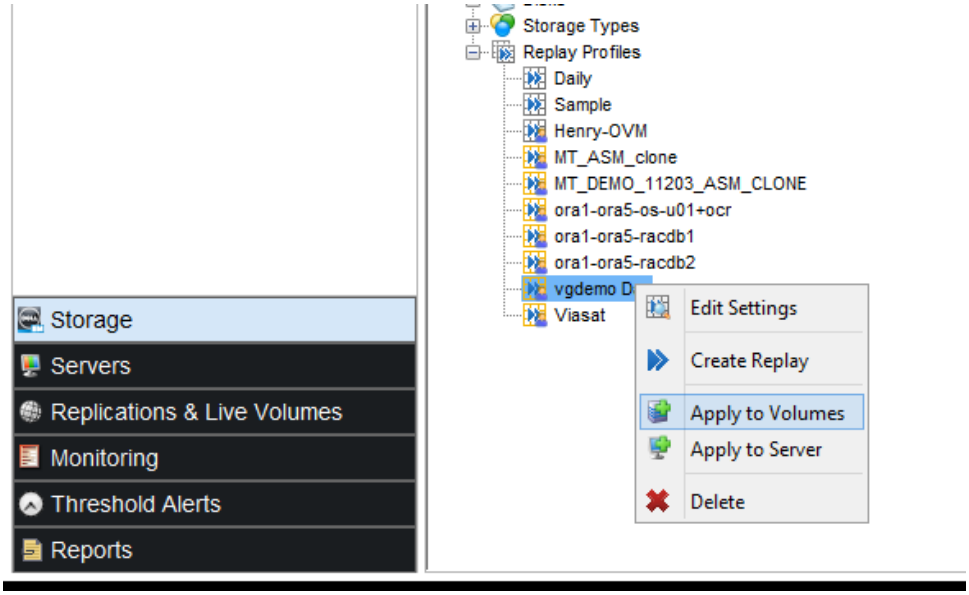

5. Select the volumes intended to be members of this replay profile and click **OK**.

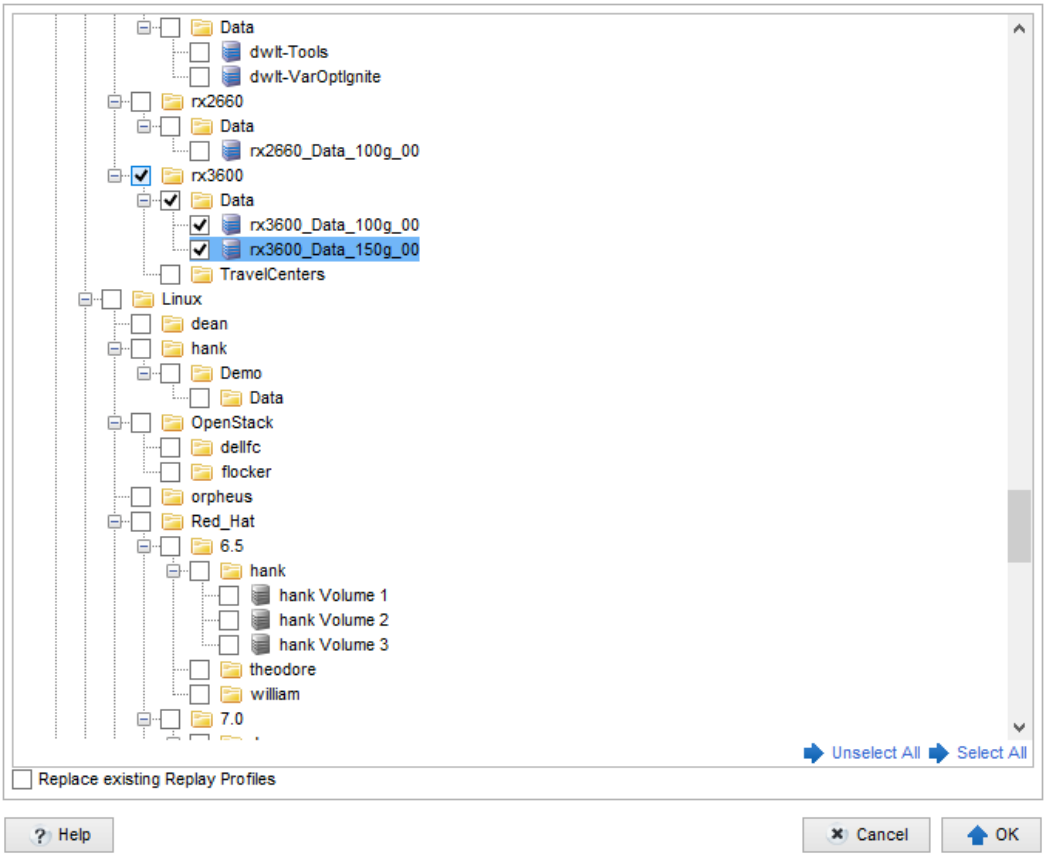

6. The volumes are now members of the replay profile. To perform a consistent replay of all volumes, right click on the replay profile and click **Create Replay**.

