

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

HP-UX 11i v3 (B.11.31)

HP Part Number: 5991-7518
Published: February 2007
Edition: 1.0



© Copyright 2007 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.

Table of Contents

Logical Volume Manager and MirrorDisk/UX Release Notes.....	5
About this Document.....	5
LVM and MirrorDisk/UX Overview.....	5
LVM.....	5
MirrorDisk/UX.....	5
Overview of Changes.....	5
New and Changed Features in This Release.....	5
New Features.....	5
Changed Features.....	7
Deprecated or Obsolete Features.....	7
Known Problems and Limitations.....	7
Installation Requirements.....	8
Required Hardware.....	8
Required Software.....	8
Required Patches.....	8
Required Disk Space.....	8
Compatibility Issues.....	9
Moving from HP-UX 11i v2 to HP-UX 11i v3.....	9
Moving Volume Groups from HP-UX 11i v3 to Previous HP-UX Releases.....	9
Related Documentation.....	11
Software Availability in Native Languages.....	11

Logical Volume Manager and MirrorDisk/UX Release Notes

About this Document

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

LVM and MirrorDisk/UX Overview

LVM

Logical Volume Manager (product BaseLVM) is the HP-UX default Volume Manager. It provides the 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

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

Overview of Changes

This release of LVM and MirrorDisk/UX is integrated with the new HP-UX 11i v3 mass storage stack, and delivers significant performance, scalability, availability, and usability enhancements. LVM has been enhanced to support larger logical volumes, temporary quiescing of volume groups, and striping with mirroring. Operations that previously required deactivation or recreation of a volume group—such as replacing a disk, resizing a LUN, and modifying volume group characteristics—can now be performed online.

New and Changed Features in This Release

New Features

The following features are new with the HP-UX 11i v3 release of LVM:

- **LVM Device Online Replacement (OLR):** The `pvchange` command has a new option, `-a`, that can be used to temporarily stop LVM accessing a device or device special file, and to reenable access.

A white paper, described in “Related Documentation”, explains how OLR simplifies the process of replacing or isolating disks.

- **Volume Group Quiesce/Resume:** Two new options, `-Q` and `-R`, have been added to the `vgchange` command. These options temporarily quiesce an activated volume group and resume it, respectively. Quiescing a volume group keeps its composite disk image in a consistent state—without deactivating the volume group—so that disk management utilities can snapshot all the disks in the volume group.

Using the `-Q` option, you can quiesce both read and write operations to the volume group, or just write operations. While the volume group is quiesced, the `vgdisplay` command reports the volume group access mode as “quiesced”. The indicated I/O operations are queued until the volume group is resumed, and commands that would modify the volume group configuration return an error.

For more information on quiescing and resuming volume groups, see the `vgchange(1M)` manpage or the *Quiescing and Resuming a Volume Group* section of the *HP-UX System Administrator’s Guide: Logical Volume Management*.



NOTE: Individual physical volumes or logical volumes cannot be quiesced using this feature. To temporarily quiesce a physical volume, use the `pvchange` command to disable all links to it. To quiesce a logical volume, deactivate the volume group; to provide a stable image of a logical volume without deactivating the volume group, mirror the logical volume, then split off one of the mirrors using the `lvsplit` command.

- **Mass Storage Stack Support:** LVM supports the agile view of mass storage introduced in HP-UX 11i v3. In particular, LVM supports the use of both legacy and persistent device special files within the same volume group. New options to the `vgscan` and `vgimport` commands, described below, affect how LVM creates its LVM configuration.

By default, `vgscan` recovers the LVM configuration information (the `/etc/lvmtab` file) using legacy device special files. If the new `-N` option is specified, then `vgscan` uses persistent device special files. If the new `-B` option is specified, then `vgscan` populates the `/etc/lvmtab` file using both legacy and persistent device special files.

By default, when importing a volume group in shared mode, `vgimport` populates the `/etc/lvmtab` file using legacy device special files. If the new `-N` option is specified, then `vgimport` uses persistent device special files.

For more information about the mass storage changes in HP-UX 11i v3, see the white paper *The Next Generation Mass Storage Stack* in the Network and Systems Management section of <http://docs.hp.com>, under Storage Area Management.

- **SLVM Single Node Online Volume Reconfiguration (SNOR):** A new option, `-x`, has been added to the `vgchange` command. This option allows an administrator to make configuration changes to a shared volume group while keeping the volume group activated on one cluster node.

A white paper, described in “Related Documentation”, explains the SNOR functionality.

- **Dynamic Volume Group Modification:** A new command, `vgmodify`, is available to dynamically modify volume group characteristics. In previous releases, the number of physical volumes, the number of logical volumes, and number of physical extents per disk were set when a volume group was created; the `vgmodify` command allows these parameters to be modified without recreating the volume group.

For more information on `vgmodify`, see the *Modifying Volume Group Parameters* section of the *HP-UX System Administrator's Guide: Logical Volume Management*.

- **Dynamic LUN Expansion:** If the administrator increases the size of a LUN, the `vgmodify` command can be used to incorporate that additional space into the volume group, without recreating the volume group.

For more information on resizing LUNs, see the *Modifying Physical Volume Characteristics* section of the *HP-UX System Administrator's Guide: Logical Volume Management*.

