Veritas Volume Manager Migration Guide

HP-UX

5.0.1

symantec™

Veritas Storage Foundation™ Veritas Volume Manager Migration Guide

Product version: 5.0.1

Document version: 5.0.1.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

www.symantec.com/techsupp/

Contacting Technical Support

Customers with a current maintenance agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/assistance_care.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system

- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

https://licensing.symantec.com

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

Customer Service is available to assist with the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and maintenance contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

clustering_docs@symantec.com

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Chapter 1

VxVM and LVM

This chapter includes the following topics:

About VxVM and LVM

This chapter provides an overview of Veritas Volume Manager by Symantec (also referred to as VxVM) and its features. It includes a brief description of the benefits of migrating from the HP-UX Logical Volume Manager (LVM) to VxVM. It also includes information about the coexistence of VxVM disks with LVM disks.

Introducing Veritas Volume Manager

Veritas Volume Manager is an alternative Volume Management product for HP-UX that includes mirroring features. It offers many capabilities that are not available with the LVM and MirrorDisk/UX products today, including the following capabilities:

- Veritas Volume Manager can coexist with LVM. Users can decide which volumes they want managed by each volume manager. For users who want to migrate LVM volume groups to VxVM disk groups, a conversion utility is included. The `vxvmconvert` utility is used to convert LVM to VxVM.

  See "About LVM to VxVM conversion" on page 19.

- Veritas Volume Manager is available for installation with all HP-UX 11i Version 3 operating environments.

See the *Storage Foundation Release Notes* for additional details of the features that are supported in this release.

- Veritas Volume Manager can be used instead of LVM to manage the root disk. Basic volume management capabilities are included in the operating system.
- Veritas Volume Manager is integrated with HP ServiceGuard and ServiceGuard OPS Edition for High Availability, but requires a specific version of the ServiceGuard products. Refer to the *Release Notes* for details about the required version number, as well as the availability of specific features in your release.

Notable features of VxVM

The Veritas Volume Manager provides many features, some of which are not available with LVM or MirrorDisk/UX. This section describes notable VxVM features.

See the *Storage Foundation Release Notes* for additional details of the features that are supported in this release.

See other Veritas Volume Manager documentation for additional details about using these features.

Veritas Volume Manager includes the following features:

- Concatenation, the combining of discontiguous disk regions into virtual devices.
- Spanning, concatenation across different physical media.
- Striping, distribution of storage mappings for a virtual device so that multi-threaded accesses tend to cause even use of all physical media.
- The Veritas Enterprise Administrator (VEA), which is a JAVA-based GUI for VxVM.
- Dynamic Multipathing (DMP) for Active/Passive and Active/Active devices. DMP provides higher availability to data on disks with multiple host-to-device pathways by providing a device path failover mechanism. If one connection to a disk is lost, the system continues to access the data over the other available connections to the disk.

  DMP can also provide improved I/O performance from disks with multiple pathways that are concurrently available. DMP can balance the I/O load uniformly across the multiple paths to the disk device.

  DMP can coexist with the native multipathing functionality that is provided in HP-UX 11i Version 3.

  DMP also supports the new persistent device file names in addition to legacy device file names and enclosure-based names.
- Free Space Management, providing simple goal-based allocation of storage.

- Task Monitor, which tracks the progress of system recovery by monitoring task creation, maintenance, and completion. The Task Monitor lets you pause, resume, and stop as desired to adjust the impact on system performance.
- Multiple mirroring with up to 32 mirror copies of a volume's address space.
- Mirrored stripes (RAID-0 + RAID-1) and striped mirrors (RAID-1 + RAID-0) combine the benefits of striping and mirroring. These layouts improve performance by spreading data across multiple disks, and provide redundancy of data by mirroring. Striped mirror volumes are more tolerant of disk failure and have a shorter recovery time than mirrored stripe volumes. More detailed information is available on these layouts.
  See *Veritas Volume Manager Administrator's Guide*.
- Hot-relocation, which allows a system to react automatically to I/O failures on redundant (mirrored or RAID-5) VxVM objects. This feature restores redundancy and access to those objects without administrative intervention. VxVM detects I/O failures on objects and relocates the affected subdisks. The vxunreloc utility can be used to restore the system to the same configuration that existed before the disk failure.
- RAID-5, which provides data redundancy by using parity, at a lower storage cost than mirroring. RAID-5 provides data redundancy by using parity. Parity is a calculated value used to reconstruct data after a failure. While data is being written to a RAID-5 volume, parity is calculated by doing an exclusive OR (XOR) procedure on the data. The resulting parity is then written in an interleaved fashion to the RAID-5 array that is established by the volume. If a portion of a RAID-5 volume fails, the data that was on that portion of the failed volume can be recreated from the remaining data and parity information.
- Online Data Migration, which allows for regions of storage on physical media to be dynamically moved to other physical devices.
- Online Relayout or Dynamic Restriping, the ability to change logical data configuration while online. For example, use this feature to change RAID-5 to a mirrored layout or to change a stripe unit size. The volume data remains available during the relayout.
- Improved RAID-5 subdisk moves, using layered volume technology where the RAID-5 subdisk move operation leaves the old subdisk in place while the new one is being synchronized, thus maintaining redundancy and resiliency to failures during the move.

Note: Additional information is available on LVM and VxVM commands.

See *HP-UX Managing Systems and Workgroups*

See LVM manual pages in *HP-UX Reference Volumes 2, 3, and 5*.

See the Veritas Volume Manager documentation.

VxVM and LVM—conceptual comparison

The following section compares the terminology that is used in LVM and VxVM at a conceptual level. For more information, refer to the glossary of this Guide for precise and detailed definitions of these terms.

Table 1-1 A conceptual comparison of LVM and VxVM

LVM term	VxVM term	Description
LVM	VxVM	Both LVM and VxVM enable online disk storage management. They both build virtual devices, called volumes, on physical disks. Volumes are not limited by the underlying physical disks, and can include other virtual objects such as mirrors. Volumes are accessed through the HP-UX file system, a database, or other applications in the same manner as physical disks would be accessed.

Table 1-1 A conceptual comparison of LVM and VxVM *(continued)*

LVM term	VxVM term	Description
Physical volume	VxVM disk	An LVM physical volume and a VxVM disk are conceptually the same. A physical disk is the basic storage device (media) where the data is ultimately stored. You can access the data on a physical disk by using a device name (`devname`) to locate the disk. In LVM, a disk that is initialized by LVM becomes known as a physical volume. A VxVM disk is one that is placed under the Volume Manager control and is added to a disk group. In addition, both LVM and VxVM have the following characteristics: ■ Volumes can be mapped to multiple VxVM disks or LVM physical volumes. ■ VxVM disks must reside in only one disk group, and LVM physical volumes must reside in one volume group.

Table 1-1 A conceptual comparison of LVM and VxVM *(continued)*

LVM term	VxVM term	Description
Logical volume	Volume	An LVM logical volume and a VxVM volume are conceptually the same. Both are virtual disk devices that appear to applications, databases, and file systems like physical disk devices, but do not have the physical limitations of physical disk devices. Due to its virtual nature, a volume (LVM or VxVM) is not restricted to a particular disk or a specific area of a disk. An LVM volume is composed of fixed length extents. LVM volumes can be mirrored or striped, but mirrored-stripe and striped-mirror layouts are not supported. VxVM volumes consist of one or more plexes/mirrors holding a copy of the data in the volume. The plexes or mirrors, in turn, are made up of subdisks with arbitrary length. The configuration of a volume can be changed by using the VxVM user interfaces. See the *Veritas Volume Manager Administrator's Guide* for more information. VxVM volumes can be concatenated, mirrored, striped, RAID-5 or combinations such as mirrored-stripe, striped-mirror, and concatenated-mirror.

Table 1-1 A conceptual comparison of LVM and VxVM *(continued)*

LVM term	VxVM term	Description
Volume group	Disk group	LVM volume groups are conceptually similar to VxVM disk groups. An LVM volume group is the collective identity of a set of physical volumes, which provide disk storage for the logical volumes. A VxVM disk group is a collection of VxVM disks that share a common configuration. A configuration is a set of records with detailed information about related VxVM objects, their attributes, and their associations.
Physical extent	Subdisk	User data is contained in physical extents in LVM and subdisks in VxVM. The LVM physical extents are of a fixed length. LVM allocates space in terms of a physical extent which is a set of physical disk blocks on a physical volume. The extent size must be the same for all physical volumes within a volume group. The extent size is usually 4 MB. VxVM allocates disk space in term of subdisks which is a set of physical disk blocks representing a specific contiguous portion of a VxVM disk and is of arbitrary size.
LVM metadata	Private region	LVM metadata and the Private Region are similar conceptually. In LVM, metadata is stored in a reserved area in the disk. In VxVM, the private region of a disk contains various on-disk structures that the Volume Manager uses for various internal purposes. Private regions can also contain copies of a disk group's configuration, and copies of the disk group's kernel log.

Table 1-1 A conceptual comparison of LVM and VxVM *(continued)*

LVM term	VxVM term	Description
Unused physical extent	Free space	VxVM can place a disk under its control without adding it to a disk group. The VxVM Storage Administrator shows these disks as “free space pool”. LVM contains unused physical extents that are not part of a logical volume, but are part of the volume group. Similarly, free space is an area of a disk under VxVM that is not allocated to any subdisk or reserved for use by any other Volume Manager object.
Mirrors	Mirrors (plexes)	Both LVM and VxVM support mirrors. Mirrors can be used to store multiple copies of a volume’s data on separate disks. In LVM, you can create mirrors using the MirrorDisk/UX product. Mirrors allow duplicate copies of the extents to be kept on separate physical volumes. A VxVM mirror consists of plexes. Each plex is a copy of the volume. A plex consists of one or more subdisks that are located on one or more disks. VxVM volumes can have up to 32 mirrors (where each plex is a copy of data).
Export	Deport	In LVM, exporting removes volume group information from /etc/lvmtab. The volume group must have already been deactivated. Similarly in VxVM, deport makes a disk group inaccessible by the system.

Table 1-1 A conceptual comparison of LVM and VxVM *(continued)*

LVM term	VxVM term	Description
Import	Import	In LVM, import adds a volume group to the system and the volume group information to /etc/lvmtab but does not make the volumes accessible. The volume group must be activated by the `vgchange -a y` command to make volumes accessible. In VxVM, import imports a disk group and makes the disk group accessible by the system.
Bad block pool	No similar term	In LVM, the bad block pool provides for the transparent detection of bad disk sectors, and the relocation of data from bad to good disk sectors. The bad block reallocation feature does not exist in VxVM because the vectoring of bad blocks is now done by most hardware.
`/etc/lvmtab`	No similar term	The `/etc/lvmtab` file contains information about the volume groups that are accessible by a system.

Coexistence of VxVM and LVM disks

Both LVM disks and VxVM disks can exist together on a system. The LVM disks are detected and displayed as such by VxVM. LVM disks are not selected by VxVM for initialization, addition, or replacement.

Both LVM and VxVM utilities are aware of the other volume manager, and will not overwrite disks that are being managed by the other volume manager unless conversion or rollback is being performed between LVM logical volumes and VxVM volumes.

Note: The above behavior is displayed on both LVM version 1 and version 2 volume groups.

The System Management Homepage (SMH) and Veritas Enterprise Administrator (VEA) graphical administrative utilities also recognize and identify all disks on the system.

See "About SMH and the VEA" on page 73.

The `vxvmconvert` command is provided to enable LVM disks to be converted to a VxVM disk format without losing any data.

See "Converting LVM volume groups to VxVM disk groups" on page 21.

Chapter 2

Converting LVM to VxVM

This chapter includes the following topics:

About LVM to VxVM conversion

This chapter explains how to convert your LVM configuration to a VxVM configuration.

The basic tools for conversion are the VxVM commands, `vxvmconvert` and `vxdiskadm`, and the LVM administrative utilities such as `pvremove` and `vgcfgbackup`. The discussion here details how to use these tools and gives some insights into how these tools work.

Note: This release only supports the conversion of LVM version 1 volume groups to VxVM. It does not support the conversion of LVM version 2 volume groups.

The disks on your system that are managed by LVM can be of two types: LVM disks in volume groups, and unused disks.

The former are disks that contain logical volumes and volume groups. Unused disks contain no user data, and are not used by any volume group, but have LVM disk headers written by `pvcreate`. Conversion is done differently for these two types of disks.

For unused LVM disks you can use a combination of `pvremove` and `vxdiskadm`. For LVM disks in volume groups, the primary tool for conversion is the `vxvmconvert` command. More information is available on the `vxdiskadm` command.

See the man page `vxdiskadm`(1M) or the *Veritas Volume Manager Administrator's Guide*.

The `vxvmconvert` utility is an interactive command. You can also use the `vxautoanalysis` command and the `vxautoconvert` command to perform non-interactive analysis and conversion of LVM volume groups. The `vxautorollback` command also lets you reverse the conversion, and turn a disk group back into a volume group.

See "Non-interactive conversion of volume groups" on page 54.

Converting unused LVM physical volumes to VxVM disks

LVM disks which are not part of any volume group, and contain no user data, are simply cleaned up, so that there are no LVM disk headers. Then the disks are given over to VxVM through the normal means of initializing disks.

Warning: Exercise caution while using this procedure to give disks over to VxVM. You must be absolutely certain that the disks are not in use in any LVM configuration. If there is any user data on these disks, it will be lost during conversion.

