

LVM Migration from Legacy to Agile Naming Model

HP-UX 11i v3



Abstract	2
Legacy and Agile Naming Models	3
LVM support for Dual Naming Models	3
Naming Model specific Commands/Options	3
Disabling Legacy Naming Model	4
Impact on LVM commands after the Disablement of Legacy Naming Model	4
Impact of Native Multi-Pathing on LVM	4
LVM Alternate Links (PVlinks)	5
Restoring LVM's Legacy Alternate Link Behavior	5
LVM Online Disk Replacement (OLR)	6
Volume Group DSF Migration	8
Migration Alternatives	8
Migration of a Clustered Volume Group	9
Migration using <i>vgdsf</i>	9
Migration Procedure	10
A Typical Migration Operation	12
Glossary	14
For More Information	15
Appendix : Using <i>vgdsf</i>	16
Call to Action	18

Abstract

This white paper discusses the migration of LVM volume group configurations from legacy to the agile naming model. The agile naming model is a new representation for mass storage devices, introduced in HP-UX 11i v3. Legacy naming model is the naming model used in releases prior to HP-UX 11i v3. This document is intended for system administrators or operators who have experience with both HP-UX and LVM. For additional information on the mass storage device naming in HP-UX 11i v3, see the [For More Information](#) section for the white paper *HP-UX 11i v3 Mass Storage Device Naming*.

Legacy and Agile Naming Models

In versions prior to HP-UX 11i v3, the Device Special Files (DSFs) of mass storage devices were bound to the path of the device. This representation of mass storage devices still exists in HP-UX 11i v3, and is referred as the **legacy view**. HP-UX 11i v3 introduces a new representation for mass storage devices called the **agile view**. In this representation, the DSFs of mass storage devices are bound to the device World Wide Identifier (WWID) making them **persistent** (that is, each LUN is uniquely identified by a single DSF, irrespective of the number of lunpaths and independent of any physical path changes). The agile view coexists with the legacy view and both types of DSFs can be used to access a given mass storage device independently and simultaneously. For additional information on the next generation mass storage subsystem in HP-UX 11i v3, see the [For More Information](#) section for the white paper *The Next Generation Mass Storage Stack*.

Typical legacy and persistent DSFs for a LUN are shown below:

Legacy DSF: `/dev/dsk/c8t1d3` and `/dev/rdisk/c8t1d3`
Persistent DSF: `/dev/disk/disk66` and `/dev/rdisk/disk66`

LVM support for Dual Naming Models

In the HP-UX 11i v3 release, LVM supports both legacy and persistent DSFs. A volume group can be configured using both styles of DSFs. On a system, some volume groups can use legacy DSFs, some can use persistent DSFs and the rest a mix of legacy and persistent DSFs. Supporting a mixed mode configuration allows a volume group that currently uses one DSF naming model to be extended (by `vgextend`) with the same or another physical volume using a different DSF naming model. Such operations result in a mixed mode volume group and facilitate a phased DSF migration from legacy to agile naming model.

LVM commands `pvccreate`, `vgcreate` and `vgextend` are backward compatible, and work transparently with both legacy and persistent DSFs. New options have been added to `vgimport` and `vgscan` to select the persistent DSF naming model.

HP highly recommends using persistent DSFs for LVM configurations.

All new volume groups should be configured with persistent DSFs only. Existing volume groups configured with legacy DSFs should be migrated (re-configured) to use corresponding persistent DSFs. The migration of volume groups from legacy DSF naming model to persistent DSF naming model can be achieved in several ways. A script named `vgdsf` is also made available to assist in the DSF migration process.

Naming Model specific Commands/Options

New options have been added to `vgimport` (option `-N`) and `vgscan` (options `-N` and `-B`), to specify the persistent naming model. Given below is a summary of these options, and for details, refer to the respective command man pages.

<code>vgimport -N</code>	Configure the volume group using persistent DSFs. In the absence of the <code>-N</code> option, legacy DSFs are used. This option can only be used together with the scan option, <code>-s</code> .
<code>vgscan -N</code>	Recover <code>/etc/lvmtab</code> file using persistent DSFs, with the following exception: Active volume groups configured with legacy DSFs. In this case, <code>vgscan</code> populates the <code>/etc/lvmtab</code> file using legacy DSFs.
<code>vgscan -B</code>	Recover <code>/etc/lvmtab</code> file using both persistent and legacy DSFs. This option can be used

to migrate a volume group configured with legacy DSFs to use corresponding persistent DSFs.

