

# HP-UX Logical Volume Manager and MirrorDisk/UX Release Notes

## HP-UX 11i v3 March 2013 Release (B.11.31)

HP Part Number: 5900-3002  
Published: May 2013  
Edition: 12.0



© Copyright 2012, 2013 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license. 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. UNIX is a registered trademark of The Open Group.

---

# Contents

Logical Volume Manager and MirrorDisk/UX Release Notes.....	4
About this document.....	4
LVM and MirrorDisk/UX overview.....	4
Overview of changes.....	4
New and changed features in this release.....	4
LVM boot support for disks greater than 2 TB in size:.....	4
Fixed issues in this version.....	6
Known problem and fix for Asian command output.....	7
Known issues and limitations.....	7
Known issues.....	7
LVM cannot handle IOs at and beyond 4TB offset on snapshots.....	7
pvmmove is coredumping.....	7
lvmput daemon log level is broken.....	7
Limitations.....	8
Installation requirements.....	8
Required hardware.....	9
Required software.....	9
Required patches.....	9
Mass Storage Critical Resource Analysis.....	9
mkfs and mkboot commands.....	9
HP Serviceguard.....	9
VxVM commands.....	10
Required disk space.....	10
Other requirements.....	10
HP Insight Remote Support Software.....	10
Compatibility issues.....	10
Version 2.x volume groups.....	11
Moving from HP-UX 11i v2 to HP-UX 11i v3.....	11
Moving volume groups from HP-UX 11i v3 to previous HP-UX releases.....	12
Sharing volume groups in an HP Serviceguard cluster.....	13
Related documentation.....	13
Software Availability in Native Languages.....	13

# Logical Volume Manager and MirrorDisk/UX Release Notes

## About this document

This document provides information about the Logical Volume Manager (LVM) and MirrorDisk/UX products in the March 2013 release of HP-UX 11i v3

## LVM and MirrorDisk/UX overview

Logical Volume Manager (bundle BaseLVM) is the HP-UX default Volume Manager. It provides user with flexibility in configuring and managing mass storage resources. In HP-UX 11i v3, the LVM kernel and commands are bundled with the core HP-UX product.

MirrorDisk/UX (bundle B2491BA) is an optionally purchased HP-UX product to enable LVM mirroring functionality.

## Overview of changes

The initial HP-UX 11i v3 release of LVM and MirrorDisk/UX was integrated with the new mass storage stack, delivering significant performance, scalability, availability, and usability enhancements. LVM was enhanced to support larger logical volumes, temporary quiescing of volume groups, and striping with mirroring. Volume group availability was improved: resizing a LUN and modifying volume group characteristics no longer required the volume group to be recreated, and replacing a disk could be done online.

The September 2010 release of LVM and MirrorDisk/UX provides enhancements to logical volume snapshots. In addition, this release provides support for cluster Device Special Files (cDSF) and an enhancement to the `vgchange` command when operating in a cluster environment.

The February 2010 release of LVM and MirrorDisk/UX is a defect fix only release; there are no new features.

This March 2013 release of LVM and MirrorDisk/UX provides defect fixes as described in [“Fixed issues in this version”](#) (page 6). In addition, it provides an enhancement to support LVM boot disks of size greater than 2 TB.

## New and changed features in this release

The following LVM feature is included in the March 2013 release of HP-UX 11i v3:

### LVM boot support for disks greater than 2 TB in size:

Beginning March 2013 release, the maximum boot disk size supported with LVM boot layout is enhanced from 2 Terabytes (TB) to 16 TB on HP Integrity systems. You can now install HP-UX on disks up to 16 TB in size.

This new boot enhancement is included to LVM starting with Volume group version 2.2. With LVM boot layout, you can create System Logical volumes (Root/Boot/Swap/Dump Logical Volumes) up to the size of 16 TB. However, the maximum system volume sizes are subject to the limits applied by Ignite-UX during the HP-UX cold installation.

For information about the limits, see the Ignite-UX release notes and administration guide available at <http://www.hp.com/go/ignite-ux> for more details on such limits.