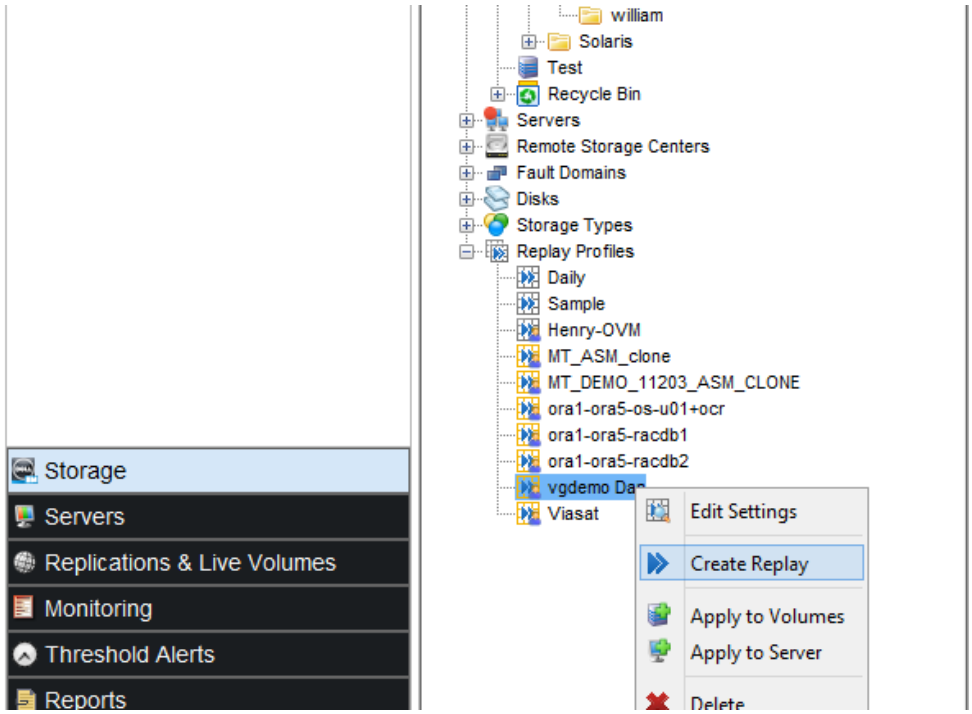

6 iSCSI with HP-UX

The use of iSCSI with HP-UX allows the use and reuse of existing Ethernet-based networks to carry storage data traffic encapsulated within IP packets. iSCSI can be used to facilitate data transfer over the internet or local LAN and to manage storage over long distances.