Disabling Legacy Naming Model

The `rmsf -L` command disables the legacy naming model by removing all legacy DSFs and related configuration information. For additional information on disabling legacy naming model in HP-UX 11i v3, see the [For More Information](#) section for the white paper *HP-UX 11i v3 Persistent DSF Migration Guide*. All LVM volume groups **must be migrated** to the agile naming model before disabling the legacy naming model. If any active volume group is still configured with legacy DSFs, the `rmsf -L` command fails with a message similar to the following:

WARNINGS

Affected Volumes

Volume manager configuration on an affected device

/dev/vg_name

(Unused volume. All paths to one or more PVs configured under this volume are affected)

PVs under this volume are: /dev/dsk/c0t2d0

Note: The `rmsf -L` command succeeds if volume groups configured with legacy DSFs are not activated. HP strongly recommends activating all the LVM volume groups before attempting to disable the legacy naming model.

Impact on LVM commands after the Disablement of Legacy Naming Model

When the legacy naming model is disabled, `vgscan` and `vgimport` commands without a `-N` option, fail with the following message:

The legacy naming model has been disabled on the system.

Try with the -N option.

The `vgscan` command when invoked with `-a` or `-B` option, fails with the following message:

The legacy naming model has been disabled on the system.

The option being used needs the legacy naming model enabled.

If the legacy naming model is disabled, when there exists inactive volume groups configured with legacy DSFs, (`/etc/lvmtab` file has legacy DSFs) any subsequent invocation of `pvcreate`, `vgimport`, `vgcreate`, `pvremove`, and `vgchgid` fails with the following message, irrespective of the physical volume/volume group on which the command is invoked:

"The legacy naming model is disabled, but the /etc/lvmtab file still contains legacy devices such as "/dev/dsk/c1t2d3". Command will not succeed until either the legacy device is removed from the /etc/lvmtab file or the legacy mode is enabled using `insf(1M) -L`."

Impact of Native Multi-Pathing on LVM

The mass storage stack in HP-UX 11i v3 provides native multi-pathing functionality, distributing I/O requests transparently across all available lunpaths to the LUN. For additional information on native multi-pathing for mass storage, see the [For More Information](#) section for the white paper *HP-UX 11i v3 Native Multi-Pathing for Mass*

Storage. This is the default behavior in the case of active/active devices, irrespective of the naming model used. Some active/passive devices may be supported by the mass storage stack but will be managed by the respective vendor delivered plug-in modules. For additional information on writing active/passive device plug-ins in HP-UX 11i v3, see the [For More Information](#) section for the white paper *HP-UX 11i v3 Writing Active/Passive Switch Plug-Ins*.

Due to the native multi-pathing functionality of the mass storage stack, the behavior of LVM alternate link and LVM online disk replacement functionalities may differ when compared with previous releases.

LVM Alternate Links (PVlinks)

LVM supports alternate links to a device to allow continued access to the device if the primary link fails. In previous HP-UX releases, (that is, HP-UX 11i v2 or earlier) LVM's multi-pathing solution increases the data availability but does not allow multiple paths to be used simultaneously. LVM continues to support this functionality even in HP-UX 11i v3, but the behavior may be very different based on how LVM and the mass storage stacks are configured.

With the introduction of mass storage stack's native multi-pathing functionality, it is **no longer required** nor recommended to configure LVM alternate links.

Listed below are the alternate link behavioral differences that could be seen, when compared with previous releases:

- Regardless of whether LVM alternate links are configured or not, the mass storage stack processes the I/O operations to the device using all available paths. That is, LVM selects the LUN and the mass storage stack selects the lunpath to be used.
- If LVM alternate links are configured, no link switching ever happens unless the mass storage stack's native multi-pathing feature is disabled using the *scsimgr* command and only legacy DSFs are used in the volume group configuration. For additional information on *scsimgr* command capabilities in HP-UX 11i v3, see the [For More Information](#) section for the white paper *scsimgr SCSI Management and Diagnostics Utility on HP-UX 11i v3*.
- Path switch options of the *pvchange* (-s and -S options) command do not switch links or stop I/O operations, as they did in earlier releases. In order to disable I/O operations on a given path, use the *scsimgr* command (*scsimgr [-f] disable Lunpath*).
- While LVM alternate link functionality supports up to 8 paths to a physical volume, mass storage stack's native multi-pathing feature can support up to 32 lunpaths.