**Table 1 Limits for various releases and platforms**

HP-UX 11i v3 OE Update	Maximum supported boot disk size	
	HP Integrity systems	HP 9000 systems
Prior to March 2011	1 TB	2 TB

**Table 1 Limits for various releases and platforms** *(continued)*

March 2011 and later	2 TB	2 TB
March 2013 and later	16 TB	2 TB

**NOTE:** This new boot capability is not available on HP 9000 (PA RISC) platforms. Also note that this feature is added only for LVM boot layout and not available with VxVM and whole disk (no volume manager) boot layouts.

Currently the feature is supported with the following I/O Cards:

- SAS HBAs: 51378-B21(P711m), AM311A(P411), AM312A(P812), Internal HBA P410i
- Fiber Channel HBAs: 403619-B21, 403621-B21, 451871-B21, 456972-B21, AD193A, AD194A, AD221A, AD222A, AD393A, AH400A, AH401A, AH402A, AH403A, AT094A

For the most recent list of cards that support the greater than 2 TB feature, see the IO cards support matrixes available at <http://www.hp.com/go/hpux-iocards-docs>

This new boot capability requires the following software, all of which is always-installed with all March 2013 (or later) Operating Environments (OEs).

- BaselvM B.11.31.1303
- PHKL\_43192 (dump cumulative patch)
- PHKL\_42916 (SCSI cumulative I/O patch)
- PHKL\_43191 (boot loader cumulative patch)
- PHKL\_43266 (mdep efi\_utils cumulative patch)
- PHKL\_43275 (USB devices, setboot cumulative patch)
- PHCO\_43237 (swapinfo(1M) cumulative patch)
- PHCO\_43400 (swapon(1M) cumulative patch)

These patches are delivered with both HWEnable11i and FEATURE11i bundles.

For more details see *LVM System Administration Guide* available at <http://www.hp.com/go/hpux-LVM-VxVM-docs>

## Fixed issues in this version

The following table lists the LVM and MirrorDisk/UX defects fixed in the March 2013 release of HP-UX 11i v3.

**Table 2 LVM Fixes in HP-UX 11i v3 March 2013 Release**

Defect ID	Description
QXCR1001211881	vgmodify -a -S when run to increase the size of a volume group results in a subset of physical volumes getting modified without modifying the volume group size itself if one or more PV's in the volume group fail modification.
QXCR1001218677	LVM IOs initiated on a failed device during the window of internal system clock timer wrap-around may fail back pre-maturely before the non-default LV timeout (configured through lvchange -t) is reached. The timer wrap-around happens on every ~497 days w.r.t. system up time.
QXCR1001214964	<p>When thin volume reclaim is tried on a VxFS file system over LVM Logical volume combination by mistake (such combination is unsupported for reclaim) by running 'fsadm -R &lt;FS_mount_point&gt;' command, it may result in the following panic:</p> <pre>panic: Fault when executing in kernel mode Stack Trace:   IP                               Function Name 0xe0000000021f6420                bad_kern_reference+0xa0 0xe000000000707ea0                \$cold_pfault+0x530 0xe000000000eccc00                vm_hdlr+0x12d0 0xe000000001ded780                bubbledown+0x0 0xe00000004e6c9df01               lvmp_start_str2+0x151 0xe00000004e6c9ea10               lvmp_strategy2+0x790 0xe00000004e6dc3350               lvmp_strategy+0x1670 0xe000000000d80be0               io_invoke_devsw+0x190 0xe000000000beb140               vx_dev_strategy+0xc0 0xe000000000a20650               vx_dummy_fsmv_strategy+0x30 0xe000000000c4a8a0               vx_ts_reclaim+0x2c0 0xe0000000022d09e0               \$cold_vx_aiocntl_common+0xf0 0xe000000000886090               vx_aiocntl+0x2d0 0xe000000000b79960               vx_ioctl+0x180 0xe000000000d82af0               vno_ioctl+0x350 0xe000000000db8450               ioctl+0x410 0xe000000000e8cde0               syscall+0x590 End of Stack Trace</pre>
QXCR1001222652	<p>vgimport 2.X fails with the following error message when the two copies of the VGDA contain different data but the directory IDs are the same :</p> <pre>\$ vgimport vggm3 /dev/disk/disk65 /dev/disk/disk60 vgimport: The Physical Volumes specified on the command line do not belong to the same Volume Group.</pre>
QXCR1001237268	<p>lvextend fails to increase the number of mirrors of a PVG strict LV despite free Physical Extents being available satisfying the policy.</p> <p>lvextend fails with the following error message:</p> <pre>\$ lvextend -m 1 /dev/vg01/lv_PVG lvextend: The LVM device driver could not extend the logical volume /dev/vg01/lv_PVG".</pre>
QXCR1001217016	'vgchange -c n -S n' returns a non-zero exit value even though the command was successful.
QXCR1001161003	Invoking 'vgchange -c n' on a cluster node when the VG is activated in shared mode on the other nodes would clear the cluster ID of the VG. This causes further activation of the VG to fail in any nodes.
QXCR1001205782	/etc/lvmtab_p file in the clients in a shared VG environment end up with persistent dsf even when cdsf was used with vgextend command in the server.
QXCR1001185604	lvmpud daemon uses LOG_ERR type of error messages for daemon start-up and daemon exit messages.