Removing LVM disk information

To remove LVM disk information from the disks use the following command:

```
# pvremove disk_name
```

The `pvremove` command does not allow the removal of disk headers which indicate a Volume Group membership for the disk. If the disk fails `pvremove` for this reason, you should first make certain that the group membership information is stale. Do this by using `pvdisplay`:

```
# pvdisplay disk_name
```

If `pvdisplay` finds no valid group information associated with the disk, you can overwrite the LVM headers using `pvcreate`:

```
# pvcreate disk_name
```

Warning: If *`disk_name`* is an alternate path to a disk that does not appear in the `lvmtab` file for this system, or is a disk that is in use on another system, but not imported onto this system, then do not `pvcreate -f` the disk. Doing so will destroy the LVM headers. You can use `pvcreate` without the `-f` option safely, as it will only succeed if the disk is not listed in the `/etc/lvmtab` file, and if the LVM headers indicate that it is not a member of any volume group. (That is, `pvcreate` has been run on the disk, but neither `vgcreate` nor `vgextend` have.)

Initializing disks for VxVM use

To initialize the disk for VxVM use, use the `vxdiskadm` command, selecting the option:

```
1) Add or initialize one or more disks
```

Or use the command:

```
# vxdisk init disk_name
```

VxVM utilities will not tamper with anydisks that are recognized as owned by LVM (by virtue of the LVM disk headers). If you attempt to use `vxdisk init`, or `vxdiskadm` on an LVM disk without using the `pvremove` command first, the command fails.

Note: The above behavior is displayed on both LVM version 1 and version 2 volume groups.

Converting LVM volume groups to VxVM disk groups

This section outlines the process for converting LVM volume groups to VxVM disk groups.

Note: Symantec recommends that you read through this section carefully before you begin any volume group conversion.

The conversion process involves many steps. Though there are tools to help you with the conversion, some of these steps cannot be automated. You should be sure to understand how the whole conversion process works, and what you will need to do in the process before beginning a volume group conversion.

The tool used for conversion is `vxvmconvert`. This interactive, menu-driven program walks you through many of the steps of the process of converting volume groups for use by VxVM. Using `vxvmconvert` can reduce the downtime associated with converting from LVM to VxVM. Without the `vxvmconvert` tool, the only possible method of conversion would be to take full backups of user data, destroy the existing LVM configuration leaving only raw disks, recreate the configuration in VxVM, and then reload the user data.

The `vxvmconvert` process converts LVM volume groups to VxVM disk groups in place. In reality, the utility changes disks within LVM volume groups to VxVM disks by taking over the areas of the disks used for LVM configuration information, and creating the equivalent VxVM volume configuration information. User data, the portions of the disks used for file systems, databases and so on, are not affected by the conversion.

The act of conversion changes the names by which your system refers to the logical storage. For this reason, the conversion process is necessarily an off-line one. There can be no application access to user data in the volume groups undergoing conversion. Access to the LVM configuration itself (the metadata of LVM) must also be limited to the conversion process.

Volume group conversion limitations

There are certain LVM volume configurations that cannot be converted to VxVM. Some of the reasons a conversion could fail are:

- A volume group with insufficient space for metadata.
  In the conversion of LVM to VxVM, the areas of the disks used to store LVM metadata are overwritten with VxVM metadata. If the VxVM metadata that needs to be written will not fit the space occupied by the LVM metadata, the group containing the disk cannot be converted. If you have just enough space for the conversion, you probably would want to have more space for future configuration changes.

  Note: The most likely scenario in which a Volume Group cannot be converted, because of insufficient private space, is when a large HP-UX system using “Extent based Striping” is being used

- A volume group containing the root volume.

The`vxvmconvert`command does not convert any volume group that contains a rootable volume, identified by the presence of the LIF area as created by `mkboot`(1M). Not only is the current root volume off limits, but any volume that might be used as an alternate root volume is rejected as well.

Note: VxVM rootability operations and rootability commands are currently not supported on the LVM version 2 volume groups.

Rootability operations involving LVM version 1 root disks like creating VM root image from LVM root disk, restoring the LVM root image from VM root, destroying the LVM root disk, continue to be supported as before.

Note: You can use the `vxcp_lvmroot`(1M) command to create a VxVM root disk on a spare physical disk from the contents of the existing LVM root disk. Similarly, you can use the `vxres_lvmroot`(1M) command to recreate an LVM root disk on a spare disk from the contents of the VxVM root disk. Additional information is available on the VxVM root disk.

See the *Veritas Volume Manager Administrator's Guide*.

- A volume group containing mirrors using the Mirror Write Cache feature for volume consistency recovery.
  Users should be aware that when converting mirrored LVM volumes to VxVM, some of these volumes will likely have the Mirror Write Cache consistency recovery method in force on the volume. The `vxvmconvert` utility can convert these volumes, but must use the Dirty Region Logging (DRL) feature to obtain the same level of functionality. However, since Dirty Region Logging requires some user space to be available for the log, a conversion could fail due to an MWC volume being full, leaving no space for the DRL log. However it is very unlikely that this situation would occur. Note that the MWC and DRL are used only when the system crashes or is improperly shut down, to quickly bring all mirrors in the volume back into a consistent state.
- A volume group containing the `/usr` file system.
  For this release, a volume group containing the `/usr` file system cannot be converted because `vxvmconvert` needs access to files in `/usr`.
- Volume groups with any `dump` or primary `swap` volumes.
  Because this release does not support rootability, `vxvmconvert` will not convert any volume group with dump or primary `swap` volumes. These are volumes known to the boot process. However, swap volumes on volumes other than the

root volume can be converted (as long as this volume is not in the same volume group as the root volume).

- Volume group disks used in MC/ServiceGuard clusters.
  The conversion process does not support conversion of any volume group that is marked as a member of a MC/ServiceGuard or OPS Edition high availability cluster. The volume group must be deactivated and removed from membership in the high availability cluster before it can be converted.
- Volume groups and logical volumes with large number of characters in the name.
  If you convert an LVM volume group with more than 31 characters in the volume group name, VxVM will retain only the first 31 characters of the LVM volume group name. Conversion of a logical volume containing more than 28 characters in the logical volume name is not supported.
  See "About conversion error messages" on page 79.
- Volume groups or logical volumes containing special characters
  Conversion of volume groups or logical volumes containing special characters like #, $, *, _ in the name, is not supported.
- Volume groups used for cluster lock disks.
  The conversion process does not support conversion of a volume group that contains a disk that is being used for a cluster lock disk for an MC/ServiceGuard cluster.
- Volume groups with any disks that have bad blocks in the bad block directory.
  Unlike LVM, VxVM does not support bad block revectoring at the physical volume level. If there appear to be any valid bad blocks in the bad block directory of any disk used in an LVM volume group, the group cannot be converted.
  See "About conversion error messages" on page 79.
- Volume groups with mirrored volumes.
  A conversion fails if the LVM volume group being converted has mirrored volumes, but the system does not have a valid license installed that enables mirroring for VxVM.

The analyze option in `vxvmconvert`, which is described in later sections, aids you in identifying which volume groups can be converted.

Conversion process summary

Several steps are used to convert LVM volume groups to VxVM disk groups. Most of these steps can be done with the `vxvmconvert` utility. All the steps are not compulsory, and some may have to be followed only if there are problems during

conversion. Some of them (e.g. backing up user data) are left to you to accomplish through your regular administrative processes.

The steps in the conversion process are:

- Identifying LVM volume groups for conversion.
- Analyzing an LVM group to see if conversion is possible.
- Taking actions to make conversion possible if analysis fails.
- Backing up your LVM configuration and user data.
- Planning for new VxVM logical volume names.
- Stopping application access to volumes in the volume group to be converted.
- Converting a volume group.
- Taking actions if conversion fails.
- Implementing changes for new VxVM logical volume names
- Restarting applications on the new VxVM volumes.
- Tailoring your VxVM configuration.

These steps are described in detail in later sections of this chapter. Also available are annotated examples on how to use `vxvmconvert`.

See "Examples" on page 37.

Additional information is available on restoring back to your original LVM configuration.

See "Restoring the LVM volume group configuration" on page 33.

Identifying LVM disks and volume groups for conversion

The obvious first step in the conversion process is to identify what you want to convert. The native LVM administrative utilities like `vgdisplay` and SMH can help you identify candidate LVM volume groups as well as the disks that comprise them.

You can also use the `vxvmconvert` command and the `vxdisk` command to examine groups and their member disks. The information presented through the `vxvmconvert` command and the `vxdisk` command and their interpretation are available.

See "Examples" on page 37.

You can also list the LVM disks with the following VxVM command:

```
# vxdisk list
```

Analyzing an LVM volume group to see if conversion is possible

After you have selected a volume group for conversion, you need to analyze it to determine if conversion for VxVM use is possible.

Use the analyze option of `vxvmconvert` to check for problems that would prevent the conversion from completing successfully. Additional information is available on all the conditions that this option checks for.

See "Volume group conversion limitations" on page 22.

The analysis calculates the space required to add the volume group disks to a VxVM disk group, and to replace any existing disks and volumes with VxVM volumes, plexes, and subdisks. If you don't have the required space to convert the disks, the conversion would fail.

Analysis can be run on a live system while users are accessing their data. To analyze LVM volume groups, choose option 1 of the `vxvmconvert` utility.

Note: The analysis option is presented as a separate menu item in `vxvmconvert`, but there is an implicit analysis with any conversion. If you simply select the "Convert LVM Volume Groups to VxVM" menu option, `vxvmconvert` will go through analysis on any group you specify. When you are using the convert option directly, you are given a chance to abort the conversion after analysis, and before any changes are committed to disk.

The analysis option is useful when you have a large number of groups/disks for conversion and some amount of planning is needed before the actual conversion. Installations with many users or critical applications can use the analyze option on a running system. Then conversion downtime can be better planned and managed. Smaller configurations may be better served by using the convert option directly while in a downtime period.

More information is available on converting disk groups.

See "Converting LVM volume groups to VxVM disk groups" on page 21.

Sample examples of the analyze option are available.

See "Examples" on page 37.

Taking actions to make conversion possible if analysis fails

Available is a list of reasons that may cause analysis to fail.

See "Volume group conversion limitations" on page 22.

Messages from `vxvmconvert` will explain the type of failure and any actions that can be taken before retrying the analysis.

Complete details on specific error messages and actions are available.

See "About conversion error messages" on page 79.

Backing up your LVM configuration and user data

After analysis you know which volume group or groups you want to convert to VxVM disk groups. Up to this point, you have not altered your LVM configuration.

By taking the next step (completing the conversion to VxVM), you are significantly changing access to your storage.

Although the conversion process does not move, or in any other way affect user data, you are strongly encouraged to back up all data on the affected disks. Similarly, you should back up the LVM configuration itself.

During a conversion, any spurious reboots, power outages, hardware errors or operating system bugs can have unpredictable and undesirable consequences. You are advised to be on guard against disaster with a set of verified backups.

Backing up an LVM configuration

Use the `vgcfgbackup`(1M) utility before running `vxvmconvert` to save a copy of the LVM configuration.

You can back up the LVM volumes using the following command:

```
# vgcfgbackup -f pathname/filename vol_grp_name
```

Be sure to use the `-f` option to save the data into a file other than the default. `vxvmconvert` uses LVM utilities which themselves save the configuration using `vgcfgbackup`. If you do not use the `-f` option when you attempt to back up the configuration, the conversion process overwrites your attempted backup.

Keep a copy of this LVM configuration offline on tape or some other medium for use in the event of a disaster during conversion.

For example, to put a copy on tape, use the following command:

```
# tar cvf /dev/rmt/c3t0d0BEST /vgbackups/vg08
```

Note: The `vxvmconvert` utility itself also saves a snapshot of the LVM metadata in the process of conversion for each disk. This data is saved in a different format from that of `vgcfgbackup`. It can only be used by the `vxvmconvert` program. With certain limitations, you can reinstate the LVM volumes after they have been converted to VxVM using this data. Even though `vxvmconvert` provides this level of backup of the LVM configuration, you are advised to use `vgcfgbackup` before running `vxvmconvert`.

See "Example: displaying the vxvmconvert menu" on page 37.

Backing up user data

To back up user data, use your regular backup processes.

Warning: Before you do the backup, you should carefully review how to implement changes for new VxVM logical volume names. Backup processes and systems themselves may have dependencies on the volume names currently in use on your system. The conversion to VxVM changes those names. You are advised to understand the implications name changes have for restoring from the backups you are about to make.

See "Implementing changes for new VxVM logical volume names" on page 32.

File system back up of user data

You can use the backup utility that you normally use to back up data on your logical volumes. For example, to back up logical volumes that contain file systems, the `fbackup`(1M) command can be used to back up the data to tape.

For example, to back up the data on `/dev/vg01/lvol3` mounted on `/foodir`, use the following command:

```
# fbackup -0i /foodir -f /dev/rmt/c0t0d0BEST
```

Non-file system back up

If a logical volume you are converting does not contain a file system, and is being used directly by an application (such as a database application), use the backup facilities provided by the application. If no such facility exists, consider using the `dd` command.

Planning for new VxVM logical volume names

When you change from LVM volumes to VxVM volumes, the device names by which your system accesses data are changed. LVM creates device nodes for its logical volumes in `/dev` under directories named for the volume group. VxVM creates its device nodes in `/dev/vx/dsk` and `/dev/vx/rdsk`. When conversion is complete, the old LVM device nodes are gone from the system, and the system will access data on the device nodes in `/dev/vx`.

This change in names can present problems. Any application that refers to specific device node names will be at risk when these names change. Similarly, any files that record specific device node names for use by applications can be problematic.

The most obvious area where this problem arises is in `/etc/fstab`. To handle this problem, `vxvmconvert` will rewrite the `fstab` with the new VxVM names when conversion is done so that `fsck`, `mount`, and related utilities will behave as they did before the conversion.

There are potentially many other applications, though, that may be put at risk by the name changes in conversion. `vxvmconvert` cannot help with these. The system administrator must examine the mechanisms used in each of the following areas to see if they reference LVM device names:

- Databases run on raw logical devices may record the name of that device node.
- Backup systems may do device level backups based on device node names recorded in private files. Also labeling of the backups may record device names.
- Scripts run by cron(1M).
- Other administrative scripts.