6.1 Installation

The HP-UX iSCSI initiator is established through the installation of the iSCSI-00* and ISCSI-SWD* filesets. This installation requires a reboot of the HP-UX host.

```
# swinstall -p -x autoreboot=true -s
rx1620.techsol.beer.town:/var/opt/ignite/depots/Misc/iSCSI \*

=======  11/10/15 15:03:28 CST  BEGIN swinstall SESSION
         (non-interactive) (jobid=rx3600-0011)

       * Session started for user "root@rx3600".

       * Beginning Selection
       * Target connection succeeded for "rx3600:/".
       * Source connection succeeded for
         "rx1620.techsol.beer.town:/var/opt/ignite/depots/Misc/iSCSI".
       * Source:
         rx1620.techsol.beer.town:/var/opt/ignite/depots/Misc/iSCSI
       * Targets:               rx3600:/
       * Software selections:
             iSCSI-00,r=B.11.31.03c,a=HP-UX_B.11.31_IA/PA,v=HP
             ISCSI-SWD.ISCSI-KRN,r=B.11.31.03c,a=HP-
UX_B.11.31_IA/PA,v=HP,fr=B.11.31.03c,fa=HP-UX_B.11.31_IA
             ISCSI-SWD.ISCSI-MAN,r=B.11.31.03c,a=HP-
UX_B.11.31_IA/PA,v=HP,fr=B.11.31.03c,fa=HP-UX_B.11.31_IA/PA
             ISCSI-SWD.ISCSI-RUN,r=B.11.31.03c,a=HP-
UX_B.11.31_IA/PA,v=HP,fr=B.11.31.03c,fa=HP-UX_B.11.31_IA
       * Selection succeeded.


       * Beginning Analysis
       * Session selections have been saved in the file
         "/.sw/sessions/swinstall.last".
       * The analysis phase succeeded for "rx3600:/".
       * Analysis succeeded.


NOTE:    More information may be found in the agent logfile using the
         command "swjob -a log rx3600-0011 @ rx3600:/".
```

```
=======  11/10/15 15:03:30 CST  END swinstall SESSION (non-interactive)
         (jobid=rx3600-0011)
```

A successful installation of the iSCSI bundle is shown.

```
# swlist -l bundle | grep iSCSI
  iSCSI-00              B.11.31.03c    HP-UX iSCSI Software Initiator
```

6.2 Configuration

Verify that the necessary iSCSI processes are running.

```
# ps -ef | grep iscsi
    root  6468     1  0 15:20:01 ?         0:00 /opt/iscsi/bin/iscsi_resolvd
    root  6470     1  0 15:20:01 ?         0:00 /opt/iscsi/bin/iswd
    root  8056  6902  0 15:23:10 pts/0     0:00 grep iscsi
```

The iSCSI Qualified Name (iqn) is required to establish a connection from the HP-UX host to the SC Series array.

```
# /opt/iscsi/bin/iscsiutil -l | grep Ini
Initiator Name            : iqn.1986-03.com.hp:rx3600.1ab0bf92-18a9-11dd-af0f-
3b63380fc243
Initiator Alias           :
Initiator CHAP Name       :
[snip]
```

The system generated iqn although adequate for setting up an iSCSI-based connection to the array, can be lengthy and unwieldy for use. It is recommended to change this iqn to a string which is more manageable and meaningful. A generally accepted rule of iqn naming is to replace the string after the semicolon to display <hostname>.<MAC address of the Ethernet interface where the iSCSI connection is bound>.

Since this iSCSI connection is bound to the lan2 interface, the MAC address is obtained.

```
# lanscan
Hardware Station        Crd Hdw   Net-Interface  NM  MAC       HP-DLPI DLPI
Path     Address        In# State NamePPA        ID  Type      Support Mjr#
0/2/1/0/6/0 0x001E0BFDA12E 0   UP    lan0 snap0     1   ETHER     Yes     119
0/2/1/0/6/1 0x001E0BFDA12F 1   UP    lan1 snap1     2   ETHER     Yes     119
0/4/2/0  0x001A4B0804B2 2   UP    lan2 snap2     3   ETHER     Yes     119
[snip]
```

The iSCSI iqn is renamed.

```
# /opt/iscsi/bin/iscsiutil -i -N iqn.1986-03.com.hp:rx3600.001A4B0804B2
iscsiutil: Initiator Name "iqn.1986-03.com.hp:rx3600.001A4B0804B2" has been
successfully updated.
```

```
bash-4.2# /opt/iscsi/bin/iscsiutil -l | grep Ini
Initiator Name            : iqn.1986-03.com.hp:rx3600.001A4B0804B2
Initiator Alias           :
Initiator CHAP Name       :
[snip]
```

This iqn is used in the Storage Center interface to establish a new iSCSI-based server object. The iSCSI initiator is logged into the target discovery IP port of the iSCSI interface on the SC Series array. If there is more than one target discovery IP port on the array, they are added with each successive run of this command.

```
# /opt/iscsi/bin/iscsiutil -a –I <IP address>
Target address “<IP address:3260.1> has been successfully added.
```