## Known problem and fix for Asian command output

On LVM version B.11.31.1103, when the `LANG` environment variable is set to any of the nine Asian locales below, output from LVM commands include extraneous “carriage return” (^M) characters at the end of message strings. These extraneous characters are not shown on output directed to the console. However, if the LVM command output is redirected to a file, these control characters are visible when the file is opened with a text editor, such as `vi` or `emacs`. Also, in the command line interface, command output in table formats might not be aligned and table labels might not display for some commands, because of these extraneous ^M characters.

This problem is limited to these nine Asian locales: `ja_JP.SJIS`, `ja_JP.eucJP`, `ja_JP.utf8`, `ko_KR.eucKR`, `ko_KR.utf8`, `zh_CN.hp15CN`, `zh_CN.gb18030`, `zh_HK.hkbig5`, and `zh_TW.big5`.

The fix for this problem is available with LVM revision “B.11.31.1103.01” (March 2011 Web Release) which can be downloaded from the LVM software download page: <http://www.hp.com/go/tsvm>

## Known issues and limitations

This section provides a list of known issues applicable to current and previous releases. It also includes limitations as known to HP at time of publication. If workarounds are available, they are included.

### Known issues

#### LVM cannot handle IOs at and beyond 4TB offset on snapshots

Defect ID	QXCR1001256843
Problem	Reading /writing IOs at and beyond 4 TB offset on a space-efficient snapshot will always return EIO (error number 5). Even fully-allocated snapshots cannot handle IOs at and beyond 4TB offsets. Though the actual IO will return success on fully-allocated snapshots, internally the data is not read or written accurately.
Severity	Critical
Corrective Action	None

#### pvmove is coredumping

Defect ID	QXCR1001254853
Problem	The <code>pvmove</code> performed to move extents from PVs belonging to same PVG and containing more than 40 PVs in a version 1 VG causes <code>pvmove</code> to dump core.
Severity	Serious
Corrective Action	Moving the destination PV to a different PVG fixes the issue.

#### lvmputd daemon log level is broken

Defect ID	QXCR1001282416
Problem	The <code>lvmputd</code> daemon logs info messages even though it is started without <code>-l</code> or <code>-d</code> option.
Severity	Medium
Corrective Action	Configure the <code>syslog.conf</code> file to redirect INFO level messages to a different file rather than to the default <code>syslog</code> file. Note that when this is done, the INFO level messages from all the daemons (not only from the <code>lvmputd</code> daemon) are redirected to the new file.