Restoring LVM's Legacy Alternate Link Behavior

If legacy LVM alternate link functionality is preferred over the mass storage stack's native multi-path functionality, the native multi-pathing feature can be disabled using *scsimgr* command. This can be accomplished by setting the device attribute *leg_mpath_enable* to false. This attribute is available on both global (system wide) and on a per-LUN basis. The value of per-LUN tunable overrides the global value. As this feature is available only on legacy DSFs, LVM volume groups must be configured using legacy DSFs.

Note: Multi-pathing is not affected by this attribute when persistent DSFs are used.

The following example shows how to enable legacy LVM alternate link functionality in a backward compatible manner. In this example, the volume group *vg01* is configured with a single persistent DSF.

1. Get a list of DSFs configured in the volume group:

```
# vgdisplay -v -F vg01 | grep pv_name
pv_name=/dev/disk/disk5:pv_status=available:total_pe=243:free_pe=242:autoswitch=On
```

2. Find legacy DSFs corresponding to the persistent DSF:

```
# ioscan -m dsf /dev/disk/disk5
Persistent DSF      Legacy DSF(s)
=====
/dev/disk/disk5    /dev/dsk/c2t1d0
                   /dev/dsk/c3t1d0
```

3. Add legacy DSFs as alternate links to the persistent DSF:

```
# vgextend vg01 /dev/dsk/c2t1d0 /dev/dsk/c3t1d0
```

4. Remove the persistent DSF from the volume group:

```
# vgreduce vg01 /dev/disk/disk5
```

5. Disable native multi-pathing through the legacy DSFs:

scsimgr command can be run in any of the two ways, listed below:

- Disable native multi-pathing for a specific LUN in a non-persistent way:

```
# scsimgr set_attr -D /dev/rdisk/disk5 -a leg_mpath_enable=false
Value of attribute leg_mpath_enable set successfully
```
- Disable native multi-pathing globally for all LUNs in the system, in a non-persistent way:

```
# scsimgr set_attr -a leg_mpath_enable=false
Value of attribute leg_mpath_enable set successfully
```

6. Verify the resulting configuration:

The *scsimgr get_attr* command should show the current setting of *leg_mpath_enable* attribute value set to false, for all associated persistent DSFs in the volume group.

```
# scsimgr get_attr -D /dev/rdisk/disk5 -a leg_mpath_enable

SCSI ATTRIBUTES FOR LUN : /dev/rdisk/disk164

name = leg_mpath_enable
current = false
default = false
saved =
```

Note:

- If the *scsimgr* command is invoked with *save_attr* command instead of *set_attr*, the corresponding change will be persistent (behavior preserved across reboots).
- To fully understand the mass storage attributes and the consequences of changing them, refer to *scsimgr(1M)* manual page.

LVM Online Disk Replacement (OLR)

In previous HP-UX releases (that is, HP-UX 11i v2 or earlier), a specific path or all paths to a physical volume can be detached (disabled) using `-n` and `-N` options of the `pvchange` command, respectively. The detached path or the physical volume will not be accessed by LVM, but remain configured in the volume group. For additional information on replacing disks configured under LVM, see the [For More Information](#) section for white papers *LVM Online Disk Replacement (LVM OLR)* and *When Good Disks Go Bad: Dealing with Disk Failures under LVM*.

In HP-UX 11i v3, detaching an entire physical volume using the `pvchange -a N` command, in order to perform an online disk replacement is still supported. The behavior is the same for both legacy and persistent DSFs and is compatible with previous releases.

However, unless native multi-pathing has been disabled and only legacy DSFs are configured for the volume group, the `pvchange -a n` command will not stop I/O operations for that path as it did in earlier releases. Instead, use the `scsimgr` command.

Volume Group DSF Migration

In HP-UX 11i v3, HP recommends the use of persistent DSFs over legacy, in LVM volume group configurations. A migration operation is required for existing volume groups already configured with legacy DSFs. Migration involves reconfiguration of existing volume groups with corresponding persistent DSFs, in place of legacy DSFs.