Workaround

`vxvmconvert` records a mapping between the names of the LVM device nodes and VxVM device nodes. This data can be used to create symbolic links from the old LVM volume to the new VxVM device names. The mapping is recorded in the file:

```
/etc/vx/reconfig.d/vgrecords/vol_grp_name/vol_grp_name.trans
```

This file provides information on how to proceed further to link the old LVM volume names to the new VxVM device names.

Warning: This method of resolving the naming problem has risks. The symbolic links can become stale. For example, if a database refers to `/dev/vx/rdsk/vol1` through a symbolic link `/dev/vg00/rvol1`("the old LVM name"), and if the underlying VxVM volume configuration is changed in any way, the database could refer to a missing or different volume.

Note: You may want to use this symbolic link approach to ease the transition to VxVM. You can set up the symbolic links after the successful conversion to VxVM. Then, you can do the investigation on a case by case basis for each volume. When you are satisfied that there are no problems introduced by the name change, the symbolic link to that volume can be removed. You must be careful to maintain a static VxVM volume configuration during this transition period.

Over time, the ultimate goal should be that the underlying VxVM naming is used by all applications, and that there are no indirect references to those volumes.

Stopping application access to volumes in the volume group to be converted

No applications can be active on the LVM volume group undergoing conversion. Before attempting to convert any volume group, you must ensure that applications using that group are down. This involves stopping databases, unmounting file systems, etc.

Note: If you are converting a volume with swap space on it, the conversion requires a reboot. The swap space cannot be taken out of control of the operating system with a shutdown to single user mode.

`vxvmconvert` tries to unmount mounted file systems during the conversion. Bear in mind though, that `vxvmconvert` makes no attempt to close down running applications on those file systems, nor does it attempt to deal with applications (e.g., databases) running on raw LVM volumes. More information is available on vxvmconvert.

See "Conversion and reboot" on page 31.

Note: It is strongly recommended that you do not rely on `vxvmconvert`'s mechanisms for unmounting file systems. Conversion will be simpler if you close applications, and unmount file systems before running `vxvmconvert`.

To unmount a file system, use the following command:

```
# umount file-system
```

Conversion and reboot

During conversion, after the analysis phase is complete, the disks to be converted are deemed to be conversion ready. The `vxvmconvert` program asks if you are ready to commit to the conversion changes. If you choose to complete the conversion, the system will try to unmount all of the associated mounted file systems, stop and export the volume group, and then install the VxVM configuration.

If `vxvmconvert` is unable to stop and export volume groups or unmount file systems, the conversion cannot be completed without rebooting the system. You will have the option of aborting the conversion or completing the conversion by rebooting the system. If you choose to reboot, `vxvmconvert` will trigger the completion of the conversion automatically, during reboot, when it can be guaranteed that no processes have access to the volumes that are being converted.

If you choose to abort rather than reboot to complete the conversion, `vxvmconvert` will return to the main menu.

Note: The LVM logical volumes to be converted must all be available to the `vxvmconvert` process. You should not deactivate the volume group or any logical volumes before running `vxvmconvert`.

To activate a volume group

If you are not certain if the LVM volumes or the corresponding volume groups are active, you can activate them with the following command:

```
# vgchange -a y vol_grp_name
```

Converting a volume group

To do the actual conversion of LVM volume groups to VxVM disk groups, choose option 2 of the `vxvmconvert` utility.

`vxvmconvert` will prompt for a name for the VxVM disk group that will be created to replace the LVM volume group you are converting. This is the only object naming that is done through `vxvmconvert`. Additional details are available on modifying VxVM volume names.

See “Tailoring your VxVM configuration” on page 33.

As described earlier, the volume groups selected for conversion are analyzed to ensure that conversion is possible.

See "Analyzing an LVM volume group to see if conversion is possible" on page 26.

After a successful analysis phase, `vxvmconvert`prompts you to commit to the change or abort the conversion. When you select to commit to conversion, the new VxVM metadata is written. Before the conversion is committed, the `vxvmconvert` operation displays the estimated required time as:

```
VxVM INFO V-5-2-4906
The expected time for convert is: 0 hrs 0 mins 7 secs.
```

Note: The time required for conversion is an estimate and is not a calculated time. The actual conversion time may differ depending on factors like number of CPUs, Memory, I/O throughput etc.

More information is available on the details of the conversion process.

See "Examples" on page 37.

Taking actions if conversion fails

There are several reasons why conversion can fail.

See "Volume group conversion limitations" on page 22.

Messages from `vxvmconvert` explain the type of failure, and any actions you can take before retrying the conversion.

Complete details of specific error messages are available.

See "About conversion error messages" on page 79.

Implementing changes for new VxVM logical volume names

You must be sure that all applications and configuration files refer properly to the new VxVM logical volumes.

See "Planning for new VxVM logical volume names" on page 29.

Restarting applications on the new VxVM volumes

After the conversion to VxVM is complete, file systems can be mounted on the new devices and applications can be restarted.

If you unmounted file systems before you ran `vxvmconvert`, you need to remount them by the new volume names. `vxvmconvert` updated `/etc/fstab` with the new

names. When you started vxvmconvert, you may have left file systems mounted that are associated with the volumes you converted. vxvmconvert remounts these with the new VxVM volume names.

Tailoring your VxVM configuration

vxvmconvert provides a default name for naming the newly formed VxVM disk group during conversion only as an option. However, you will be given the choice of choosing your own VxVM disk group name. By default, vxvmconvert renames the LVM volume group by replacing the prefix vg in the volume group name with the prefix dg. For example, vg08 would become dg08. If there is no vg in the LVM volume group name, vxvmconvert simply uses the same volume group name for its disk group.

The disks in the new VxVM disk group are given VxVM disk media names () based on this disk group name. Additional information is available on VxVM disk media names.

See vxintro(1M).

If your new VxVM disk group is dg08, it will have VxVM disks with names like dg0801, dg0802, etc. The VxVM plexes within the logical volumes will be dg0801-01, dg0801-02, etc.

If you do not like the default object names generated by the conversion, use the standard VxVM utilities to rename these objects. See the rename option in the vxedit(1M) man page for more details on renaming the disk groups.

Note: You must only rename objects in the VxVM configuration after you are fully satisfied with that configuration. In particular, you should never use menu option 3 of vxvmconvert (Roll back) after name changes. If you have chosen to set up symbolic links to the VxVM volumes, avoid renaming VxVM objects. Additional information is available on setting up symbolic links. These symbolic links are made invalid if the underlying VxVM device node name changes.

See “Planning for new VxVM logical volume names” on page 29.

Restoring the LVM volume group configuration

In some circumstances, you may need to restore the LVM configuration that existed before you converted to VxVM with vxvmconvert. For example:

- If something went wrong during the conversion, such as a system crash or a disk crash that caused the conversion to be unworkable.

- If during a conversion only some of a set of volume groups converted successfully, then you may want to restore the LVM configuration for the entire set.

It is possible to restore the original LVM configuration in one of two ways, but both have limitations and restrictions. The method you use depends on if any changes have been made to the VxVM configuration since the conversion occurred. Any of the following actions changes the VxVM configuration:

- adding or removing disks
- adding or removing volume groups
- changing the names of VxVM objects

Restoration methods include:

- rollback using `vxvmconvert`

  Use rollback only if the VxVM configuration has not changed since the conversion. This method restores the LVM configuration without the need for user data restoration. Additional information is available on the rollback method.

  See "Rollback to LVM using vxvmconvert" on page 35.

- restore user data using `vgrestore` and `frecover`

  This method is a full LVM restoration which is used to restore your user data from backup when the VxVM configuration has changed since the conversion was made. First of all, this method restores the original LVM configuration information (`vgrestore`), and then restores the original user data from the backup that was made before the conversion was done (`frecover`). Additional information is available on full LVM restoration.

  See "Full LVM restoration" on page 36.

Note: Restoring user data using the `vgrestore` and `frecover` method will result in the loss of all user data changes made since the conversion, and the loss of all new volumes created since the conversion.

In other words, this method of restoring data will take you back to exactly where you were before the conversion was done.

However, if no new volumes have been created, and if none of the existing volumes have been resized, you can use the `vxvmconvert rollback` option to restore the original LVM configuration. If you use this method, any user data changes made since the conversion will be retained, and you will not need to carry out a user level data restore (`frecover`).

The name changes that `vxvmconvert` makes as part of the conversion are managed by rollback, and do not count as VxVM configuration changes for the purposes of choosing a restoration method.

The `vgrestore` command should not be confused with the LVM command, `vgcfgrestore`. `vgcfgrestore` is used to restore the LVM configuration information saved by `vgcfgbackup`, but it will not restore your device files and `/etc/fstab` entries. It also will not import and activate the volume group, nor will it clean up any VxVM information left around. However, `vgrestore` will do all of this for you.

Rollback to LVM using vxvmconvert

Rollback replaces the VxVM disk groups with the original LVM volume groups. During conversion, vxvmconvert saves a “snapshot” of the original LVM metadata and associated configuration files, such as /etc/fstab and LVM device files. It restores only the LVM metadata and configuration files from this snapshot; user data is not changed. This method can only be used if no changes have been made to the configuration since the conversion.

For example, if a disk has been added to the disk group or if the names of any logical volumes have changed, you cannot use the rollback method.

Note: In many cases, if you choose the rollback method and the configuration has changed, you receive an error and must use the full restore method..

If you used the workaround of creating symbolic links from the old LVM names to the new VxVM names, you must remove the symbolic links you created before beginning the rollback. Additional information is available on creating symbolic links.

See "Planning for new VxVM logical volume names" on page 29.

This "snapshot" is kept on the root file system. The presence of this snapshot should not be taken as assurance that full off-line backups will not be needed. Specific information is available on backups.

See "Backing up your LVM configuration and user data" on page 27.

To rollback to LVM from the VxVM conversion, run `vxvmconvert` and choose option 3. An illustration of this method is available.

See "Example: VxVM to LVM rollback" on page 52.

Warning: Do not use this option unless you are certain that you want to restore LVM volume groups. After this is run, the VxVM disks that were created as a result of the original conversion from LVM to VxVM no longer exists. This option is not a full complement to `vxvmconvert`. It simply writes the saved LVM metadata back on top of the disks. Those data can only be considered valid for the period of time when the logical volumes are offline. If the VxVM configuration has been brought online, the metadata in the rollback snapshot should be considered obsolete. Specific information on full LVM restoration is available.

See "Full LVM restoration" on page 36.

Full LVM restoration

If you need to restore the original LVM configuration, but changes have been made to the VxVM configuration, you cannot use the rollback option of `vxvmconvert`. In this case, you must restore the user data in addition to restoring the old LVM metadata and associated configuration files. You may need to use this method if the disks in use by the LVM/VxVM volumes were corrupted during or after conversion.

Note: The snapshot of LVM internal data is kept on the root file system.

To use this method, you must have backed up data located on all the volume groups' logical volumes before conversion to VxVM.

Restoration of LVM volume groups is a two-step process consisting of a restoration of LVM internal data (metadata and configuration files), and restoration of user or application data.

The process is limited to restoring the state of the logical volumes as they existed before conversion to VxVM disks. If the data has changed on the volumes during the time they were VxVM volumes, those changes are lost after you restore the LVM configuration and saved user data.

To do a full restoration of the original LVM configuration, do the following:

1. Use `vgrestore` to restore LVM internal data.

   ```
   # vgrestore vol_grp_name
   ```

2. Use the recovery method to restore user or application data. In preparation for conversion, the recovery method should have been done with the standard backups you made in preparation for conversion. The following example shows an `frecover` from the `fbackup` example. Additional information on the `fbackup` example is available.

   See "Backing up your LVM configuration and user data" on page 27.

   ```
   # mount -F vxfs /dev/vg01/lvol3 /foodir
   # frecover -r -f /dev/rmt/c0t0d0BEST
   ```

Examples

The following are annotated examples on how to use `vxvmconvert` command.

Example: displaying the vxvmconvert menu

To display the `vxvmconvert` menu, use the following command:

```
# vxvmconvert
```

The following menu is displayed:

```
Volume Manager Support Operations
Menu: Volume Manager/LVM_Conversion

1       Analyze LVM Volume Groups for Conversion
2       Convert LVM Volume Groups to VxVM
```

```
3      Roll back from VxVM to LVM
list   List disk information
listvg List LVM Volume Group information
 ?     Display help about menu
 ??    Display help about the menuing system
 q     Exit from menus
```

Example: listing disk information

The list option of `vxvmconvert` displays information about the disks on a system. Select the `list` option from the `vxvmconvert` Main Menu:

```
Menu: Volume Manager/LVM_Conversion/list

# list

Use this menu option to display a list of disks. You can also choose
to list detailed information about a disk by entering a specific disk
device address.

Enter disk device or "all" [<address>,all,q,?](default: all) x

DEVICE        DISK       GROUP       STATUS
disk10        -          -           online
disk11        -          -           LVM
disk12        -          -           LVM
disk13        mydg01     mydg        online
disk14        -          -           online

Device to list in detail [<address>,none,q,?] (default: none)
none
```

Example: listing LVM volume group information

To list LVM volume group information, use the `listvg` option of `vxvmconvert`.