Example:

1. Add the `daemon.info` configuration to the `syslog.conf` file:

```
# more /etc/syslog.conf
...
*.info;mail.none      /var/adm/syslog/syslog.log
daemon.info           /var/adm/syslog/daemon.log
*.alert               /dev/console
...
```

2. Touch the `/var/adm/syslog/daemon.log` file:

```
# touch /var/adm/syslog/daemon.log
```

3. Send a HUP signal to `syslogd` to make it re-read the configuration file:

```
# kill -HUP `pgrep syslogd`
```

4. Now, INFO messages from daemons are redirected to the `/var/adm/syslog/daemon.log` file:

```
# more /var/adm/syslog/daemon.log
..
lvmpud[3886]: Read task ticket=283 type=1
lvmpud[3886]: Process task ticket=283 type=1
lvmpud[3886]: Complete task ticket=283 type=1
status=0
..
```

## Limitations

This section provides a list of limitations as known to HP at time of publication.

- Logical volumes that have snapshots associated with them might experience an increase in latencies associated with writes. This is also applicable for writes on (writable) snapshots themselves.
- With the HP-UX 11i v3 September 2010 Update release, when the number of extents remaining in a space-efficient snapshot logical volume's pre-allocated pool falls beyond a certain internally computed threshold, by default, LVM tries to increase the pre-allocated pool size by threshold value. This does not guarantee that snapshot will not become over-commit. For example, if the I/O rate is faster than auto pre-allocation, the snapshot becomes over-commit.

If auto pre-allocation is manually disabled, a message is logged in the `syslog`. This message is displayed to inform you that further unsharing of data between the snapshot and its successors might end up depleting the extents in the pre-allocated pool and lead to the snapshot and its predecessors being marked as inoperative. As soon as you see this message in the `syslog`, you must increase the pre-allocated extent pool size for the snapshot logical volume.

- Automatic increase of pre-allocated extents does not prevent snapshots from becoming inoperative. If the I/O rate is faster than pre-allocation, it is possible that the free extents are picked before pre-allocation happens and the subsequent I/O marks the snapshot as inoperative.
- Under very low memory conditions, the removal of a single snapshot using `lvremove` can hang if data unsharing is required.
- For more limitations on snapshots, see the *Using LVM Logical Volume Snapshots* white paper.

## Installation requirements

This section describes the installation requirements for this release.



## Required hardware

LVM and MirrorDisk/UX have no hardware requirements beyond the requirements of the HP-UX 11i v3 operating system. Both products run on all supported HP 9000 and HP Integrity servers.

## Required software

Because LVM is installed with the HP-UX 11i v3 operating environment, there are no software requirements beyond the requirements of the HP-UX 11i v3 operating system. MirrorDisk/UX requires the BaseLVM™ product.

LVM and MirrorDisk/UX require the following software:

- Software 1
- Software 2
- Software 3

You can download software from the following Web site:

<http://www.hp.com/go/softwaredepot>

Enter <Keyword> into the Search box at the top of the page.

## Required patches

In HP-UX 11i v3, LVM and MirrorDisk/UX do not require any software patches.

For each of the following subsystems, you can download patches from the following website:

<http://www2.itrc.hp.com/service/patch/mainPage.do>

## Mass Storage Critical Resource Analysis

The Mass Storage Critical Resource Analysis (MS CRA) tool checks for a maximum of two mirror copies of a logical volume. However, Version 2.0 and higher volume groups allow up to five mirror copies. To enable MS CRA to check for a maximum of five mirror copies of a volume, you must install the following patch, or a superseding patch:

- PHCO\_37562

MS CRA checks for a maximum of 511 physical volumes in a volume group. However, Version 2.1 volume groups allow up to 2048 physical volumes. To enable MS CRA to check for a maximum of 2048 physical volumes in a volume group, you must install the following patch:

- PHCO\_38145

PHCO\_38145 supersedes PHCO\_37562. Installing PHCO\_38145 enables both checks.

## mkfs and mkboot commands

The `mkfs` and `mkboot` commands check whether a specified disk device is being used by LVM; if so, they display an error message. For these commands to correctly check whether a device is part of a Version 2.x volume group, you must install the following patches:

- PHCO\_37328
- PHCO\_37340
- PHCO\_37394

These patches are delivered in the `FEATURE11i` bundle.