DSF migration should be performed one volume group at a time, and should be completed for all volume groups. All volume groups configured with legacy DSFs can be migrated, irrespective of the volume group state (activated/de-activated, shared/standalone, with/without alternate links). Volume groups can be migrated using any of the following methods:

Note: The migration changes only the DSFs. Volume group information and layout is not altered in any way.

Note: `/etc/lvmpvg` is an optional ASCII file (refer to `lvmpvg(4)` manual page) that stores the physical volume group (PVG) information. `/etc/lvmtab` is a binary file that stores the association between physical volumes and the corresponding volume group. Both these files support legacy and persistent DSFs and must be updated as part of the migration operation.

Migration Alternatives

All of the migration alternatives listed below automatically update the `/etc/lvmtab` file to reflect the newer configuration and no explicit user action is required. However, some alternatives require that the PVG information in `/etc/lvmpvg` file be updated manually. Please refer to the information on `/etc/lvmpvg` file, listed in each of the alternatives.

Any one of the below alternatives can be used to migrate a volume group. HP recommends using the `vgdsf` script ([alternative 3](#)) for migration.

Note: Each alternative updates one volume group at a time. The steps must be repeated in the alternative chosen for all volume groups to be migrated.

1. Backup the `/etc/lvmpvg` file, if it exists. Deactivate and export the volume group. Re-import the same volume group using persistent DSFs in place of legacy DSFs. Persistent DSFs, corresponding to legacy DSFs, can be found using `ioscan -m dsf <legacy_dsf>` command. Restore the backed-up `/etc/lvmpvg` file and replace legacy DSFs listed under the volume group being migrated with corresponding persistent DSFs (In case of multi-path devices, a single persistent DSF can correspond to multiple legacy DSFs). The disadvantage of this method is the necessity of volume group deactivation.
2. Run `vgscan -B` command followed by volume group activation.

For each of the physical volumes configured in the volume group, `vgscan -B` command populates the `/etc/lvmtab` file with both persistent and legacy DSFs. However, `/etc/lvmtab` file supports a maximum of 8 paths per physical volume. Therefore, if a physical volume has more than 8 paths configured, only 7 are retained in order to allow space for the addition of persistent DSF.

Note: The `vgscan` command will not update existing volume group entries in `/etc/lvmtab` file unless the `-f` option is used. The `-f` option restricts the operation to the specified volume group (overwrite existing volume group entries in `/etc/lvmtab`). If the `/etc/lvmtab` file is moved to a new location before running the `vgscan -B` command, full effect of the command is taken on all configured volume groups and `/etc/lvmtab` file gets re-created.

Following the invocation of `vgscan -B` command, use `vgchange` to reactivate the corresponding volume group. This reconfigures the volume group with the DSFs (both legacy and persistent DSFs corresponding to each of the physical volume) in the `/etc/lvmtab` file. Legacy DSFs can later be removed using `vgreduce`, leaving behind corresponding persistent DSFs. This facilitates a phased DSF migration from legacy to agile naming model, while the volume group continues to remain active.

Backup the `/etc/lvmpvg` file, if it exists. After the `vgreduce` operations, restore the backed-up `/etc/lvmpvg` file and replace legacy DSFs listed under the volume group being migrated with corresponding persistent DSFs (In case of multi-path devices, a single persistent DSF can correspond to multiple legacy DSFs). Persistent DSFs, corresponding to legacy DSFs, can be found using `ioscan -m dsf <legacy_dsf>` command.

3. A shell script `/usr/contrib/bin/vgdsf`, can be used to perform the DSF migration of a given volume group. `vgdsf` uses existing LVM commands - `vgdisplay`, `pvdisplay`, `vgextend` and `vgreduce`, in addition to the `ioscan` command to perform the DSF migration. It automatically updates both the `/etc/lvmtab` and `/etc/lvmpvg` files to reflect the newer configuration and no additional user action is required. The use of the `vgdsf` script simplifies the DSF migration operation and the volume group can remain active and in use.
4. The `vgdsf` script is a wrapper script around existing LVM and I/O commands. Hence, the same functionality can also be achieved by the direct use of these commands. The operation and the steps can be followed by viewing the `vgdsf` script (briefly explained in the later sections of this paper). This alternative provides a user, control over the DSF migration operation. (Remember to use `vgextend` with `-g` option in order to preserve the PVG configuration.)