Note: Ensure that the network interfaces on the HP-UX host are properly connected and configured using the proper switches and VLAN as needed.

```
# /opt/iscsi/bin/iscsiutil –p -D
Discovery Target Information
----------------------------

Target # 1
-----------
        IP Address                : <IP address>
        iSCSI TCP Port            : 3260
        iSCSI Portal Group Tag    : 1
[snip]
```

The ioscan command is used to initiate the iSCSI login into the SC Series array.

```
# ioscan -fnNC iscsi
Class     I  H/W Path   Driver S/W State   H/W Type     Description
=================================================================
iscsi     0  64000/0x2  iscsi   CLAIMED     VIRTBUS      iSCSI Virtual Root
```

Note: The example above has only setup a single target discovery IP address and bound to a single LAN interface. It is recommended in production use to setup more than one target discovery IP address across multiple VLAN on the HP-UX host to ensure proper failover and redundancy in case of failure.

Volumes are created on the array and mapped to the newly created iSCSI-based server object. These new volumes are discovered on the HP-UX host.

```
# ioscan –fnNH 6400
iscsi        0  64000/0x2      iscsi          SCAN        VIRTBUS      iSCSI
Virtual Root
escsi_ctlr   2  64000/0x2/0x0  isvctlr        SCAN        INTERFACE    iSCSI
Virtual Controller
```

```
tgtpath     26  64000/0x2/0x0.0x0                    estp          SCAN
TGT_PATH     iscsi target served by isvctlr driver, target port id 0x0
lunpath     26  64000/0x2/0x0.0x0.0x0                eslpt         SCAN
LUN_PATH     LUN path for ctl50
lunpath     27  64000/0x2/0x0.0x0.0x4001000000000000  eslpt         SCAN
LUN_PATH     LUN path for disk12
tgtpath     27  64000/0x2/0x0.0x1                    estp          SCAN
TGT_PATH     iscsi target served by isvctlr driver, target port id 0x1
lunpath     23  64000/0x2/0x0.0x1.0x0                eslpt         SCAN
LUN_PATH     LUN path for ctl48
lunpath     24  64000/0x2/0x0.0x1.0x4001000000000000  eslpt         SCAN
LUN_PATH     LUN path for disk12
tgtpath     29  64000/0x2/0x0.0x2                    estp          SCAN
TGT_PATH     iscsi target served by isvctlr driver, target port id 0x2
lunpath     25  64000/0x2/0x0.0x2.0x0                eslpt         SCAN
LUN_PATH     LUN path for ctl49
lunpath     28  64000/0x2/0x0.0x2.0x4001000000000000  eslpt         SCAN
LUN_PATH     LUN path for disk9
tgtpath     28  64000/0x2/0x0.0x3                    estp          SCAN
TGT_PATH     iscsi target served by isvctlr driver, target port id 0x3
lunpath     21  64000/0x2/0x0.0x3.0x0                eslpt         SCAN
LUN_PATH     LUN path for ctl47
lunpath     22  64000/0x2/0x0.0x3.0x4001000000000000  eslpt         SCAN
LUN_PATH     LUN path for disk9
```

The device files are built to support these new devices and an `ioscan` command is shown of these new iSCSI-based devices.

```
# insf -e; ioscan -fnNC disk
Class     I  H/W Path  Driver S/W State   H/W Type     Description
===================================================================
[snip]
disk      8  64000/0xfa00/0x18  esdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                      /dev/disk/disk8   /dev/rdisk/disk9
disk     12  64000/0xfa00/0x1a  esdisk   CLAIMED     DEVICE
COMPELNTCompellent Vol
                      /dev/disk/disk12   /dev/rdisk/disk12
```

7 Performance

Proper configuration and tuning of kernel, SCSI, network and filesystem parameters may be required to achieve optimized I/O with the operating system. The process of configuration and tuning is not a science but an art of understanding the needs of the business and application workload in which the HP-UX host(s) are to serve.

The following recommendations are for using Storage Center with HP-UX. It is not meant to be a complete nor comprehensive representation of all possible tuning scenarios within managed HP-UX environments.

7.1 Kernel

HP-UX host performance baselining can be performed with the freely available HP Tune N Tools bundle. This bundle is meant for HP-UX 11i v3 use only on Itanium or PA-RISC architecture. The installation of this bundle does not require a reboot.