Select the `listvg` option from the `vxvmconvert` Main Menu:

```
Menu: Volume Manager/LVM_Conversion/ListLVMVolumeGroups
# listvg
Use this menu option to display a list of LVM volume groups. You
can also choose to list detailed information about the LVM
volume groups at a specific disk device address.
Select the Volume Group as follows:
Enter Volume Group (i.e.- vg08) or "all" [<address>,all,q,?]
     default: all)
```

```
LVM VOLUME GROUP INFORMATION
NAME   VERSION    TYPE     PHYSICAL VOLUME
vg00   1.0        ROOT     disk10
vg09   2.0        Non-Root disk11
vg08   1.0        Non-Root disk12

Volume Group to list in detail
[<address>,none,q,?] (default: none) none
To display detailed information about a volume group, select any
of the volume groups from the above list.
Volume Group to list in detail \
   [<address>,none,q,?] (default: none) vg08

--- Volume groups ---
VG Name                            /dev/vg08
VG Write Access                    read/write
VG Status                          available
Max LV                             255
Cur LV                             2
Open LV                            2
Max PV                             16
Cur PV                             1
Act PV                             1
Max PE per PV                      1016
VGDA                               2
PE Size (Mbytes)                   4
Total PE                           250
Alloc PE                           250
Free PE                            0
Total PVG                          0

--- Logical volumes ---
   LV Name                            /dev/vg08/lvol1
   LV Status                          available/syncd
   LV Size (Mbytes)                   500
   Allocated PE                       125
   Used PV                            1

--- Physical Extent ---
  LV Name                             /dev/vg08/lvol2
  LV Status                           available/syncd
  LV Size (Mbytes)                    500
```

```
  Current LE                          125
  Allocated PE                        125
  Used PV                             1

--- Physical volumes ---
 PV Name                              /dev/disk/disk12
 PV Status                            available
 Total PE                             250
 Free PE                              0
List another LVM Volume Group? [y,n,q,?] (default: n)
Select an operation to perform:
```

Note: The volume groups you want to convert must not be a root volume group or have bootable volumes in the group.

Example: analyzing LVM volume groups

To analyze one or more LVM volume groups:

```
# vxvmconvert

Volume Manager Support Operations
Menu: VolumeManager/LVM_Conversion
1      Analyze LVM Volume Groups for Conversion
2      Convert LVM Volume Groups to VxVM
3      Roll back from VxVM to LVM
list   List disk information
listvg List LVM Volume Group information
 ?     Display help about menu
 ??    Display help about the menuing system
 q     Exit from menus

Select an operation to perform: 1
Analyze one or more LVM Volume Groups
Menu: Volume Manager/LVM_Conversion/Analyze_LVM_VGs
Use this operation to analyze one or more LVM volume groups for
possible conversion using the VxVM Volume Manager. This
operation checks for problems that would prevent the conversion
from completing successfully. It calculates the space required
to add the volume groups disks to a Volume Manager disk group,
and to replace any existing partitions and volumes with Volume
Manager volumes, plexes, and sub-disks.
```

```
More than one volume group or pattern may be entered at the
prompt.
Here are some LVM volume group selection examples:
all:    analyze all LVM Volume Groups (all except Root VG)
listvg: list all LVM Volume Groups
list:   list all disk devices
vg_name:a single LVM Volume Group, named vg_name
<pattern>:   for example vg08 vg09 vg05

Select volume groups to analyze:
[<pattern list>,all,list,listvg,q,?] vg08
Name a new disk group [<group>,list,q,?] (default: dg08)
Each volume group will be analyzed one at a time. If there are
any in this list that you do not want to analyze, you can either
abort now or wait until a later time when you will be given an
opportunity to skip the analysis of any group(s) in this list.
The following disk has been found in the vg08 volume group and
will be analyzed for VxVM conversion.
  disk12
To allow analysis, a new VxVM disk group, dg08, will be
fabricated and the disk device disk12 will be added to the disk
group with the disk name dg0801.
The disk12 disk has been configured for conversion.
The first stage of the Analysis process has completed
successfully.
Second Stage Conversion Analysis of vg08
Analysis of vg08 found sufficient Private Space for conversion
Conversion Analysis of disk12 indicates that the Volume Group is
still in use, which may prevent the completion of the conversion
without having to reboot the system. You may want to double
check that none of the volumes in the volume group are in use
before continuing with the conversion.
Volume Group Analysis Completed
Hit RETURN to continue.
```

Example of a failed analysis:

```
# vxvmconvert

Volume Manager Support Operations
Menu: VolumeManager/LVM_Conversion
1      Analyze LVM Volume Groups for Conversion
2      Convert LVM Volume Groups to VxVM
3      Roll back from VxVM to LVM
```

```
list   List disk information
listvg List LVM Volume Group information
 ?     Display help about menu
 ??    Display help about the menuing system
 q     Exit from menus

Select an operation to perform: 1
Analyze one or more LVM Volume Groups
Menu: Volume Manager/LVM_Conversion/Analyze_LVM_VGs
Use this operation to analyze one or more LVM volume groups for
possible conversion using the VxVM Volume Manager. This
operation checks for problems that would prevent the conversion
from completing successfully. It calculates the space required
to add the volume groups disks to a Volume Manager disk group,
and to replace any existing partitions and volumes with Volume
Manager volumes, plexes, and sub-disks.
More than one volume group or pattern may be entered at the
prompt.
Here are some LVM volume group selection examples:

all:    analyze all LVM Volume Groups (all except Root VG)
listvg: list all LVM Volume Groups
list:   list all disk devices
vg_name:a single LVM Volume Group, named vg_name
<pattern>:   for example vg08 vg09 vg05

Select Volume Groups to analyze :
[<pattern-list>,all,list,listvg,q,?] vg08

Name a new disk group [<group>,list,q,?] (default: dg08)

The following disk has been found in the vg08 volume group and
will be analyzed for VxVM conversion.
  disk12
To allow analysis, a new VxVM disk group, dg08, will be
fabricated and the disk device disk12 will be added to the disk
group with the disk name dg0801.
The disk12 disk has been configured for conversion.
The first stage of the Analysis process has completed
successfully.
Second Stage Conversion Analysis of vg08

Analysis of vg08 found insufficient Private Space for conversion
```

```
SMALLEST VGRA space =  176
RESERVED space sectors =   78
PRIVATE SPACE/FREE sectors =   98
AVAILABLE sector space =   49
AVAILABLE sector bytes =   50176
RECORDS neededs to convert = 399
MAXIMUM records allowable = 392

The smallest disk in the Volume Group (vg08) does not have
sufficient private space for the conversion to succeed. There is
only enough private space for 392 VM Database records and the
conversion of Volume Group (vg08) would require enough space to
allow 399 VxVM Database records. This would roughly translate to
needing an additional 896 bytes available in the private space.
This can be accomplished by reducing the number of volumes in
the (vg08) Volume Group, and allowing that for every volume
removed, the number of Database records required would be
reduced by three. This is only a rough approximation, however.

Hit RETURN to continue.
```

Example: converting LVM volume groups to VxVM disk groups

To convert LVM volume groups to VxVM disk groups:

```
# vxvmconvert

Volume Manager Support Operations
Menu: VolumeManager/LVM_Conversion
1      Analyze LVM Volume Groups for Conversion
2      Convert LVM Volume Groups to VxVM
3      Roll back from VxVM to LVM
list   List disk information
listvg List LVM Volume Group information
 ?     Display help about menu
 ??    Display help about the menuing system
 q     Exit from menus

Select an operation to perform: 2
Convert one or more LVM Volume Groups
Menu: VolumeManager/LVM_Conversion/Convert_LVM_VGs
Use this operation to convert one or more LVM Volume Groups to
```

```
one or more VxVM disk groups. This adds the disks to a disk group
and replaces existing partitions with volumes. LVM-VxVM Volume
Group conversion may require a reboot for the changes to take
effect. For this release, only Non-root LVM Volume Groups are
allowed to be converted.
More than one Volume Group or pattern may be entered at the
prompt.
Here are some LVM Volume Group selection examples:
all:    analyze all LVM Volume Groups (all except Root VG)
listvg: list all LVM Volume Groups
list:   list all disk devices
vg_name:a single LVM Volume Group, named vg_name
<pattern>:   for example vg08 vg09 vg05

Select Volume Groups to convert :
[<pattern-list>,all,list,listvg,q,?] listvg

LVM VOLUME GROUP INFORMATION
NAME     VERSION     TYPE       PHYSICAL VOLUME

vg00     1.0         ROOT       disk10
vg05     1.0         Non-Root   disk11
vg03     2.0         Non-Root   disk14 disk15
vg08     1.0         Non-Root   disk12
Select Volume Groups to convert :
[<pattern-list>,all,list,listvg,q,?] vg08
  vg08
Convert this Volume Group? [y,n,q,?] (default: y)
Name a new disk group [<group>,list,q,?] (default: dg08)
The following disk has been found in the vg08 volume group and
will be configured for conversion to a VxVM disk group.
  disk12
A new disk group dg08 will be created and the disk device disk12
will be converted and added to the disk group with the disk name
dg0801.
The disk12 disk has been configured for conversion.

The first stage of the conversion operation has completed
successfully. If you commit to the changes hereafter, the system
will attempt to umount all of the associated file systems, stop
and export each Volume Group, and then attempt to complete the
conversion without having to reboot the system. If we are unable
to stop and export any of the Volume Groups, then the conversion
```

```
process will not be able to complete without a reboot.You would
then be given the choice to either abort the conversion, or
finish the conversion by rebooting the system.

The conversion process will update the /etc/fstab file so that
volume devices are used to mount the file systems on this disk
device. You will need to update any other references such as
backup scripts, databases,or manually created swap devices. If
you do not like the default names chosen for the corresponding
logical volumes, you may change these to whatever you like using
vxedit.

Second Stage Conversion Analysis of vg08
Analysis of vg08 found sufficient Private Space for conversion

Conversion Analysis of disk12 indicates that the Volume Group is
still in use, which may prevent the completion of the
conversion without having to robot the system. You may want to
double check that none of the volumes in the volume group are
in use before continuing with the conversion.

Hit RETURN to continue.
VxVM  INFO V-5-2-4906
The expected time for convert is: 0 hrs 0 mins 7 secs.

Are you ready to commit to these changes? [y,n,q,?] (default: y)
Saving LVM configuration records for Volume Group vg08

Volume Group configuration for /dev/vg08 has been saved in
/etc/vx/reconfig.d/vgrecords/vg08/vg08.backup
Beginning the export process on Volume Group "/dev/vg08".
Volume group "/dev/vg08" is still active.
/dev/disk/disk12
/dev/vg08/lv1 will convert to /dev/vx/dsk/dg08/dg08lv1
/dev/vg08/rlv1 will convert to /dev/vx/rdsk/dg08/dg08lv1
/dev/vg08/lv2 will convert to /dev/vx/dsk/dg08/dg08lv2
/dev/vg08/rlv2 will convert to /dev/vx/rdsk/vg08dg/dg08lv2
LVM Volume Group vg08 Records Saved
Unmounting vg08 file systems
Volume group "/dev/vg08" has been successfully changed.

The Volume Manager is now reconfiguring (partition phase)...
Volume Manager: Initializing disk12 as a converted LVM disk.
```

```
The system reconfiguration will now be done without rebooting.
The Volume Manager is now reconfiguring (initialization
phase)...
Volume Manager: Adding dg0801 (disk12) as a converted LVM disk.
Adding volumes for disk12...
Starting new volumes...
Updating /etc/fstab...
The system will now Convert the LVM Volume Groups over to VxVM
disk groups.
Convert other LVM Volume Groups? [y,n,q,?] (default: n)
```

Example of a failed conversion:

```
# vxvmconvert

Volume Manager Support Operations
Menu: VolumeManager/LVM_Conversion

1      Analyze LVM Volume Groups for Conversion
2      Convert LVM Volume Groups to VxVM
3      Roll back from VxVM to LVM
list   List disk information
listvg List LVM Volume Group information
 ?     Display help about menu
 ??    Display help about the menuing system
 q     Exit from menus

Select an operation to perform: 2
Convert one or more LVM Volume Groups
Menu: VolumeManager/LVM_Conversion/Convert_LVM_VGs
Use this operation to convert one or more LVM Volume Groups to
one or more VxVM disk groups. This adds the disks to a disk group
and replaces existing partitions with volumes. LVM-VxVM Volume
Group conversion may require a reboot for the changes to take
effect. For this release, only Non-root LVM Volume Groups are
allowed to be converted.
More than one Volume Group or pattern may be entered at the
prompt.
Here are some LVM Volume Group selection examples:
all:    analyze all LVM Volume Groups (all except Root VG)
listvg: list all LVM Volume Groups
list:   list all disk devices
vg_name:a single LVM Volume Group, named vg_name
<pattern>:   for example vg08 vg09 vg05
```

```
Select Volume Groups to convert :
[<pattern-list>,all,list,listvg,q,?] listvg
LVM VOLUME GROUP INFORMATION
NAME     VERSION     TYPE       PHYSICAL  VOLUME
vg00     1.0         ROOT       disk10
vg05     1.0         Non-Root   disk11
vg03     2.0         Non-Root   disk14 disk15
vg08     1.0         Non-Root   disk12
Select Volume Groups to convert :
[<pattern-list>,all,list,listvg,q,?] vg08
  vg08
Convert this Volume Group? [y,n,q,?] (default: y)
Name a new disk group [<group>,list,q,?] (default: dg08)

The following disk has been found in the vg08 volume group and
will be configured for conversion to a VxVM disk group.
  disk12
A new disk group dg08 will be created and the disk device disk12
will be converted and added to the disk group with the disk
name dg0801.
The disk12 disk has been configured for conversion.
The first stage of the conversion operation has completed
successfully. If you commit to the changes hereafter, the system
will attempt to unmount all of the associated file systems, stop
and export each Volume Group, and then attempt to complete the
conversion without having to reboot the system. If we are unable
to stop and export any of the Volume Groups, then the conversion
process will not be able to complete without a reboot. You would
then be given the choice to either abort the conversion, or
finish the conversion by rebooting the system.

The conversion process will update the /etc/fstab file so that
volume devices are used to mount the file systems on this disk
device. You will need to update any other references such as
backup scripts, databases, or manually created swap devices. If
you do not like the default names chosen for the corresponding
logical volumes, you may change these to whatever you like using
vxedit.

Second Stage Conversion Analysis of vg08

Analysis of vg08 found insufficient Private Space for conversion
SMALLEST VGRA space              =   176
```

```
RESERVED space sectors =   78
PRIVATE SPACE/FREE sectors =   98
AVAILABLE sector space =   49
AVAILABLE sector bytes = 50176
RECORDS neededs to convert= 399
MAXIMUM records allowable   = 392

The smallest disk in the Volume Group (vg08) does not have
sufficient private space for the conversion to succeed. There is
only enough private space for 392 VM Database records and the
conversion of Volume Group (vg08) would require enough space to
allow 399 VxVM Database records. This would roughly translate to
needing an additional 896 bytes available in the private space.
This can be accomplished by reducing the number of volumes in
the (vg08) Volume Group, and allowing that for every volume
removed, the number of Database records required would be
reduced by three. This is only a rough approximation, however.
Hit RETURN to continue.
```