## HP Serviceguard

To support Version 2.0 volume groups, HP Serviceguard 11.18 requires the April 2008 patch. Customers requiring Version 2.1 volume group support or Version 2.0 and higher cluster lock disks

must install the HP Serviceguard September 2008 patch. For more information, see the *HP Serviceguard Version A.11.18 Release Notes*.

Cluster device special files (cDSFs) require the installation of the following patches on HP-UX 11i v3 September 2010 OE:

- PHSS\_41225 11.31 Serviceguard A.11.20.00 patch
- PHCO\_41235 11.31 iocdsfd(1M) and io\_cdsf\_config(1M) patch

For more information, see the *New Features for the September 2010 Patch* section of the *HP Serviceguard Version A.11.20 Release Notes*.

## VxVM commands

Several VxVM commands check whether a specified disk device is being used by LVM before overwriting it. For these commands to correctly check whether a device is part of a Version 2.x volume group, you must install the following patch:

- PHCO\_37836 for VxVM 4.1
- PHCO\_40294 for VxVM 5.0

## Required disk space

Because LVM is installed with the HP-UX 11i v3 operating environment, it consumes no additional disk space.

MirrorDisk/UX consumes no additional disk space aside from a license key.

## Other requirements

### HP Insight Remote Support Software

HP strongly recommends that you install HP Insight Remote Support software to complete the installation or upgrade of your product and to enable enhanced delivery of your HP Warranty, HP Care Pack Service or HP contractual support agreement. HP Insight Remote Support supplements your monitoring, 24x7 to ensure maximum system availability by providing intelligent event diagnosis, and automatic, secure submission of hardware event notifications to HP, which will initiate a fast and accurate resolution, based on your product's service level. Notifications may be sent to your authorized HP Channel Partner for on-site service, if configured and available in your country. The software is available in two variants:

- **HP Insight Remote Support Standard:** This software supports server and storage devices and is optimized for environments with 1-50 servers. Ideal for customers who can benefit from proactive notification, but do not need proactive service delivery and integration with a management platform.
- **HP Insight Remote Support Advanced:** This software provides comprehensive remote monitoring and proactive service support for nearly all HP servers, storage, network, and SAN environments, plus selected non-HP servers that have a support obligation with HP. It is integrated with HP Systems Insight Manager. A dedicated server is recommended to host both HP Systems Insight Manager and HP Insight Remote Support Advanced.

Go to <http://www.hp.com/go/insightremotesupport> for details on both versions.

To download the software, go to Software Depot: <http://hp.com/go/softwaredepot>

Select **Insight Remote Support** from the menu on the right.

## Compatibility issues

This release is specific for HP-UX 11i v3 (B.11.31). The following compatibility issues exist in the March 2010 release of HP-UX 11i v3.

This release is specific for HP-UX 11i v3 (B.11.31). The following compatibility issues exist in the September 2010 web release of HP-UX 11i v3.

This release is specific for HP-UX 11i v3 (B.11.31). The following compatibility issues exist in the October 2010 web release of HP-UX 11i v3.