The bundle, usage and installation instructions are obtained on the HP Inc. Software Depot.

The .depot file is downloaded and installed with the native `swinstall` command toolset. A correctly installed bundle is:

```
# swlist -l bundle | grep Tune
  Tune-N-Tools          B.11.31.0909   Optimized Kernel Tunables and Tools for
Database and Application Servers
```

The first command to run is shown below. This provides an assessment of the HP-UX, taking into consideration all host-based resources (such as CPU, disk, RAM and swap) and provides a recommendation towards each kernel parameter setting.

```
# /opt/tuneserver/bin/tuneserver -l

The optimal tunable settings for server applications are:
    max_async_ports     should be at least 27000
    max_thread_proc     should be at least 3000
    maxdsiz             should be at least 2147483647
    maxdsiz_64bit       should be at least 274877906944
    maxfiles            should be at least 8192
    maxfiles_lim        should be at least 8192
    maxssiz             should be at least 134217728
    maxssiz_64bit       should be at least 1073741824
    maxtsiz             should be at least 1073741824
    maxtsiz_64bit       should be at least 8589934592
    maxuprc             should be at least 27000
[snip]
```

A system performance should be captured at this time using any preferred toolset (such as vdbench or fio) during nominal operating business and application workload.

The following command is run to automatically apply all the recommended kernel changes shown above. A reboot is recommended afterwards to ensure all the parameters are applied.

```
# /opt/tuneserver/bin/tuneserver -o
```

The following command can be used to revert the HP-UX host to its original kernel parameter state as required. A reboot is recommended thereafter to ensure all parameters are applied.

```
# /opt/tuneserver/bin/tuneserver -r
```

An additional resource regarding HP-UX performance management is the *HP-UX Performance Cookbook*.

7.2 SCSI

The application, understanding and management of SCSI timeout and queue depth parameters can greatly optimize I/O performance on a HP-UX host. These discussed parameters should be applied and used as a recommended starting point and adapted accordingly to the varying needs of the business and application workload landscape.

7.2.1 Queue depth

Run the following command against each presented SC Series volume to increase the stored I/O command queue depth to the recommended value of 32.

```
# /usr/sbin/scsimgr save_attr -D /dev/rdisk/diskX -a max_q_depth=32
```

The above command comes into effect on the next reboot of the HP-UX host. This command below will apply the queue depth change to the volume on the fly.

```
# /usr/sbin/scsictl –m queue_depth=32 /dev/rdisk/diskX
```

7.2.2 SCSI timeout

SCSI timeout on HP-UX 11i v3 is managed through several commands discussed in this section.

The `pvchange` command is used to define and set the limit in seconds for LVM operations to a volume. The default value is 30 and should be left at this value. However, if the value can be changed if needed by using:

```
# pvchange -t 30 /dev/rdisk/diskX
```

The `scsimgr` command is used to manipulate the esd_secs and path_fail_secs attributes of a volume. The first attribute esd_secs defines the esdisk driver timeout which defaults to 30 seconds and should be left at this value.

The second attribute (path_fail_secs) is especially relevant to multipath volumes and defines the time to wait before declaring a volume path as unavailable and moving I/O to the next available volume path. The default value is 120 seconds. It is recommended to change this value on all SC Series volumes to 60 seconds and thereby reduce latency in I/O should a path become unavailable.

```
# /usr/sbin/scsimgr save_attr -D /dev/rdisk/diskX -a path_fail_secs=60
```

This attribute can also be changed on the fly on a per volume basis with the following command.

```
# /usr/sbin/scsictl –m path_fail_secs=60 /dev/rdisk/diskX
```

7.3 File system

Use the following command when creating a file system on a new logical volume. This command will apply an 8K block size as part of the filesystem build options.

```
# newfs –F vxfs –o largefiles –b 8192 /dev/vgdemo/rlvolX
```

8 Dell Compellent Command Utility (CompCU)

The Storage Center can have many of its daily functions managed through a remote command utility called the Dell Compellent Command Utility (CompCU). This allows for scripting and automation integration of SAN tasks between Linux® and Storage Center. CompCU is a java-packaged application and therefore requires the installation of Java on the Linux host. CompCU can be used to script common administrative tasks that can be tremendous time savers and provide a consistent framework for managing Storage Center volumes and Replays.