- **Boot resiliency:** If during boot time the LVM subsystem detects an inconsistency between the firmware boot path and the LVM root volume group configuration, LVM scans all the disk devices to find the physical volumes belonging to the root volume group and continues the boot sequence. In previous releases, the administrator had to boot in LVM Maintenance Mode to resolve this inconsistency.

The following feature is new with the HP-UX 11i v3 release of MirrorDisk/UX:

- **Striped Mirrors:** In previous releases, LVM supported a limited extent-based striped mirror functionality, as described in the `lvcreate(1M)` manpage. This type of striped mirror required the stripe size to be a multiple of extent size. In HP-UX 11i v3, LVM supports mirroring of striped logical volumes with the entire range of stripe sizes. The `lvcreate` command options `-m` and `-i/-I` can be used together, and the `-m` option of the `lvextend` and `lvreduce` commands can be applied to striped logical volumes. Note that logical volumes

configured with such striped mirrors cannot be imported on releases before HP-UX 11i v3; see the “Compatibility Issues” section.

To maximize data integrity for striped mirrors, LVM enforces a strict allocation policy; that is, mirrored physical extents must be allocated on different physical volumes. This forces the physical extents used by all the replicas of a strip onto different physical volumes.

Changed Features

The following features have changed from the previous release of LVM:

- **Increased logical volume size:** LVM now supports logical volumes up to 16 terabytes (16 TB) in size, an increase from 2 TB in previous releases. Logical volumes larger than 2 TB may not be usable on previous releases; see the “Compatibility Issues” section.
- **Display enhancements:** The `lvdisplay`, `pvdisplay`, `vgdisplay`, and `vgscan` commands all support long hostnames. Those commands also support a new `-F` option to print in a format more easily parsed by user scripts. The `pvdisplay` command has a new `-l` option to display if a disk is under LVM control; the existing `-d` option displays if a physical volume is a bootable physical volume.
- **vgscan enhancements:** In addition to support of the mass storage stack, `vgscan` has new `-f` and `-k` options.

By default, the `vgscan` command does not modify or supplement the `/etc/lvmtab` file entries for volume groups that already have entries. The new `-f` option forces an update of the existing entries for the specified volume group.

By default, the `vgscan` command scans the I/O configuration searching for LVM physical volumes, which can be a time-consuming operation. The new `-k` option reads the LVM data structures in kernel memory and populates the `/etc/lvmtab` file based on that data.

Deprecated or Obsolete Features

The following features are deprecated or obsolete with the HP-UX 11i v3 release of LVM:

- **Bad Block Relocation:** LVM no longer performs software bad block relocation, as modern disks and disk arrays handle such relocation in their own hardware. Existing software relocation information is honored, unless the physical volume is larger than 256 GB.
- **Kernel Tunable maxvgs:** The tunable parameter `maxvgs` is obsolete. LVM can now create up to 256 volume groups dynamically. Any attempts to modify `maxvgs` using the `kctune` command cause the following error message:

```
ERROR: There is no tunable named 'maxvgs'.
```

Known Problems and Limitations

No known problems were identified for this release of LVM and MirrorDisk/UX.

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

Since 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.

Required Disk Space

Since 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.

Compatibility Issues

This release is specific for HP-UX 11i v3 (B.11.31).

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, please refer to the LVM migration white paper described in the “Related Documentation” section. 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:

- **Alternate Links (PVLINKS):** In HP-UX 11i v3, management of multi-pathed 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 multi-pathed disk, even when using legacy device special files. Using LVM with persistent or legacy device special files may cause I/O requests to be sent across alternate links, even if the links are not configured as PVLINKS; this should not introduce any errors, but it does differ from PVLINK behavior in previous releases.

HP recommends converting volume groups with multi-pathed disks to persistent device special files and native multi-pathing, 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 multi-pathing for those physical volumes. To disable multi-pathing on legacy device special files, use the `scsimgr` command to configure a global device tunable called `leg_mpath_enable`.

For each multi-pathed disk, enter the following command:

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

Alternatively, you can disable multi-pathing for all legacy device files with this command:

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

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

- **Parallel Resynchronization:** A HP-UX 11i v1 patch (PHCO_35313) introduced new command line options to the `lvextend`, `lvmerge`, and `lvsync` commands which allowed the user to control the synchronization of stale extents. As of this writing, those enhancements are not available in HP-UX 11i v3.

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 may encounter these issues:

- **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.
- **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.

- **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.

As of this writing, these features were not available on previous releases of HP-UX 11i. Contact your Hewlett-Packard support representative for up-to-the-moment information, 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.

Related Documentation

The latest documentation for LVM is available in English at <http://docs.hp.com> in the *HP-UX Operating Environment* collection. Available documents include:

- *HP-UX System Administrator's Guide: Logical Volume Management* (this document is localized)
- *LVM Limits*
- *LVM Migration from Legacy to Agile Naming Model: HP-UX 11i v3*
- *LVM Online Disk Replacement (LVM OLR)*
- *LVM Volume Group Quiesce/Resume*
- *SLVM Single-Node Online Reconfiguration (SLVM SNOR)*
- *When Good Disks Go Bad: Dealing with Disk Failures under LVM*

Software Availability in Native Languages

LVM and MirrorDisk/UX are available only in the English language.