

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

HP-UX 11i v3 September 2012 Release (B.11.31)

HP Part Number: 5900-2442
Published: September 2012
Edition: 12.0



© Copyright 2012 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
pvcreate enhancement to detect Oracle Disks.....	4
Fixed issues in this version.....	5
Fixed issues in LVM Provider product.....	6
Known issues and limitations.....	6
Known issues.....	6
“vgchange -c [y n]” does not safe guard against cluster-ID corruption.....	6
vgmodify -a -S modifies a subset of the PVs.....	7
vgimport 2.X fails saying a PV doesn't belong to the same VG.....	7
vgchange returns non-zero with option -c n -S n.....	7
VG Quiesce on shared vg hangs while resync is going on.....	7
Limitations.....	7
Installation requirements.....	8
Required hardware.....	8
Required software.....	8
Required patches.....	8
Mass Storage Critical Resource Analysis.....	8
mkfs and mkboot commands.....	8
HP Serviceguard.....	9
VxVM commands.....	9
Required disk space.....	9
HP Insight Remote Support Software.....	9
Compatibility issues.....	10
Version 2.x volume groups.....	10
Moving from HP-UX 11i v2 to HP-UX 11i v3.....	10
Moving volume groups from HP-UX 11i v3 to previous HP-UX releases.....	11
Sharing volume groups in an HP Serviceguard cluster.....	11
Related documentation.....	12
Software Availability in Native Languages.....	12

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 September 2012 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.

This September 2012 release of LVM and MirrorDisk/UX provides defect fixes as described in [“Fixed issues in this version”](#) (page 5). In addition, it provides an enhancement to detect Oracle disks.

New and changed features in this release

pvcreate enhancement to detect Oracle Disks

Beginning with September 2012 release, the pvcreate command has been enhanced to detect disk or disk partitions used by Oracle OCR or Vote in addition to ASM (Enhancement CR QXCR1001176636). The pvcreate command will now fail with an appropriate message if the input disk or disk partition contains a signature for any of the above known Oracle disk types.

Fixed issues in this version

The following table lists the known LVM and MirrorDisk/UX defects fixed in the September 2012 release of HP-UX 11i v3.

Table 1 LVM Fixes in HP-UX 11i v3 September 2012 Release

Defect ID	Description
QXCR1001197843	<p>mkfs command will fail with following errors when performed on LVM 2.x VG with LV minor 2047 (7ff hexadecimal).</p> <pre># mkfs -F hfs /dev/vg01/lvol2047 mkfs (hfs): read of lvm record failed mkfs (hfs): No such device # mkfs -F vxfs /dev/vg01/lvol2047 vxfs mkfs: V-3-21203: read of lvm record failed # ll /dev/vg01/lvol2047 brw-r... 128 0x0007ff /dev/vg01/lvol2047</pre> <p>The last three digits of the device minor is LV minor as highlighted above.</p>
QXCR1001178600	<p>Every time the cluster is stopped (cmhaltcl) and restarted (cmruncl), LVM sub system makes memory leak of about 0.5625 MB on all the nodes. This memory leak happens irrespective of whether any LVM Volume group is activated on any of the node.</p>
QXCR1001125477	<p>lvchange -s g on a 1.0 LV fails with following error messages if lvmpvg contains any 2.X VGs:</p> <pre># lvchange -s g /dev/vg01/lvol1 Error: Information in "/etc/lvmtab" or "/etc/lvmtab_p" differs from that in "/etc/lvmpvg". VG "/dev/anothervg" is not defined in file "/etc/lvmtab". VG "/dev/anothervg" is not defined in file "/etc/lvmtab_p". lvchange: Error detected when reading from file "/etc/lvmpvg".</pre>
QXCR1001197577	<p>When pvchange -T is used to handle media corrupted PVs by specifying threshold time, the physical disk may get detached prematurely even before the threshold is reached. This will occur only when the first occurrence of media error occurs before the system clock wraps (497 days since UP time) and the next occurrence is after the wrap.</p> <p>When this happens, 'Media Error Elapsed Time' field of pvdisplay shows unusually huge value as below:</p> <pre># pvdisplay -v /dev/disk/disk1 ... Media error threshold time 2000 Media error elapsed time 18446744073666602006</pre>
QXCR1001204527	<p>Running multiple vgimport in parallel causes usage of same minor number for more than one VG.</p>
QXCR1001196582	<p>In a cluster with 3 or more nodes, when VG recovery is in progress for at least 5 VGs on each node following a previous server node crash, one or more nodes may get TOC'ed by SG. On the TOC'ed node, most of the lvmkd daemons will have the following stack trace:</p> <pre>sleep+0x3f0 () slvm_send_vgsuspend_resume_req_common+0x330 () slvm_send_vgsuspend_cstate_request+0x40 () slvm_vgsuspend_recovery+0xb0 () slvm_do_remaining_recovery_asynchronously+0xb0 () lvmkd_daemon+0xe0 () lvmkd_fork1+0x170 () lvmkd_fork+0x50 () lvmkd_init+0x50 ()</pre>
QXCR1001212048	<p>Performance Collection Component services like midaemon that are part of HP Operations Agent experience hang when volume groups are in deactivated state.</p> <p>Export the environment variable LVML_SKIP_INACTVG=1 in the midaemon start up scripts to avoid this condition with the 1209 version of the liblvm library.</p>