Example: converting LVM verion 2 volume groups to VxVM disk groups

To convert LVM volume groups to VxVM disk groups:

```
#vxvmconvert

Volume Manager Support Operations
Menu: VolumeManager/LVM_Conversion
1      Analyze LVM Volume Groups for Conversion
2      Convert LVM Volume Groups to VxVM
3      Roll back from VxVM to LVM
list   List disk information
listvg List LVM Volume Group information
 ?     Display help about menu
 ??    Display help about the menuing system
 q     Exit from menus

Select an operation to perform:listvg
List LVM Volume Group information
Menu: VolumeManager/LVM_Conversion/ListLVMVolumeGroups

  Use this menu operation to display a list of LVM volume Groups.  You can
  also choose to list detailed information about the LVMVG at
  a specific disk device address.
```

```
Enter Volume Group (i.e.- vg04) or "all"
[<address>,all,q,?] (default:all)

LVM VOLUME GROUP INFORMATION
Name         Version    Type         Physical Volumes

vg00         1.0        ROOT         c2t0d0
vg01         2.0        Non-Root     c9t3d4


Volume Group to list in detail
[<address>,none,q,?] (default:none)

Volume Manager Support Operations
Menu: VolumeManager/LVM_Conversion

 1      Analyze LVM Volume Groups for Conversion
 2      Convert LVM Volume Groups to VxVM
 3      Roll back from VxVM to LVM
 list   List disk information
 listvg List LVM Volume Group information


 ?      Display help about menu
 ??     Display help about the menuing system
 q      Exit from menus

Select an operation to perform:2

Convert one or more LVM Volume Groups
Menu: VolumeManager/LVM_Conversion/Convert_LVM_VGs

  Use this operation to convert LVM Volume Groups to VxVM disk groups.
  This adds the disks to a disk group and replaces existing LVM volumes
  with VxVM volumes.  LVM-VxVM Volume Group conversion may require a
  reboot for the changes to take effect. For this release, only
  Non-root LVM Volume Groups can be converted.

  More than one Volume Group or pattern may be entered at the prompt.
  Here are some LVM Volume Group selection examples:

  all:          convert all LVM Volume Groups (all except Root VG)
```

```
  listvg:       list all LVM Volume Groups
  list:         list all disk devices
  vg_name:      a single LVM Volume Group, named vg_name
  <pattern>:    for example:  vg04 vg08 vg09

Select Volume Groups to convert :
[<pattern-list>,all,list,listvg,q,?]vg01
 Convert LVM Volume Group vg01

Convert this Volume Group? [y,n,q,?] (default:y)

VxVM  NOTICE V-5-2-0
LVM Volume Group vg01 reports a version 2.0
LVM Volume Group other than version 1.0 is not supported for conversion.

Hit RETURN to continue.
```

Example: list, listvg, and vxprint outputs of an LVM volume group before and after conversion

The examples below show the `vxvmconvert listvg`, `list`, and `vxprint` output for an LVM volume group `vg08` converted to a VxVM disk group `dg08`.

Example of vxvmconvert listvg output before conversion of volume group vg08

```
LVM VOLUME GROUP INFORMATION
NAME      VERSION     TYPE       PHYSICAL VOLUME
vg00      1.0         ROOT       c0t5d0
vg08      1.0         Non-Root   c0t8d0
vg09      2.0         Non-Root   c0t9d0
```

Example of the vxvmconvert list output which shows the disk devices on a system

```
DEVICE     DISK     GROUP     STATUS
c0t5d0     -        -         online
c0t8d0     -        -         LVM
c0t9d0     -        -         LVM
c0t10d0    disk01   rootdg    online
c0t11d0    -        -         online
Disk group: rootdg
```

What does vxvmconvert list display?

The device indicates a physical disk, a disk with a name indicates if the disk is under VxVM control, a group shows the disk group name, and the status indicates if it is an LVM disk. If the status is online, that means VxVM acknowledges the disk but doesn't have it under its control.

Example vxprint output before conversion

The `list` and `listvg` output is from within the `vxvmconvert` command. `vxprint` is a command line command.

```
TY NAME       ASSOC   KSTATE  LENGTH  PLOFFS  STATE TUTIL0   PUTIL0
dg rootdg     rootdg      -       -      -      -      -       -
dm disk01     c0t10d0     -   2079468    -      -      -       -
```

Example of the listvg output after conversion of vg08

```
LVM VOLUME GROUP INFORMATION
Name       Version      Type       Physical Volumes
vg00       1.0          ROOT       c0t5d0
vg09       2.0          Non-Root   c0t9d0
ourvg      2.0          Non-Root   c5t3d5
Volume Group to list in detail: None
```

Note: Note that `vg08` is no longer listed under LVM information.

Example of the vxvmconvert list output after conversion of volume group vg08 to dg08

```
DEVICE   DISK   GROUP   STATUS
c0t5d0   -      -       online invalid
c0t8d0   dg0801 dg08    online
c0t9d0   -      LVM
c0t10d0  disk01 rootdg  online
c0t11d0  -      -       online
Disk group: rootdg
```

Example of the vxprint output after conversion

```
TY NAME       ASSOC      KSTATE  LENGTH  PLOFFS STATE  TUTILO PUTIL0
dg dg08       dg08       -       -       -      -      -      -
dm dg0801     c0t8d0     -       2080768 -      -      -      -
```

```
v  dg08lv1    fsgen      ENABLED 102400  -      ACTIVE -      -
pl dg08lv1-01 dg08lv1    ENABLED 102400  -      ACTIVE -      -
sd dg0801-01  dg08lv1-01 ENABLED 102400  -      -      -      -
```

The `vxprint` output provides the following information:

- The disk group `dg08` contains the VxVM disk `dg0801` and the volume `dg08lv1`. The VxVM disk `dg0801` is associated with disk device `c0t8d0` and is 2080768 blocks in length. The volume `dg08lv1` is of type `fsgen`, is enabled in the VxVM kernel driver, is of `length` 102400, and is in the ACTIVE state. This means that the volume is started, and the plex is enabled. Operations to the volume such as recovery and data access will be governed by the usage type `fsgen`.
- The plex `dg08lv1-01` is associated with volume `dg08lv1`, and maps the entire address range of the volume. Associated with the plex is one subdisk, `dg0801-01` which maps the plex address range from 0 to the entire length of the plex, i.e. 102400 blocks. As implied by the root of its name, the subdisk `dg0801-01` uses an extent from the VxVM disk `dg0801`.

Example: VxVM to LVM rollback

```
Volume Manager Support Operations
Menu: VolumeManager/LVM_Conversion

1      Analyze LVM Volume Groups for Conversion
2      Convert LVM Volume Groups to VxVM
3      Roll back from VxVM to LVM
list   List disk information
listvg List LVM Volume Group information
?      Display help about menu
??     Display help about the menuing system
q      Exit from menus

Select an operation to perform: 3
Rollback one or more LVM Volume Groups
Menu: VolumeManager/LVM_Conversion/Rollback_LVM_VG
```

Use this operation to rollback from a conversion of an LVM Volume Group. This operation will tear down the VxVM disk group and recreate the LVM volume group in its original form. User data is untouched by rollback.

Warning: If the VxVM configuration has changed since the conversion you should NOT use this operation. Additional information is available on restoring configuration.

See "Restoring the LVM volume group configuration" on page 33.

```
More than one Volume Group or pattern may be entered at the
prompt. Here are some LVM Volume Group selection examples:
all:    Rollback all converted LVM Volume Groups
listvg: list all LVM Volume Groups eligible for rollback
list:   list all disk devices
vg_name:   a single LVM Volume Group, named vg_name
<pattern>:   for example vg08 vg09 vg05

Select Volume Group(s) to rollback :
[<pattern-list>,all,list,listvg,q,?]
vg08
Roll back this Volume Group? [y,n,q,?] (default: y)
Rolling back LVM configuration records for Volume Group vg08
Selected Volume Groups have been restored.
Hit RETURN to continue.
Rollback other LVM Volume Groups? [y,n,q,?] (default: n)
```

General information regarding conversion speed

The speed of the process of converting an existing LVM volume group to a similar VxVM disk group is largely dependent upon the size of the volume group being converted, as well as on the complexity of the volumes within that volume group.

Note: The `vxvmconvert` operation displays an estimated time before the actual conversion is committed.

Factors affecting conversion speed include:

- Size of volume groups. The larger the volume groups, the larger the LVM metadata on each disk. A copy must be made of the LVM metadata for each physical disk. Some areas are greater than 2MB; therefore, a 50-disk volume group requires 50 2MB reads and writes (i.e., 100 large I/Os) to complete.
- Individual size of a logical volume in a volume group, and the complexity of the logical volume layout. For example, for a system with 50 9GB drives, a simple 50GB logical volume of the first 5 1/2 disks can be created. But a 50GB striped logical volume that takes the first 1GB of all 50 disks can also be created.

The first and simple logical volume takes less time to convert than the striped volume. However, for the striped volume, 50 disks need to be checked. Also, the complexity of reproducing the VxVM commands to set up the striped volumes requires more VxVM commands to be generated to represent more smaller sub-disks representing the same amount of space.

Another factor in converting stripes is that stripes create more work for the converter. In some cases, stripes require 1GB volume, although only the metadata is being changed. In other cases, where there are more physical disks in one volume than another, there is more metadata to deal with. The converter has to read every physical extent map to ensure there are no holes in the volume; if holes are found, the converter maps around them.

- Number of volumes. While it takes longer to convert one 64GB volume than one 2GB volume, it also takes longer to convert 64 1GB volumes than one 64GB volume, providing that the volumes are of similar type.
- Mirrored volumes. Mirrored volumes typically do not take more time to convert than simple volumes. Volumes that are mirrored and striped at the same time would take longer, but LVM currently does not allow this.

  Currently, after conversion, mirrored volumes are not automatically synchronized because a large mirror could take hours to complete.

  For example, in tests, a 150GB volume group consisting of 20 simple logical volumes takes approximately 35-40 minutes to convert. In contrast, the same volume group (150GB) consisting of mirrored volumes that need to be synchronized can take 30-40 hours to convert.

Note: If you convert mirrored volumes, you must synchronize them in a separate step.

Non-interactive conversion of volume groups

The `vxvmconvert` utility is an interactive command. You can also use the `vxautoanalysis` command and the `vxautoconvert` command to perform non-interactive analysis and conversion of LVM volume groups. The `vxautorollback` command also lets you reverse the conversion, and turn a VxVM disk group back into an LVM volume group.

Note: This release only supports the conversion of LVM version 1 volume groups to VxVM. It does not support the conversion of LVM version 2 volume groups.

Analyzing volume groups for conversion

The `vxautoanalysis` utility analyses one or more LVM volume groups, and reports on their suitability for conversion to VxVM disk groups. Only analysis of the suitability of the volume groups for conversion is performed. The actual conversion can be performed by using the `vxautoconvert` utility or the `vxvmconvert` utility.

Note: The VxVM configuration daemon (`vxconfigd`) must be running in order for the analysis to succeed.

To analyze volume groups for conversion

- Run the `vxautoanalysis` command:

  ```
  # /usr/sbin/vxautoanalysis [-f] [vgname ...]
  ```

  The volume groups may be specified by their names or full pathnames. If no volume groups are specified, analysis of all volume groups on the system is attempted.

  If the value of the system tunable, `nproc`, is too low, the analysis will report that the conversion analysis of the volume groups cannot be performed in parallel. In that case, you can use the `-f` option to specify that the volume groups are to be analyzed one at a time.

  See the `vxautoanalysis`(1M) manual page.

Converting volume groups to disk groups

The `vxautoconvert` utility converts one or more LVM volume groups to VxVM disk groups. Any LVM extent-based striped volumes are converted to stripe-mirror (non-layered) VxVM volumes.

Note: The VxVM configuration daemon (`vxconfigd`) must be running in order for the conversion to succeed. Conversion of the root volume group is not permitted. An error results if this is attempted.

To convert volume groups to disk groups

- Run the `vxautoconvert` command:

  ```
  # /usr/sbin/vxautoconvert [-f] [vgname ...]
  ```

  The volume groups may be specified by their names or full pathnames. If no volume groups are specified, conversion of all volume groups on the system is attempted.

  If the value of the system tunable, `nproc`, is too low, the analysis will report that the conversion of the volume groups cannot be performed in parallel. In that case, you can use the `-f` option to specify that the volume groups are to be converted one at a time.

  See the `vxautoconvert`(1M) manual page.

Converting disk groups back to volume groups

The `vxautorollback` utility converts one or more VxVM disk groups back to the LVM volume groups from which they had previously been converted.