Note: The above methods (except for the `vgdsf` script in [alternative 3](#)) can also be used to migrate volume group configurations back to the legacy naming model.

Migration of a Clustered Volume Group

To migrate volume groups used in a Serviceguard cluster, each volume group must be either inactive, or active on only one node in the cluster. The volume group must be migrated on each node. It is possible, though not a best practice, to use legacy DSFs on some nodes after migrating to agile addressing on others; this allows you to migrate different nodes at different times, if necessary.

On nodes where the volume group is not active, alternatives [1](#) and [2](#) can be used to migrate the volume group (shared and non-shared).

Additional steps need to be taken before `vgdsf` can be used to migrate volume groups shared by Serviceguard cluster. The activation mode must be switched to "exclusive" on the node where the volume group is active. `vgdsf` can then be run on this node to migrate the volume group. Then, method [1](#) or [2](#) should be used to migrate the volume group on the remaining nodes whilst still deactivated. Then the volume group is switched back to "shared" mode on this node and reactivated on all nodes.

Migration using `vgdsf`

To simplify migration of volume groups, a user script – `vgdsf`, is made available.

The `vgdsf` script performs the following tasks. Read the script for a more complete understanding.

1. Extend the volume group with corresponding persistent DSFs:

- Identify all legacy DSFs configured in the volume group, using `vgdisplay -v`. Ignore alternate links.
 - For each legacy DSF found, find its corresponding persistent DSF, using `ioscan -m dsf <legacy_dsf>`. If the persistent DSF is already part of the volume group configuration, continue for this particular DSF, with the step 2 - "Reduce the volume group of all legacy DSFs".
 - If a physical volume is configured with 8 paths¹ (alternate links), remove one using `vgreduce`. This makes room for the additional persistent DSF.
 - Extend the volume group to add the persistent DSF, keeping it in the same physical volume group (PVG) configuration, if any (`/etc/lvmpvg` file gets automatically updated).
2. Reduce the volume group of all legacy DSFs:
 - Identify all legacy DSFs configured in the volume group, using `vgdisplay -v`.
 - For each of the legacy DSF that also has a corresponding persistent DSF configured, reduce the legacy DSF from the volume group. If the corresponding persistent DSF is not configured in the volume group, a message is displayed on screen and the command continues with other DSFs.
 3. Back up the resulting configuration:
 - Back up the volume group configuration, using `vgcfgbackup` command.

Migration Procedure

Explained below are the steps to be performed for migration (using the `vgdsf` script).

1. Activate all configured volume groups in the system and identify the ones requiring a migration, using `vgchange -a y`
2. For reference, make a note of the volume group configuration. Save the outputs of LVM display commands – `vgdisplay`, `pvddisplay` and `lvdisplay`, invoked in verbose mode.
3. Ensure all physical volumes configured in the volume group are online using `ioscan -P health` command.
4. Run `/usr/contrib/bin/vgdsf -c <vg_name>` to migrate each volume group. Ensure that no failures are reported.
5. Verify the resulting volume group configuration using LVM display commands. Use `ioscan -m dsf <dsf_name>` command to validate the DSF mappings.

Note: `vgdsf` script must be invoked on all configured volume groups in the system and can easily be automated. Given below is a sample script.

```
#!/usr/bin/ksh
#
# This script converts legacy DSFs to persistent DSFs in
# all active volume groups.
#
VGS=`/usr/sbin/vgdisplay -F | grep vg_name | cut -d: -f1 | cut -d= -f2`
for VG in $VGS
do
  if vgdisplay -F $VG | grep -q "vg_status=available"; then
    /usr/contrib/bin/vgdsf -c $VG
```

¹ LVM supports a maximum of 8 paths to a physical volume.

```
else
  echo Volume Group $VG is not active and cannot be migrated.
fi
done
```

A Typical Migration Operation

1. Identify volume groups for migration:

- a) Activate deactivated volume groups:

In this example, volume group `vg01` is not active. Activate it.

```
# vgdisplay -F | grep -E "vg_status=deactivated"
vg_name=/dev/vg01:vg_status=deactivated
```

```
# vgchange -a y vg01
Activated volume group
Volume group "vg01" has been successfully changed.
```

- b) Identify volume groups configured with legacy DSFs. Only the ones configured with legacy DSFs must be migrated.