Table 1 LVM Fixes in HP-UX 11i v3 September 2012 Release (continued)

Defect ID	Description
QXCR1001136971	<p>If pstat_dynamic() is invoked while vgcreate command is running in parallel, a panic occurs with following stack trace:</p> <pre>lv_pbufs_usage_body+0x2a0 () lv_pbufs_usage+0x30 () pstat_dynamic() pstat+0x10b0 () syscall+0x530 ()</pre>
QXCR1001182973	<p>2.x volume groups can fail activation with the following error message if they were migrated from their 1.0 version using vgversion prior to the 1108 release or vgversion from the 1109 release and they had one or more logical volumes on which an lvextend operation had failed:</p> <pre># vgchange -a y myvg vgchange: Couldn't activate volume group "myvg": On disk LVM metadata structures are corrupted.</pre> <p>The pre-existing vgchange -l option has been enhanced to deal with this problem.</p> <p>Use -l option with vgchange -a to determine if the activation failed due to inconsistencies in the logical volume size and extent properties.</p> <p>Messages in syslog similar to the following will help confirm this:</p> <pre>vmunix: LVM:Warning:No PE for LV [major: 128, minor:0x004001] mirror 0 LE 10 vmunix: LVM:Warning:2037 of 2047 extents have no PE for LV [major: 128, minor:0x004001] mirror 0</pre> <p>Use lvextend+lvreduce to fix such logical volumes and retry activation without the -l option.</p>

Fixed issues in LVM Provider product

Table 2 Fixed issues in LVM Provider product

Defect ID	Description
QXCR1001214321	<p>LVMP: Enumerating HPUX_PhysicalVolume reports error during vgchange.</p> <p>Activation and deactivation of volume groups in a continuous loop throws incorrect exception while enumerating the class HPUX_PhysicalVolume. This causes good physical volumes not to be displayed by CIMUtil.</p>

Known issues and limitations

This section provides a list of known issues and limitations as known to HP at time of publication. If workarounds are available, they are included.

Known issues

“vgchange -c [y | n]” does not safe guard against cluster-ID corruption

Defect ID:	QXCR1001161003
Problem:	When the volume group is activated in shared mode in other nodes, running ‘vgchange -c n’ on any node in SLVM environment would leave the cluster in a state that prevents the newly joining node to activate the volume group.
Severity:	Serious
Corrective Action:	Change the physical volume’s properties on the other node in the cluster after activating the volume group in exclusive mode. Before performing this step, the volume group must be deactivated in all the other nodes except this node.

vgmodify -a -S modifies a subset of the PVs

Defect ID: **QXCR1001211881**
Problem: Under certain conditions a subset of the PV's in the vg are changed in size by vgmodify but this has no advantage and leaves the vg susceptible to further problems.
Severity: Serious
Corrective Action: If the failure occurred on a PV due to lack of free PEs, then use pvmove as suggested in the output of vgmodify to free up extents and retry the vgmodify operation.
To prevent this condition from occurring, use the preview mode of vgmodify (-r option) to make sure the modification is possible.

vgimport 2.X fails saying a PV doesn't belong to the same VG

Defect ID: **QXCR1001222652**
Problem: Under rare conditions, vgimport fails with following failure message because of a metadata ID mismatch.
vgimport vg22 /dev/dsk/vrd05[01]p00
vgimport: The Physical Volumes specified on the command line do not belong to the same Volume Group
Severity: Serious
Corrective Action: Run vgchgid on all the PVs in the VG. vgchgid will change the VGID on all copies.

vgchange returns non-zero with option -c n -S n

Defect ID: **QXCR1001217016**
Problem: vgchange -c n -S n <vgname> would return non-zero value in a non-clustered system.
Severity: Serious
Corrective Action: Skip -S n since it is implicit when -c n is specified. Just use vgchange -c n.

VG Quiesce on shared vg hangs while resync is going on

Defect ID: **QXCR1001154451**
Problem: In a cluster with 3 or more nodes, suspending the VG (vgchange -Q) on the first client node hangs if resync is in progress for that VG and also the server of the VG crashes at the same time.
Severity: Serious
Corrective Action: None.

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.

- 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.

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.

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

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 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 (PVLlinks)**:

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 PVLlinks; this does not introduce any errors, but it does differ from PVLlink 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 PVLlink 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 10), 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.

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>.