## Version 2.x volume groups

- Version 2.2 volume groups are not recognized on previous releases of HP-UX, including versions of HP-UX 11i v3 prior to the March 2010 Update release. Version 2.1 volume groups are not recognized on previous releases of HP-UX, including versions of HP-UX 11i v3 prior to the September 2008 release. Version 1.0 volume groups are supported on all supported versions of HP-UX, including 11i v1, 11i v2, and 11i v3.
- The administration of boot disks bigger than 2 TB should not be attempted on a system running versions of LVM prior to the March 2013 release. The administration of boot disk here includes any usage of LVM commands that change boot configuration (`lvlnboot(1M)`, `lvrmboot(1M)`) or changes data layout (`lvextend(1M)`, `lvreduce(1M)`, `pvmove(1M)` etc).
- Downgrade of LVM software from version March 2013 or later to a release prior to March 2013 should not be attempted on a boot disk of size greater than 2TB. Doing so may result in boot failure when Maintenance Mode Boot (MMB) is attempted subsequently with the downgraded kernel. Also the downgrade of entire OE version from March 2013 or later to a release prior to March 2013 may result in regular boot failure if attempted on a boot disk of size greater than 2TB.  
See "HP-UX System Administrator's Guide: Logical Volume Management" for more details.
- There is currently no method for converting a volume group in place from one version to another. To migrate a volume group to a different version, you must create a new volume group and copy the data.
- The following HP-UX product does not currently support Version 2.x volume groups:
  - HP Process Resource Manager (HP PRM)  
This product plans to add support of Version 2.x volume groups. For the most recent information on this products, see the IT Resource Center (ITRC) at <http://itrc.hp.com>, or consult the release notes for the specific product.
- Encrypted Volume and File System (EVFS) v1.1 and greater supports Version 2.x volume groups.

## Moving from HP-UX 11i v2 to HP-UX 11i v3

If you are migrating a system from HP-UX 11i v2 to HP-UX 11i v3, see the LVM migration white paper described in "[Related documentation](#)". It contains information on migrating an LVM configuration from the legacy naming model to the agile naming model.

Existing LVM configurations created on HP-UX 11i v2 continue to work on HP-UX 11i v3 under the legacy naming model. However, there is a change in behavior for **Alternate Links (PVLinks)**:

In HP-UX 11i v3, management of multipathed devices is available outside of LVM using the next generation mass storage stack. By default, the next generation mass storage stack distributes I/O requests across all available paths to a multipathed disk, even when using legacy device special files. Using LVM with persistent or legacy device special files might cause I/O requests to be sent across alternate links, even if the links are not configured as PVLinks; this does not introduce any errors, but it does differ from PVLink behavior in previous releases.

When using LVM configuration commands on legacy device special files, LVM does not select an alternate path if the path corresponding to the specified device special file is unavailable, unless the unavailable path and the alternate path are configured as part of an active volume group.

When using LVM configuration commands on persistent device special files, LVM succeeds if at least one of the paths to the device is available.

HP recommends converting volume groups with multipathed disks to persistent device special files and using native multipathing, as described in the migration white paper in [“Related documentation”](#).

However, if you want backward-compatible PVLink behavior, you must use legacy device special files for physical volumes, and disable the mass storage stack multipathing for those physical volumes. To disable multipathing on legacy device special files, use the `scsimgr` command to configure a global device tunable called `leg_mpath_enable`.

For each multipathed disk, enter the following command:

```
# scsimgr save_attr -D /dev/rdisk/diskn -a leg_mpath_enable=false
```

Alternatively, you can disable multipathing for all legacy device files with this command:

```
# scsimgr save_attr -a leg_mpath_enable=false
```

Note that this has no effect on multipathing through persistent device special files. For more information, see `scsimgr(1M)`.

## Moving volume groups from HP-UX 11i v3 to previous HP-UX releases

If a volume group used on HP-UX 11i v3 is accessed from a system running a previous release of HP-UX 11i, you might encounter these compatibility issues:

- **Version 2.x volume groups:** As noted in [“Version 2.x volume groups”](#) (page 11), Version 2.x volume groups are not recognized on previous releases of HP-UX.
- **Logical volumes larger than 2 TB:** Releases prior to HP-UX 11i v3 can only access data within the first 2 TB of a logical volume. If a logical volume larger than 2 TB is created on HP-UX 11i v3, its use is not recommended on any previous HP-UX release. The volume group can be activated and the logical volume can be used, but any data in that logical volume beyond 2 TB will be inaccessible.

---

**NOTE:** Patches PHKL\_36745, PHCO\_36744, and PHCO\_37939 resolve this compatibility issue for HP-UX 11i v2. Installing these patches enables the creation and use of logical volumes up to 16 TB.

---

- **Striped Mirrors:** Releases prior to HP-UX 11i v3 only support extent-based striping via the `-D` option to `lvcreate`. If a logical volume using simultaneous mirroring and non-extent-based striping is created on HP-UX 11i v3, attempts to import or activate its associated volume group will fail on a previous HP-UX release. To import the volume group, you must remove the incompatible logical volumes or reduce them to a single mirror.