In this example, volume group `vg03` does not require a migration. Only the volume group `vg01` needs to be migrated as it is configured with legacy DSFs.

```
# vgdisplay -Fv /dev/vg03 | grep -E "pv_name=/dev/dsk/c[0-9]+t[0-9]+d[0-9]"
```

```
# vgdisplay -Fv /dev/vg01 | grep -E "pv_name=/dev/dsk/c[0-9]+t[0-9]+d[0-9]"
pv_name=/dev/dsk/c8t0d1,/dev/dsk/c10t0d1,/dev/dsk/c19t0d1,/dev/dsk/c21t0d1:pv_status=available:total_pe=1279:free_pe=1279:autoswitch=On
pv_name=/dev/dsk/c8t0d4:pv_status=available:total_pe=1279:free_pe=1279:autoswitch=On
pv_name=/dev/dsk/c8t0d5:pv_status=available:total_pe=1279:free_pe=1279:autoswitch=On
pv_name=/dev/dsk/c8t0d6:pv_status=available:total_pe=1279:free_pe=1279:autoswitch=On
```

- c) Ensure all physical volumes configured in the volume group are online.

```
# iocan -P health -C /dev/disk/disk59
Class   I H/W Path health
=====
disk   153 64000/0xfa00/0x40 online
```

2. Make a note of the volume group configuration:

For reference, make a complete note of volume group configuration. This information helps to revert the volume group to its original configuration, if required. LVM display commands like `vgdisplay`, `pvdisplay` and `lvdisplay` can be invoked in verbose mode to gather the required information. A backup of `/etc/lvmtab` and `/etc/lvmpvg` files should also be taken using `cp` command.

```
# cp /etc/lvmtab /etc/lvmtab.backup
# cp /etc/lvmpvg /etc/lvmpvg.backup
```

3. Run `vgdsf` to migrate volume group `vg01`:

```
# /usr/contrib/bin/vgdsf -c /dev/vg01
Converting legacy DSFs to persistent DSFs in VG /dev/vg01
Persistent DSF /dev/disk/disk56 added to VG /dev/vg01
Persistent DSF /dev/disk/disk59 added to VG /dev/vg01
```

Persistent DSF /dev/disk/disk60 added to VG /dev/vg01
Persistent DSF /dev/disk/disk61 added to VG /dev/vg01
Legacy DSF /dev/dsk/c10t0d1 removed from VG /dev/vg01
Legacy DSF /dev/dsk/c19t0d1 removed from VG /dev/vg01
Legacy DSF /dev/dsk/c21t0d1 removed from VG /dev/vg01
Legacy DSF /dev/dsk/c8t0d1 removed from VG /dev/vg01
Legacy DSF /dev/dsk/c8t0d4 removed from VG /dev/vg01
Legacy DSF /dev/dsk/c8t0d5 removed from VG /dev/vg01
Legacy DSF /dev/dsk/c8t0d6 removed from VG /dev/vg01
Volume Group configuration for /dev/vg01 has been saved in /etc/lvmconf/vg01.conf

4. **Verify the resulting volume group configuration:**

Make a note of the resulting volume group configuration using LVM display commands - `vgdisplay`, `lvdisplay` and `pvdisplay`. DSF mapping can be checked using `ioscan -m dsf <dsf_file>`.

Glossary

Active/Active device

An active/active device has two or more ports to the same LUN, both of which can be accessed simultaneously.

Active/Passive device

Also known as asymmetric active/active device. Though the same LUN has multiple ports, unlike an active/active device, only the primary port should be used for LUN access at any time.

Agile Addressing

The ability to address a LUN with the same device special file regardless of the physical location of the LUN or the number of paths leading to it. The DSF for a LUN remains the same even if the LUN is moved from one HBA to another, moved from one switch or hub port to another, presented using a different target port to the host, or configured with multiple hardware paths. This is also referred to as persistent LUN binding.

Agile View

The representation of LUNs using lunpath hardware paths, LUN hardware paths, and persistent DSFs, introduced in HP-UX 11i v3.

Legacy DSF

A DSF with the hardware path information such as SCSI bus, target, and LUN embedded in the file's minor name and file name, such as `/dev/dsk/c2t3d4`.

Legacy Naming Model

The representation of legacy hardware paths and legacy DSFs as in releases prior to HP-UX 11i v3.