Note: The VxVM configuration daemon (`vxconfigd`) must be running in order for the analysis to succeed. Reverse conversion is performed on each disk group in turn. Parallel conversion is not supported. If the configuration or layout of the volumes in the converted disk group has been changed since conversion from LVM to VxVM, reverse conversion is not possible. The records of the converted volume groups are stored in the directory `/etc/vx/reconfig.d/vgrecords`. Reverse conversion is not possible if these records are removed or modified.

To convert disk groups back to volume groups

- Run the `vxautorollback` command:

  ```
  # /usr/sbin/vxautorollback [dgname ...]
  ```

  The disk groups may be specified by their names or full pathnames. If no disk groups are specified, conversion is attempted for all disk groups on the system that have volume group records stored in the directory `/etc/vx/reconfig.d/vgrecords`.

  See the `vxautorollback`(1M) manual page.

Chapter 3

Command differences

This chapter includes the following topics:

About LVM and VxVM command differences

This chapter describes the differences between LVM and VxVM commands, and tasks. It includes a task comparison chart which lists some of the tasks performed using LVM with a near equivalent task performed using VxVM. It also provides a list of VxVM tasks which are not available with LVM, and the LVM features currently not supported in VxVM.

Additional information is available on LVM and VxVM commands.

See *HP-UX Managing Systems and Workgroups*.

See theLVM manual pages in *HP-UX Reference Volumes 2, 3, and 5*.

See the Veritas Volume Manager documentation.

LVM and VxVM command equivalents

The table below lists the LVM commands and a near equivalent command to use in VxVM. For more information, refer to the Task Comparison chart. Additional information is available on VxVM commands.

Refer to the Veritas Volume Manager documentation package.

Table 3-1 Command comparison

LVM	Description/action	VxVM	Description/action

`lvchange`	Changes the characteristics of logical volumes.	`vxedit`
		or
		`vxvol set`
			Creates, removes, and modifies Volume Manager records.

	There is no single equivalent LVM command.	`vxresize`	Resizes a file system and its underlying volume at the same time.
`lvlnboot`	Creates root, primary and secondary swap and dump volumes. It also creates boot areas on the disk.		There is no equivalent command for this release.
`lvcreate`	Creates a logical volume.	`vxassist`	Creates volumes with the `make` parameter. Example: `vxassist make` *`vol_name`* `100M layout=stripe`

`lvextend`	Increases disk space allocated to a logical volume.	`vxassist`	Increases a volume in size with the `growto` or `growby` parameter.
			Example:
			`vxassist growto` *`vol_name`* `200M,`
			
			`vxassist growby` *`vol_name`* `100M`
			`vxassist` creates and modifies volumes.

`lvreduce`	Decreases disk space allocated to a logical volume.	`vxassist`	Decreases a volume in size with the `shrinkto` or `shrinkby` parameters. Example: `vxassist shrinkto` *`vol_name`* `200M` Make sure you shrink the file system before shrinking the volume.

Table 3-1 Command comparison *(continued)*

LVM	Description/action	VxVM	Description/action
`lvremove`	Removes one or more logical volumes from a volume group.	`vxedit` `vxassist`	Removes volumes with the `-rf rm` parameters. Example: `vxedit -rf rm` *`vol_name`* Removes a volume with the `remove volume` parameters. Example: `vxassist remove volume` *`vol_name`*
`lvsplit`	Splits a mirrored logical volume into two logical volumes.	`vxassist snapshot`	The `snapshot` operation takes one of the attached temporary mirrors and creates a new volume with the temporary mirror as its one plex. Example: `vxassist snapshot` *`vol_name new_volume`*
`lvmerge`	Reverses and converts the `lvsplit` logical volumes to a single logical volume.	`vxassist snapback`	The snapback operation returns the snapshot plex to the original volume from which it was snapped. Example: `vxassist snapback new_volume`
`lvsync`	Synchronizes mirrors that are stale in one or more logical volumes.	`vxrecover` `vxvol start`	The `vxrecover` command performs resynchronize operations for the volumes, or for volumes residing on the named disks (medianame or the VxVM name for the disk). Example: `vxrecover` *`vol_name media_name`*

Table 3-1 Command comparison *(continued)*

LVM	Description/action	VxVM	Description/action
`pvcreate`	Makes a disk an LVM disk.	`vxdisksetup`	Brings a disk under VxVM control. Example: `vxdisksetup c0t3d0` Option 1 in the `vxdiskadm` menu adds or initializes one or more disks.
`pvdisplay`	Displays information about physical volumes in a volume group.	`vxdisk list`	Lists information about VxVM disks. Example: `vxdisk list` *`disk_name`*

`pvchange`	Sets physical volume characteristics to allow/deny allocation of additional physical extents from this disk.	`vxdisk`
		`vxdisk set`
		`vxedit`
			The `vxdisk` utility performs basic administrative operations on VxVM disks. Operations include initializing and replacing disks, as well as taking care of some book-keeping necessary for the disk model presented by the Volume Manager.

`pvmove`	Moves allocated physical extents from source to destination within a volume group.	`vxevac`
		`vxsd mv`
		`vxdiskadm`
			Moves volumes off a disk.
			Performs volume operations on a subdisk. Moves the contents of old subdisk onto the new subdisks and replaces old sub disk with the new subdisks for any associations.
			The `vxdiskadm` script presents a menu of possible operations to the user.
			Option 7 in the `vxdiskadm` menu moves volumes.

`pvremove`	Removes the LVM header information and releases the disk from LVM control.	`vxdiskunsetup`	Removes the VxVM header information and releases the disk from VxVM control.
`vgcreate`	Creates a volume group.	`vxdiskadd` `vxdg init`	Creates a new disk group and/or adds disks to a disk group.

Table 3-1 Command comparison *(continued)*

LVM	Description/action	VxVM	Description/action
`vgdisplay`	Displays information on all volume groups.	`vxdg list` `vxprint`	Displays the contents of a disk group. Displays information about all objects or a subset of objects.
`vgchange`	Activates or deactivates one or more volume groups.	*`vxdg -g diskgroup set activation=mode`*	Activates a shared disk group.
`vgextend`	Extends a volume group by adding one or more disks to it.	`vxdiskadd` `vxdiskadm`	Adds a disk to the disk group. Option 1 in the `vxdiskadm` menu adds disks to the disk group.

`vgreduce`	Reduces a volume group by removing one or more disks from it.	`vxdg rmdisk`
		`vxdisk rm`
		`vxdiskadm`
			Removes disks from a disk group.
			Removes the specified disk access record by disk access name.
			Option 3 in the `vxdiskadm` menu removes disks.

`vgscan`	Scans all disks and looks for logical volume groups.	`vxinfo`
		`vxprint`
		`vxdiskadm`
			Displays information about volumes.
			Displays complete or partial information from records in VxVM disk group configurations.
			Option `list` in the `vxdiskadm` menu displays disk information.

`vgsync`	Synchronizes mirrors that are stale in one or more logical volumes.	`vxrecover`	Starts resynchronization and recovery of volumes.
`vgremove`	Removes the definition of a volume group from the system.	`vxdg deport` `vxdiskadm`	Deports a disk group from the system. Option 9 in the `vxdiskadm` menu removes a disk group.
`vgexport`	Removes a volume group from the system.	`vxdg deport` `vxdiskadm`	Deports a disk group from the system. Option 9 in the `vxdiskadm` menu removes a disk group.
`vgimport`	Adds a volume group to the system by scanning physical volumes which have been exported using `vgexport`.	`vxdg import` `vxdiskadm`	Imports a disk group. Option 8 in the `vxdiskadm` menu imports a disk group.
	No LVM command	`vxplex`	Operates on plex objects.

Table 3-1 Command comparison *(continued)*

LVM	Description/action	VxVM	Description/action
`lvchange`, `lvextend`, `lvcreate`, `lvreduce`	Performs operations on logical volumes.	`vxvol`	Operates on volume objects.
	No LVM command	`vxsd`	Operates on subdisk objects.
	No LVM command	`vxmend`	Fixes simple misconfigurations.

Comparison of LVM and VxVM tasks

This section contains a list of tasks which you can perform using LVM, and near equivalent tasks which you can perform using Veritas Volume Manager. You can perform the LVM tasks by using SMH or the command line interface. Similarly, you can choose to perform VxVM tasks by using the Veritas Enterprise Administrator (VEA) or the command line interface. This document focuses on the command line interface.

Note: The following features in VxVM require an additional license: Mirroring, Mirroring and Striping, Dynamic Multipathing of Active/Active Devices, Hot-relocation, Online Migration, and RAID-5.

All the VxVM tasks listed in the task comparison chart can be performed by the Veritas Enterprise Administrator.

See the *Veritas Enterprise Administrator User's Guide*.

Additional information is available on LVM and VxVM commands.

See *HP-UX Managing Systems and Workgroups*.

See theLVM manual pages in *HP-UX Reference Volumes 2, 3, and 5*.

See the Veritas Volume Manager documentation.

Note: Mirroring of a VxVM root disk is supported in this release.

Mirroring in LVM is supported only if you have MirrorDisk/UX already installed as an add-on product. In addition, mirroring in VxVM requires an additional license.

Table 3-2 LVM and VxVM task comparisons

Task type	Description	Example
LVM	Create an LVM disk.	`pvcreate /dev/rdsk/disk_name`
VxVM	Bring a disk under Volume Manager control	`vxdiskadd device_name` Option 1 in the `vxdiskadm` menu adds a disk and initializes it
LVM	Create a volume group	`vgcreate /dev/vol_grp /dev/dsk/disk_name`
VxVM	Create a disk group	`vxdg init disk_group disk_name` Option 1 in the `vxdiskadm` menu performs this task
LVM	Add a new disk to the existing volume group.	`vgextend /dev/vol_grp /dev/dsk/disk_name`
VxVM	Add a new disk to the existing disk group.	`vxdg -g disk_group` `adddisk disk=devicename`
LVM	Extend a logical volume or increase space allocated to a logical volume.	`lvextend -l 50 /dev/vol_grp/lvol_name` l– indicates the number of logical extents in the logical volume

VxVM	Increase the volume by or to a given length.	`vxresize -g disk_group -F vxfs vol_name length`
		`vxassist growto vol_name new_length`
		`vxassist growby vol_name length_change`
		Grow the file system after growing the volumes.

LVM	Reduce a logical volume.	`lvreduce -L to_size /dev/vol_grp/lvol_name` `-L` indicates the number of megabytes.

VxVM	Reduce a volume by or to a given length.	`vxresize -g disk_group -F vxfs vol_name to_length`
		`vxassist -b shrinkby vol_name length`
		`vxassist -b shrinkto vol_name newlength`
		Shrink the file system before reducing the volume.

Table 3-2 LVM and VxVM task comparisons *(continued)*

Task type	Description	Example
LVM	Import and activate a volume group.	`vgimport -v /dev/vol_grp\ /dev/dsk/disk_name` `vgchange -a y /dev/vol_grp`
VxVM	Import a disk group to make the specified disk group accessible on the local machine.	`vxdg -tfC -n newname import disk_group` Option 8 in the `vxdiskadm` menu performs this task
LVM	Export and deactivate an LVM volume group, and its associated logical volumes.	`vgchange -a n vol_group` `vgexport /dev/vol_group`
VxVM	Deport a disk group to disable access to the specified disk group. A disk group cannot be deported if any volumes in the disk group are currently open.	`vxdg deport disk_group` Option 9 in the `vxdiskadm` menu performs this task
LVM	Back up volume group configuration information.	`vgcfgbackup -f /pathname/filename vol_grp`
VxVM	Back up volume group configuration information.	`dgcfgbackup -f /pathname/filename vol_grp`
LVM	Restore volume group configuration to a particular physical volume.	`vgrestore -n /dev/vol_grp /dev/rdsk/disk_name`
VxVM	Restore volume group configuration to a particular physical volume.	`dgcfgrestore -n /dev/vol_grp /dev/rdsk/disk_name`
LVM	Increase or decrease secondary swap space. Enlarge an existing swap logical volume, or add a new swap logical volume.	`lvextend`–to increase swap spac `lvreduce`–to decrease swap space
VxVM	Not supported for the current release. Add a new swap volume (HP-UX 11i Version 1.5 only).	`vxassist make swapvol2 size` (HP-UX 11i Version 1.5 only)
LVM	Remove a volume group. This destroys a volume group by removing its last disk and removing it from `/etc/lvmtab`.	`vgremove /dev/vol_grp` This is preceded by `lvremove` and `vgreduce` down to the last disk.
VxVM	Destroy a disk group.	`vxdg deport disk_group` `vxdg init disk_group`

Table 3-2 LVM and VxVM task comparisons *(continued)*

Task type	Description	Example
LVM	Extend a volume group by adding LVM disks to the volume group.	`vgextend /dev/vol_grp/\ /dev/dsk/disk_name`
VxVM	Add one or more disks to the disk group.	`vxdiskadd disk_name` Option 1 in the `vxdiskadm` main menu performs this task.
LVM	Reduce a volume group by reducing the number of disks in a volume group	`vgreduce /dev/vol_grp /dev/dsk/disk_name`
VxVM	Remove a disk from disk group.	`vxdg -g disk_group -k rmdisk disk_name`
LVM	Mirroring a disk involves several steps.	1 `pvcreate /dev/rdsk/second_disk` 2 `vgextend /dev/vol_grp \ /dev/dsk/second_disk` 3 `lvextend -m no_of_mirrors \ /dev/vol_grp/lvol_name \ /dev/dsk/second_disk`

VxVM	Mirroring a disk
	To mirror volumes on a disk or control default mirroring and causes a disk to have its contents mirrored to available space on another disk.
		`vxmirror -g disk_group -d yes|no \ disk_name [new_disk_name]`
		`vxmirror -d yes disk_name`
		Option 6 in the `vxdiskadm` menu performs this task.