---

**NOTE:** Patches PHKL\_36745, PHCO\_36744, and PHCO\_37939 resolve this compatibility issue for HP-UX 11i v2. Installing these patches enables the creation and use of striped mirrors.

---

- **Mirror Write Cache (MWC):** When a volume group containing a logical volume using the Mirror Write Cache is activated on HP-UX 11i v3, its Mirror Write Cache format is converted to a new format. Importing or activating the volume group on a previous HP-UX release does not recognize the new format and triggers a full resynchronization of the mirrors.

---

**NOTE:** Patch PHKL\_36244 adds support for the new MWC format to HP-UX 11i v2. Installing this patch avoids the unnecessary resynchronization.

---

In the future, HP might add support of larger logical volumes and striped mirrors to previous releases of HP-UX 11i. However,

As of this writing, larger logical volumes and striped mirrors are only available on HP-UX 11i v3 and HP-UX 11i v2 (with patches). For up-to-the-moment information, contact your Hewlett-Packard

support representative or consult the Hewlett-Packard IT Resource Center site: <http://itrc.hp.com> (Americas and Asia Pacific) or <http://europe.itrc.hp.com> (Europe) for LVM patches.

## Sharing volume groups in an HP Serviceguard cluster

To share a volume group in an HP Serviceguard cluster where the cluster nodes have different versions of the LVM product installed, use the following guidelines:

- To share a version 2.1 volume group, you must install the HP-UX 11i v3 September 2008 release or above on all the cluster nodes.
- To share a 2.2 volume group, you must install the HP-UX 11i v3 March 2010 release or above on all the cluster nodes.
- To use the Multi Node Online Reconfiguration (MORE) feature, you must install the HP-UX 11i v3 September 2009 release or above on all nodes of the cluster. Further, you can use MORE only on version 2.1 or above volume groups.

---

**NOTE:** MORE does not operate in a cluster with both HP-UX 11i v3 September 2009 and HP-UX 11i v3 March 2010 nodes. During a rolling update from HP-UX 11i v3 September 2009 to HP-UX 11i v3 March 2010, you cannot use MORE on the shared volume group.

---

## Related documentation

LVM and MirrorDisk/UX manpages are installed with the product and provide the latest documentation for all LVM commands and options. These manpages can be displayed via the man command.

In addition, LVM user documentation is available in English at <http://www.hp.com/go/hpux-core-docs>. Available documents include:

- *HP-UX Logical Volume Manager and MirrorDisk/UX Release Notes* (this version and previous versions)
- *HP-UX System Administrator's Guide: Logical Volume Management* (this document is localized)
- *LVM 2.0 Volume Groups in HP-UX 11i v3*
- *LVM New Features in HP-UX 11i v3*
- *LVM Limits*
- *LVM Migration from Legacy to Agile Naming Model: HP-UX 11i v3*
- *LVM Online Disk Replacement (LVM OLR)*
- *Using the vgmodify command to Perform LVM Volume Group Dynamic LUN Expansion (DLE) and Contraction (DLC)*
- *LVM Volume Group Quiesce/Resume*
- *SLVM Single-Node Online Reconfiguration (SLVM SNOR)*
- *SLVM Online Reconfiguration*
- *When Good Disks Go Bad: Dealing with Disk Failures under LVM*
- *Using the vgversion Command to Perform LVM Volume Group Version Migration*
- *Using LVM Logical Volume Snapshots*

## Software Availability in Native Languages

The commands delivered with LVM and MirrorDisk/UX support localized message catalogs. The kernel components which generate messages directly to the console and the system log are available only in the English language.

LVM and MirrorDisk/UX manpages are available in English and Japanese.

The *HP-UX System Administrator's Guide: Logical Volume Management* is available in English, Japanese, and Simplified Chinese. These documents are on <http://www.hp.com/go/hpux-core-docs>.