A Linux platform is used for the purpose of this CompCU demo, however every command discussed in this section may be used in similar fashion on any UNIX/Linux platform as long as it meets the required Java runtime configurations.

CompCU requires the host to have the proper Java release installed. The CompCU.jar object can be downloaded from the Dell Storage Center support site. Once installed on the Linux host, this tool can be used to perform Storage Center tasks from the shell prompt, which can be incorporated into new or existing user management scripts. Outlined below are some common use cases for CompCU:

- Creating Volumes, mapping to the server.
- Taking Replays, recovering Replays, and other related tasks.

The examples below do not cover the full breadth of the usefulness of CompCU by any means; they are designed to give an initial insight into the types of tasks that may be automated with CompCU.

8.1 Verifying Java, configuring and testing CompCU functions

First, install Java (RTE v1.6.x or newer) on the Linux host. The Java runtime may have already been installed with the OS and can be verified with the command shown below.

```
# /usr/bin/java –version
java version "1.7.0_07"
Java(TM) SE Runtime Environment (build 1.7.0_07-b10)
Java HotSpot(TM) Server VM (build 23.3-b01, mixed mode)
```

Download the CompCU package from the Storage Center support site (login required). The package includes a PDF user guide as well as the required CompCU.jar file. Save this CompCU.jar file to a logical filesystem location. Verify that CompCU is working with Java by executing the command below to display the help and usage syntax.

```
# /usr/bin/java -jar ./CompCU.jar -h
Compellent Command Utility (CompCU) 6.4.1.1

usage: java -jar CompCU.jar [Options] "<Command>"
 -c <arg>              Run a single command (option must be within
                       quotes)
 -default              Saves host, user, and password to encrypted file
 -defaultname <arg>    File name to save default host, user, and password
```

```
                        encrypted file to
 -file <arg>            Save output to a file
 -h                     Show help message
 -host <arg>            IP Address/Host Name of Storage Center Management
                        IP
 -password <arg>        Password of user
 -s <arg>               Run a script of commands
 -user <arg>            User to log into Storage Center with
 -verbose               Show debug output
 -xmloutputfile <arg>   File name to save the CompCU return code in xml
                        format.  Default is cu_output.xml.
[snip]
```

To facilitate the ease of access in using CompCU, the tool can be initially run with the -default switch to configure an encrypted password file as shown below. A file named default.cli is created in the local directory. This file may be renamed as required for clarity and usage.

```
# /usr/bin/java -jar ./CompCU.jar  -default -host 172.16.2.109 -user Admin -
password XXX
Compellent Command Utility (CompCU) 6.4.1.1

================================================================================
User Name:              Admin
Host/IP Address:        172.16.2.109
================================================================================
Connecting to Storage Center: 172.16.2.109 with user: Admin
java.lang.IllegalStateException: TrustManagerFactoryImpl is not initialized
Saving CompCu Defaults to file [default.cli]...
The “default.cli” file may then be referenced in other commands to login to the
same Storage Center and perform tasks. A separate .cli file may be created for
each Storage Center under management with each containing the appropriate login
credentials for the respective Storage Center array. The example below
demonstrates a “volume show” command applied to the Storage Center located at IP
address 172.16.2.109.
# /usr/bin/java -jar ./CompCU.jar  -defaultname default.cli -host 172.16.2.109 -
user Admin -password XXX -c "volume show"
Compellent Command Utility (CompCU) 6.4.1.1

================================================================================
User Name:              Admin
Host/IP Address:        172.16.2.109
Single Command:         volume show
================================================================================
Connecting to Storage Center: 172.16.2.109 with user: Admin
java.lang.IllegalStateException: TrustManagerFactoryImpl is not initialized
Running Command: volume show
```

```
Index  Name                                Status   ConfigSize      ActiveSize
ReplaySize      Folder
StorageProfile                   DeviceI              D
SerialNumber      ConfigSizeBlock ActiveSizeBlock ReplaySizeBlock
MaxWriteSizeBlo ReadCache  WriteCache
------ -------------------------------- -------- -------------- --------------
-------------- ------------------------------------------------------------------
-------------------------------- -------              ------------------------
----------------- --------------- --------------- --------------- --------------
- ---------- ----------
283    Fal-asm-mirror-test-failgroup1   Up       100.00 GB       31.33 GB
0.00 KB         Oracle/11gR2/ASM-Mirror
Recommended                      6000d31              00000650000000000000012f
00000065-0000012f 209715200       65712128        0               0
Enabled    Enabled
290    Fal-asm-mirror-sp-fg1            Up       100.00 GB       7.94 GB
0.00 KB         Oracle/11gR2/ASM-Mirror
Recommended                      6000d31              000006500000000000000136
00000065-00000136 209715200       16658432        0               0
Enabled    Enabled
824    ibmsvc00-managed-mdisk1          Up       500.00 GB       98.02 GB
0.00 KB         IBMSVC
Recommended                      6000d31              00000650000000000000034c
00000065-0000034c 1048576000      205561856       0               0
Enabled    Enabled
[snip]
Successfully finished running Compellent Command Utility (CompCU) application.
```