LVM	Mirroring an LVM root disk involves several steps.	1 `pvcreate -B /dev/rdsk/second_disk` 2 `mkboot -l /dev/dsk/second_disk` 3 `vgextend /dev/vol_grp \ /dev/dsk/second_disk` 4 `lvextend -m no_of_mirrors \ /dev/vol_grp/root_lvol \ /dev/dsk/second_disk` 5 `lvlnboot -r /dev/vol_grp/lvol_name`

Table 3-2 LVM and VxVM task comparisons *(continued)*

Task type	Description	Example
VxVM	Mirroring the VxVM root disk.	`vxrootmir [-v] [-t tasktag] disk_access_name | disk_media_name`
LVM	Create a logical volume in LVM volume group.	`lvcreate -L vol_size /dev/vol_grp`

VxVM	Create a volume of one of these layout types:
	A concatenated volume
	A striped mirror volume
	A RAID-5 volume
		`vxassist make vol_name length`
		`vxassist make vol_name length layout=mirror, stripe`
		`vxassist make vol_name length layout=raid5`

LVM	Display information about logical volumes.	`lvdisplay /dev/vol_grp/lvol_name`
VxVM	Display all volume information. Display information about a specific volume.	`vxprint -vt` `vxprint -ht vol_name`
LVM	Display information about volume groups.	`vgdisplay -v /dev/vol_grp`

VxVM	Display disk group information.
	Display information about a specific disk group.
		`vxdisk list`
		`vxprint -g disk_group`
		`vxdg list`
		`vxdisk list disk_group`

LVM	Display information about physical volumes.	`pvdisplay /dev/dsk/disk_name`
VxVM	Display information about Volume Manager volumes.	`vxinfo` or `vxprint`
LVM	Remove a logical volume.	`lvremove /dev/vol_grp/lvol_name`
VxVM	Remove a volume.	`vxedit rm vol_name`
LVM	Remove disks from a volume group or reduce the number of disks in the volume group.	`vgreduce /dev/vol_grp /dev/dsk/disk_name`

Table 3-2 LVM and VxVM task comparisons *(continued)*

Task type	Description	Example

VxVM	Remove disks from a disk group.	`vxdisk rm disk_group`
		
		`vxdg rmdisk disk_name`
		
		`vxdg -g group_name rmdisk disk_name`

LVM	Remove an entire volume group. Before attempting to remove the volume group, you must remove the logical volumes using `lvremove`, and all physical volumes except the last one using `vgreduce`.	`vgremove /dev/vol_grp`
VxVM	Deport a disk group. You must unmount and stop any volumes in the disk group first.	`vxdg deport disk_group`

LVM	Set up alternate links to a physical volume.
	
	If a disk has two controllers, you can make one primary and the other an alternate link.
		`vgcreate /dev/vol_grp\ /dev/dsk/disk_name /dev/dsk/disk_name_2`
		
		To remove the link:
		
		`vgreduce /dev/vol_grp /dev/dsk/disk_name`

VxVM	The Multipathing disk devices in the Volume Manager represent virtual devices with one or more physical access paths to a particular physical disk. Dynamic Multipathing provides reliability of disk access by dynamically switching to another physical path in the event of failure of a path.	The DMP feature in VxVM sets up links automatically. It is not required to set up links separately. See manual page `vxdmp` (7). See the *Veritas Volume Manager Administrator's Guide* for more information on DMP.

LVM	Create a mirrored logical volume.	`lvcreate -l num_log_extents -m 1 \`
		
		`-n mirr_lv /dev/vol_grp`

VxVM	Create a mirrored volume/plex or add a mirror to an existing volume.	`vxplex att vol_name plex_name`

LVM	Reduce a single/double mirrored logical volume to an unmirrored logical volume.
	
	Remove a mirrored logical volume.
		`lvreduce -m 0 /dev/vol_grp/mirr_lv`
		
		`lvremove /dev/vol_grp/mirr_lv`

Table 3-2 LVM and VxVM task comparisons *(continued)*

Task type	Description	Example
VxVM	Remove mirrors or reduce the number of plexes/mirrors. Remove a volume with the plexes associated with it.	`vxplex -o rm dis plex_name` `vxedit -rf rm vol_name`
LVM	Increase the number of mirror copies.	`lvextend -m 2 /dev/vol_grp/lvol_name`
VxVM	Add mirrors to a volume or increase the number of plexes.	`vxassist mirror vol_name`
LVM	Convert a mirrored logical volume into two logical volumes. Split a logical volume.	`lvsplit -s backup /dev/vol_grp/lvol_name`
VxVM	Snapshot a volume and create a new volume.	`vxassist snapshot vol_name new_vol_name`
LVM	Combine two logical volumes back into a mirrored logical volume	`lvmerge /dev/vol_grp/split_vol_name\ /dev/vol_grp/lvol_name` `split_vol_name= active logical volume`
VxVM	Returns the snapshot plex to the original volume from which it was snapped.	`vxassist snapback new_vol_name`
LVM	Move a mirrored logical volume from one disk to another.	`pvmove -n /dev/vol_grp/lvol_name\ / dev/dsk/disk_name /dev/dsk/disk_name2`
VxVM	Move a plex.	`vxplex mv orig_plex new_plex`
LVM	Synchronize a mirrored logical volume. Synchronize extents within a mirrored logical volume.	`lvsync /dev/vol_grp/lvol_name`
VxVM	Resynchronize operations for the given volumes.	`vxvol resync`
LVM	Synchronize extents within mirrored logical volumes in a volume group.	`vgsync /dev/vol_grp`

Table 3-2 LVM and VxVM task comparisons *(continued)*

Task type	Description	Example
VxVM	Resynchronize operations for the named volumes, or for volumes residing on the named disks. If no medianame or volume operands are specified, then the operation applies to all volumes.	*vxrecover -s vol_name*
LVM	Start a volume.	lvchange -a y /dev/vol_grp/lvol_name
VxVM	Start a volume.	vxrecover -s *vol_name* vxvol start *vol_name*
LVM	Stop a volume.	lvchange -a n /dev/vol_grp/lvol_name
VxVM	Stop a volume.	*vxvol stop vol_name*
LVM	Make a disk available as a hot spare.	*pvchange -z y /dev/dsk/disk_name*
VxVM	Make a disk available as a hot spare.	*vxedit set spare=on disk_name*

Example for a *disk_group = veritasdg*, *medianame = disk01*, *vol_name -= veritasvol*, *plex name = veritasvol-01*, *subdisk -= disk01-01*, *devicename = c0t0d0*

Tasks with no direct LVM equivalents

The following table lists tasks which have no direct LVM equivalent. Most of these tasks can be performed either with the Veritas Enterprise Administrator (VEA) GUI, or the command line interface. Additional information is available.

See the *Veritas Enterprise Administrator User's Guide*.

See the *Veritas Volume Manager Administrator's Guide*.

Table 3-3 Additional VxVM tasks with no LVM equivalents

Task descriptions	Examples
Hot-relocation: in addition to using disks as hot spares, the hot relocation facility can also use any available free space in the disk group. If no disks have been designated as spares when a failure of a redundant object occurs, VxVM automatically uses any available free space in the disk group in which the failure occurs. If there is not enough spare disk space, a combination of spare space and free space is used. After a disk is repaired, you can move all the hot-relocated subdisks back to the original disk using the `vxunreloc` (1M) utility.	No action needed for hot relocation. To move hot-relocated subdisks back to the original disk: *`vxunreloc disk_name`*
Rename a disk	*`vxedit rename old_diskname new_diskname`*
Offline a disk.	*`vxdisk offline disk_name`* Alternatively, menu option 12 of `vxdiskadm` performs this task.
Online a disk.	*`vxdisk online disk_name`* Select menu option 10 of `vxdiskadm`.
Evacuate a disk.	*`vxevac -g disk_group medianame new_medianame`*
Replace a disk.	Select menu option 4 of `vxdiskadm`.
Recover volumes on a disk.	*`vxrecover -g disk_group vol_name medianame`*
Display a DMP node.	*`vxdisk list meta_device`*
Rename a disk group.	*`vxdg -tC -n newdg_name`*
Rename a volume.	*`vxedit -v rename name newname`* Update the `/usr/fstab` file with the new name.
Add a DRL log to a volume.	*`vxassist addlog vol_name`*
Create a snapshot copy of a volume.	*`vxassist snapshot vol_name temp_vol_name`*
Recover a volume.	`vxrecover -g` *`disk_group volume medianame`* `vxmend fix clean` *`plex_name`*
Repair a mirror	*`vxplex att plex_name`*
Disable a mirror	*`vxplex det plex_name`*

Table 3-3 Additional VxVM tasks with no LVM equivalents *(continued)*

Task descriptions	Examples
Remove a log from a volume.	`vxassist remove log vol_name`
Move a subdisk.	`vxsd mv old_subdisk new_subdisk`
Join subdisks.	`vxsd join subdisk1 subdisk2 new_subdisk`

Existing features in LVM not supported in VxVM

Some of the existing features in LVM are not supported in the current release of VxVM. Given below is a table with the unsupported LVM features, and possible workarounds in VxVM.

Table 3-4 LVM features and VxVM equivalents

LVM Feature	VxVM Equivalent
Physical volume groups	VxVM has no equivalent feature. The disk group feature of VxVM combines the logical volume group (VG) and physical volume group (PVG) of LVM.
Powerfail timeout feature: Automatically re-enable a disk or a path to a disk, after temporary error condition (resulting inEPOWERF error on I/Os) disappears on that disk or path.	Powerfail timeout feature: After the EPOWERF error condition disappears, the reconfiguration command must be run manually to re-enable the paths and the disks which were disabled due to EPOWERF error. Additional information is available. See the `pfto` feature in the `vxdctl`(1M) manual page.
Logical Volume Timeout (LVTO). If LVTO on a logical volume is set to zero, which is the default, an I/O is retried forever.	VxVM does not support the LVTO feature. However, VxVM supports the powerfail timeout feature to handle transient error conditions. By default, the use of PFTO is disabled in the HP-UX native multipathing devices. In case of DMP devices the use of PFTO is enabled. However, you can change the PFTO settings. See the powerfail timeout feature and the `pfto` feature in the `vxpfto` manual pages.
Bad media block relocation.	VxVM relocates whole subdisks. Smaller granularity relocation is not supported. The bad block reallocation feature does not exist in VxVM because the vectoring of bad blocks is now done by most hardware.

Chapter 4

SMH and the VEA

This chapter includes the following topics:

About SMH and the VEA

This chapter describes the Veritas Enterprise Administrator (VEA) graphical user interface (GUI), and its relationship with the System Management Homepage (SMH).

SMH and the VEA coexist as independent entities. The VEA recognizes and labels LVM volumes and disks, but does not manage them. Similarly, SMH recognizes and labels VxVM disks, but does not manage them. To manage VxVM disks graphically, you must use the VEA.

Note: If you run the VEA client from the command line, refresh SMH to see VxVM changes reflected in the SMH screen. Only privileged users can run the VEA client. It is possible to launch the VEA client from SMH by selecting **Home** > **System Configuration** > **Disks and File Systems** > **Veritas Enterprise Administrator**.

For information about the VEA, see the *Veritas Enterprise Administrator User's Guide* and the online help that is available from within the VEA.

Displaying disk devices in SMH

To display disk devices in SMH, select **Tools** > **Disks and File Systems** > **Disks**. The **Disks** tab of the **HP-UX Disks and File Systems Tool** screen lists the system's disk devices. To switch between legacy device names and new agile device names, click on **Toggle Global Device View**.

In the legacy view, all paths to a device are listed using the old-style `c#t#d#` naming convention that is used in the `/dev/dsk` and `/dev/rdsk` directories.

In the agile view, a device is listed according to the name of its persistent device file in the `/dev/disk` and `/dev/rdisk` directories, using the `disk##` naming convention.

The number of paths to a device is indicated in the "Paths" column in both views.

When VxVM is installed on the system, the "Use" column in SMH indicates whether a disk is under LVM or VxVM control, or whether it is unused.

If a VxVM disk is online and part of a disk group, the disk group name is listed under the "VG Name" column. If a VxVM disk is initialized, but not yet part of a disk group, it is listed as "unused" under the "Use" column.

A more detailed description of a disk's properties can be obtained by selecting the radio button to the left of a listed device.

Figure 4-1 shows a sample Disk Devices screen for a system with several VxVM disks, one LVM disk, several unused disks, and a DVD-ROM device. The VxVM disks are all in the `mydg` disk group. The detailed view of disk30 shows that it has two paths with the legacy names c5t0d7 and c7t0d7.

Figure 4-1 Displaying disk devices in SMH

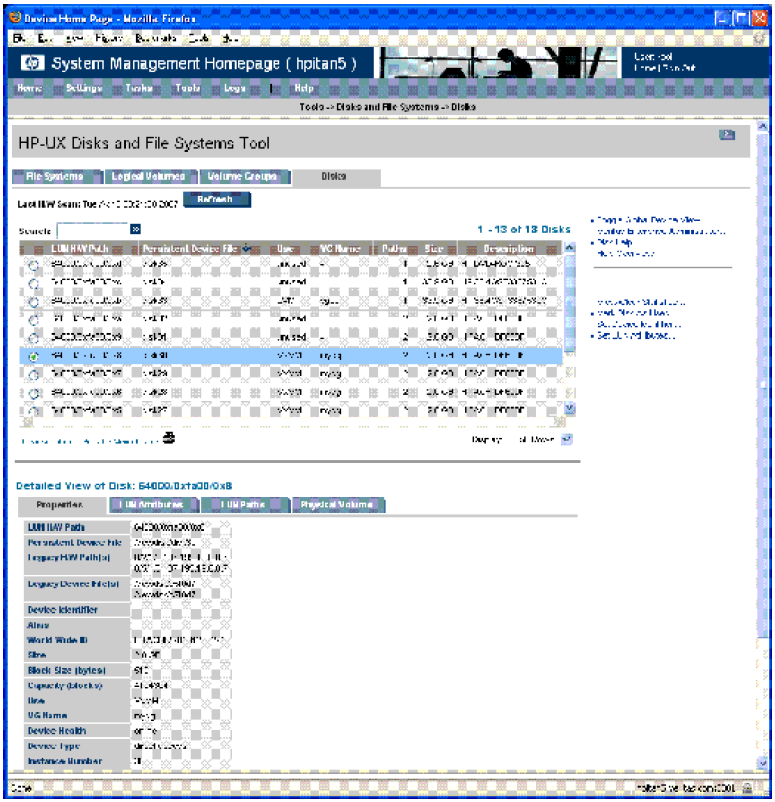

Displaying volume groups and disk groups in SMH

To display volume groups and disk groups in SMH, select **Tools** > **Disks and File Systems** > **Volume Groups**. The Volume Groups screen lists all the LVM volume groups and VxVM disk groups that are on the system.