DSF

Device Special File. A file associated with an I/O device. DSFs are read and written the same as ordinary files, but requests to read, write or control are sent to the associated device.

Hardware Path

A series of numbers representing the physical or virtualized location of a device. The path is a sequence of I/O addresses that share a hierarchical relationship. The address elements may not correspond to physical hardware addresses, and may represent only a "handle" to a device rather than a physical path to it.

LUN

A SCSI logical unit. This refers to an end storage device such as a disk, tape, floppy, or CD. This is the logical unit itself and does not represent the path to the logical unit.

Lunpath

The physical hardware path leading to a SCSI logical unit. A SCSI LUN can have more than one lunpath.

Multi-Pathing

The detection, correlation, and coordinated usage of multiple hardware paths leading to the same LUN.

Persistent DSF

A DSF conforming to the naming model introduced in HP-UX 11i v3 to support agile addressing. The device file name contains an instance number, such as `/dev/disk/disk10`, and the minor number has no hardware path information.

For More Information

To learn more about the agile view, mass storage, DSF migration and LVM OLR on HP-UX, see the following documents on the HP documentation website:

<http://docs.hp.com/>

- The Next Generation Mass Storage Stack
- HP-UX 11i v3 Mass Storage Device Naming
- HP-UX 11i v3 Persistent DSF Migration Guide
- HP-UX 11i v3 Native Multi-Pathing for Mass Storage
- HP-UX 11i v3 writing Active/Passive Switch Plug-Ins
- scsimgr SCSI Management and Diagnostics Utility on HP-UX 11i v3
- LVM Online Disk Replacement (LVM OLR)
- When Good Disks Go Bad: Dealing with Disk Failures under LVM

Most of these documents are available in the Network and Systems Management section of <http://docs.hp.com/>, under [Storage Area Management](http://docs.hp.com/en/netsys.html#Storage%20Area%20Management) (<http://docs.hp.com/en/netsys.html#Storage%20Area%20Management>).

Appendix : Using vgdsf

NAME

vgdsf – Script to migrate a LVM volume group from legacy to agile naming model.

SYNOPSIS

```
/usr/contrib/bin/vgdsf {-a | -d | -c} vg_name
```

DESCRIPTION

Migrate the volume group from legacy to the persistent DSF naming model.

The volume group `vg_name` should be active and can be in use during the migration operation. All legacy DSFs in the volume group `vg_name`, will be replaced by their corresponding persistent DSFs. This script is aware of physical volume groups (PVG) and maintains the same PVG configuration even after the migration. Both `/etc/lvmtab` and `/etc/lvmpvg` configuration files are automatically updated. An automatic backup of the volume group is performed after the migration is complete.

OPTIONS AND ARGUMENTS

`vgdsf` recognizes the following options and arguments:

- a For each of the legacy DSF configured in the volume group `vg_name`, a corresponding persistent DSF will be added to the configuration, as an alternate link. Legacy DSFs continue to remain configured as is. If a physical volume has more than seven alternate links, one of them will be removed to make room for its corresponding persistent DSF. This option should precede the invocation of `-d` option.
- d All legacy DSFs configured in the volume group `vg_name` are removed if their corresponding persistent DSFs are also found configured. Legacy DSFs in the volume group without corresponding persistent DSFs are not removed. This option should follow the invocation of `-a` option.
- c Do a complete migration of the volume group from legacy to agile naming model. The functionality is equivalent to the invocation of `vgdsf -a` followed by `vgdsf -d`.

`vg_name` volume group to be migrated.

NOTE

Use of `-a` and `-d` options is generally not required.

EXAMPLES

Migrate volume group `vg_name`, configured using legacy DSFs, to persistent DSFs.

```
/usr/contrib/bin/vgdsf -c /dev/vg_name
```

An equivalent of the above can also be achieved as below:

```
/usr/contrib/bin/vgdsf -a /dev/vg_name
```



```
/usr/contrib/bin/vgdsf -d /dev/vg_name
```

SEE ALSO

intro(7), lvm(7), vgcreate(1M), vgextend(1M), vgdisplay(1M), pvdisplay(1M)

Call to Action

HP welcomes your input. Please give us comments about this white paper, or suggestions for mass storage or related documentation, through our technical documentation feedback website:

<http://docs.hp.com/en/feedback.html>

© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Updated March 2007 (to add migration of /etc/lmpvg)