8.2 Using CompCU to automate common tasks

This section illustrates some use cases for managing Storage Center tasks with CompCU on Linux. As mentioned above, these examples are indicative of the types of tasks which can easily be accomplished from the Linux shell prompt using CompCU. They are only meant as a starting point to familiarize the system administrator with this powerful tool set.

8.2.1 Creating a single volume with CompCU

This example demonstrates using CompCU to create a single 100GB Storage Center volume named hadrian_100g_00 from the Linux host and then placed in the Storage Center Linux folder. The volume is mapped to the Linux hadrian host.

```
# /usr/bin/java -jar ./CompCU.jar  -defaultname default.cli -host 172.16.2.109 -
user Admin -password XXX -c "volume create -name hadrian_100g_00 -folder Linux -
server hadrian -size 100g"
```

8.2.2 Creating a Replay and a Replay View with CompCU

This example demonstrates the following with a single CompCU command:

- Creating a Replay, hadrian_100g_00_Replay of the existing hadrian_100g_00 volume on Storage Center,
- Creating Replay View, hadrian_100g_00_View from this mentioned Replay, and
- Mapping the Replay View to the Linux host, maximus

```
# /usr/bin/java -jar ./CompCU.jar  -defaultname default.cli -host 172.16.2.109 -
user Admin -password XXX -c "replay create -volume 'hadrian_100g_00' -name
'hadrian_100g_00_Replay' -view 'hadrian_100g_00_RpView' -server 'maximus'"
```

8.2.3 Rapid deployment of multiple volumes with CompCU

This final example demonstrates using CompCU for rapid volume creation and deployment from Storage Center and mapping these volumes to the Linux host, maximus.

```
# for i in 0 1 2 3 4 5 6 7 8 9; do /usr/bin/java -jar ./CompCU.jar  -defaultname
default.cli -host 172.16.2.109 -user Admin -password XXX -c "volume create –name
maximus_10g_0${i} –folder Linux -server 'maximus' –size 10g"; done
```

A Configuration details

Table 1 Component table

Component	Description
OS	HP-UX 11i v3 0903
Driver version	NA
Firmware version	NA
Switch	Brocade
Cabling	Fiber Channel
Server	HP Itanium rx36xx
Storage	Dell Storage Center OS (SCOS) 6.5.30 @ Virtual port mode
Enterprise Manager	2015 R1

B Additional resources

B.1 Technical support and resources

For Copilot support of Dell SC Series products:

- Global online support
- Email: support@compellent.com (non-emergency business hours)
- Phone: 866-EZ-STORE (866-397-8673) (United States only)

The Dell SC Series Portal is an online portal for existing customers. A valid portal account is required to access the Knowledge Center. Once login to the portal, go to "Knowledge center".

Dell TechCenter is an online technical community for IT professionals and is a great resource to discover and learn about a wide range of technologies such as storage, servers, networking, software, and cloud management.

B.2 Related documentation

Table 2 lists the referenced or recommended resources related to this document.

Table 2 Referenced or recommended resources

Vendor	Resource
HP	http://www8.hp.com/us/en/products/servers/hp-ux.html#.VnCLenmhpaQ
HP Roadmap	http://www8.hp.com/h20195/v2/GetPDF.aspx/4AA3-9071ENW.pdf
HPE	https://www.hpe.com/us/en/home.html?jumpid=ps_529xnyrcq3::AID-510018643&gclid=CKGdouzj-cgCFSdcMgodB1wHZg&gclsrc=ds