A more detailed description of a volume group's properties can be obtained by selecting the radio button to the left of a listed volume group or disk group. However, you should use the VEA to obtain full details of VxVM disk groups.

Figure 4-2 shows an example Volume Groups screen. `vg00` is an LVM volume group with eight logical volumes; `mydg` is a VxVM disk group with three volumes.

Figure 4-2 Displaying LVM volume groups and VxVM disk groups in SMH

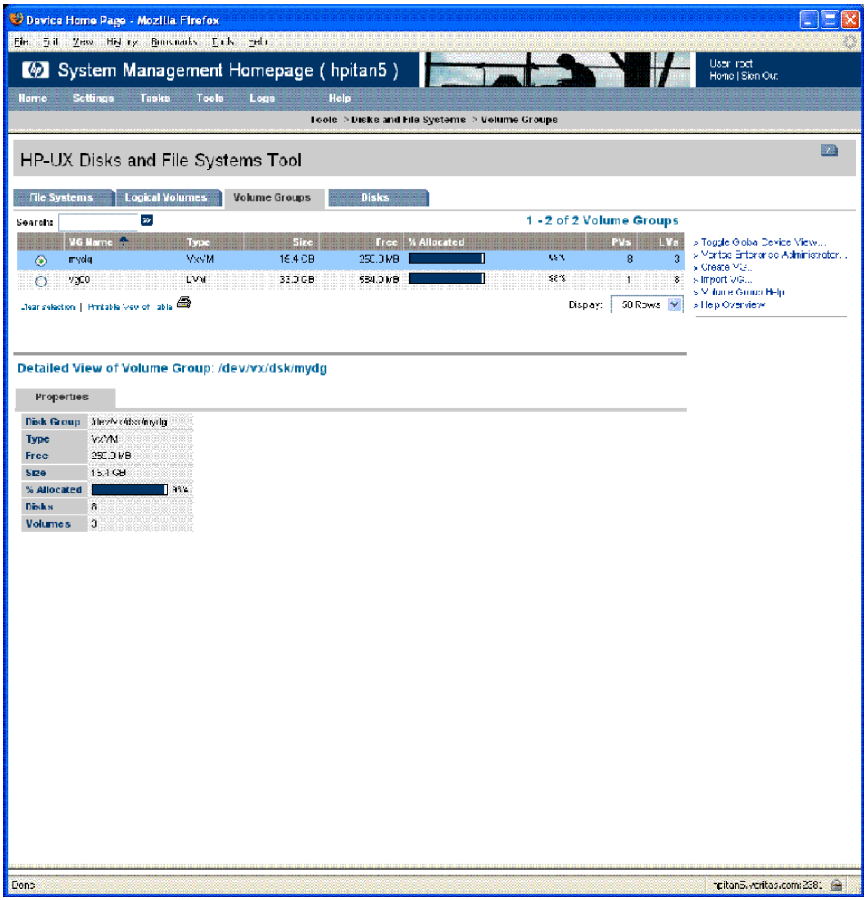

Displaying logical volumes in SMH

To display logical volumes in SMH, select **Tools** > **Disks and File Systems** > **Logical Volumes**. The Logical Volumes screen lists the LVM logical volumes and VxVM volumes on the system. The “Type” column indicates whether a volume is controlled by LVM or VxVM. The “Use” column shows whether a volume is in use and if so, what it is used for.

A more detailed description of a logical volume’s properties can be obtained by selecting the radio button to the left of a listed volume. However, you should use the VEA to obtain full details of VxVM volumes.

Figure 4-3 shows an example Logical Volumes screen. The LVM logical volumes in the `vg00` volume group are being used for HFS and VxFS file systems and for swap and dump. The `myvol1` and `myvol2` VxVM volumes in the `mydg` disk group are being used for VxFS and HFS file systems. The remaining VxVM volume, `myvol3`, is not currently in use.

Figure 4-3 Displaying LVM logical volumes and VxVM volumes in SMH

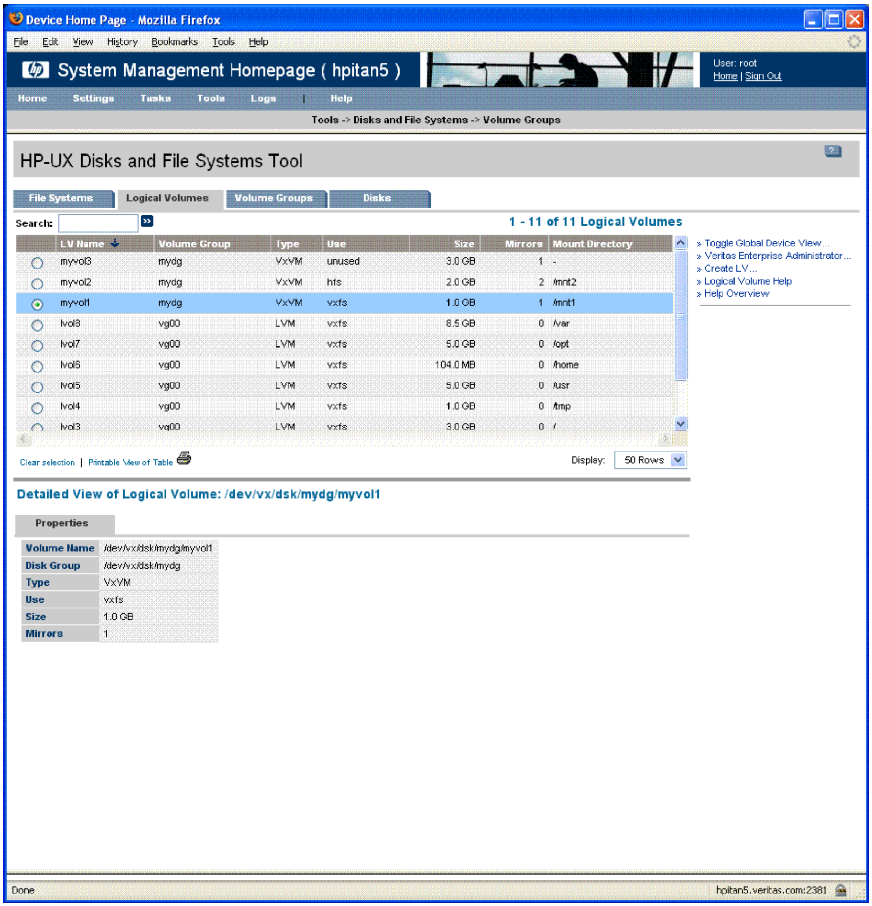

Appendix A

Conversion error messages

This appendix includes the following topics:

About conversion error messages

This appendix lists the error messages that you may encounter when converting LVM volume groups to VxVM disk groups and volumes. For each error message, a description is provided of the problem, and the action that you can take to troubleshoot it.

Table A-1 shows the error messages that you may encounter during conversion.

Table A-1 Conversion error messages

Message	Description
`Analysis indicates that this volume group cannot be converted because not all of the disks and/or volumes in the LVM volume group are currently accessible`	For successful conversion, all physical volumes in a volume group must be on-line, and all logical volumes must be active and accessible. Make sure the physical volumes in a volume group are on-line and the logical volumes are active and not in use.
`Analysis shows that there is insufficient private space available to convert this volume group`	The error message indicates the maximum amount of records that can be stored in the private space, and how many records are needed to convert this particular volume group. You can reduce the number of records needed by reducing the number of logical volumes in volume group by combining some of the logical volumes together.

Table A-1 Conversion error messages *(continued)*

Message	Description

`Device device_name has the following bad blocks... Cannot convert LVM Volume Group`	Unlike LVM, VxVM does not support bad block revectoring at the physical volume level. If there appear to be any valid bad blocks in the bad block directory (BBDIR) of any disk used in an LVM volume group, the group cannot be converted.
	To clear the BBDIRcomplete the following steps in the order shown:
	■ Stop all access to the physical volume, and back up all the data in thevolume group.
	■ If the errors are persistent or the disk has not been replaced, have a certified engineer revector the blocks or replace the disk.
	■ Back up the LVM headers, for example:
	`# vgcfgbackup /dev/vg01`
	■ Close all the logical volumes associated with the volume group, and then deactivate the volume group.
	■ Create a dummy link to the raw device, for example:
	`# ln /dev/rdsk/c0t0d0 \`
	`/dev/rdsk/temp`
	■ Destroy the LVM headers on the dummy link:
	`# pvcreate -f /dev/rdsk/temp`
	■ Remove the dummy link to the raw device:
	`# rm /dev/rdsk/temp`
	■ Restore the headers from the `lvmconf` backup, for example:
	`# vgcfgrestore -n /dev/vg01 \`
	`/dev/rdsk/c0t0d0`
	■ Activate the volume group. Check the integrity of the data, or recover the data from a backup.

`The conversion process was unable to deactivate the volume group vol_grp_name`	This indicates that the conversion process cannot deactivate the volume group. The conversion cannot be completed without rebooting the machine. If you cannot afford to reboot, then choose abort and try again later.

Table A-1 Conversion error messages *(continued)*

Message	Description
`This Volume Group contains one or more logical volumes with mirrored data`	If you attempt to convert a Mirrored LVM Volume Group without a valid VxVM license installed, the conversion is not allowed. Install the required license before attempting the conversion.
`Too many LVM Volumes to convert in this LVM Volume Group`	If there is insufficient private space, the conversion is not allowed to continue. Also, the conversion records already generated are removed such that in the event of an unexpected crash and reboot, the conversion cannot proceed automatically. You can reduce the number of logical volumes in volume group by combining some of the logical volumes together, or by aborting. You can restart the conversion process later with fewer volumes in the group.
`vgchange: Couldn't deactivate volume group /dev/vol_grp`	The conversion process was unable to deactivate the volume group. The conversion cannot proceed without reboots being done. If you choose to not reboot your system, the conversion is aborted. The system responds with an option to complete the conversion by rebooting the system.
`vxdiskadm or vxvmconvert is already being run and these programs cannot run concurrently`	The system detects that the `vxdiskadd` or `vxvmconvert` program is already running. Retry at a later time. Otherwise, if you are certain that no other users are running either of these programs, remove the file `.DISKADD.LOCK` from the `/var/spool/locks` directory to allow you to run `vxvmconvert`.

`LVM Volume Group Volume Group name reports a version 2.0`
 	
`LVM Volume Group other than version 1.0 is not supported for conversion.`
	This message occurs while converting an LVM version 2 volume group to VxVM. You can only convert LVM version 1 volume groups.

Table A-1 Conversion error messages *(continued)*

Message	Description
`WARNING!: The LVM Volume Group name is too long for longvGname. VxVM will not allow such long name for VxVM Disk Group. Default VxVM Volume group will have only first 31 characters of LVM Volume Group name.`	If you convert an LVM volume group with more than 31 characters in the volume group name, VxVM will retain only the first 31 characters of the LVM volume group name.
`VxVM ERROR V-5-2-0 The LVM Volume Group (longvGname) has Logical Volume (1234567890123456789012345678901) having > 28 characters. VxVM doesn't allow Volume names that long. Please reduce the name of the Logical Volume and retry the conversion.`	The conversion of a logical volume containing more than 28 characters in the logical volume name is not supported.

Glossary

block	A unit of space for data on a disk, typically having a size of 1024-bytes.
Dirty Region Logging	Dirty Region Logging (DRL) is an optional property of a volume, used to provide a speedy recovery of mirrored volumes after a system failure. DRL keeps track of the regions that have changed due to I/O writes to a mirrored volume.
file system	The organization of files on storage devices. The term file system can refer either to the entire file system or to a subsection of that file system, contained within a disk section or a logical volume that can be mounted or unmounted from that tree.
logical volume	A logical structure that is a map of storage areas on physical volumes. A logical volume can be conceptualized as a storage device of flexible size. The data in a logical volume can be mapped to one or more physical volumes. A virtual disk device that represents a contiguous virtual disk space that maps to single or multiple areas on a single or multiple physical volumes.
LVM	The Logical Volume Manager (LVM) is a subsystem for managing disk space. LVM is an HP-UX product, similar to VxVM.
LIF	The Logical Interchange Format (LIF) is a HP mass-storage format that can be used for interchange of files among various HP computer systems. Each boot disk has a LIF directory that contains boot programs.
logical Extent	A set of logical blocks that maps to one physical extent and is a basic unit of access in a logical volume.
LVMREC	The LVM record, which is an identifier that is set on each disk when it is initialized the first time by LVM. There are two copies of the LVMREC; one is at sector 8, while the other is at sector 72.
mirror disk/UX	HP-UX software product that allows disk mirroring as part of LVM functionality. Mirror Disk/UX allows up 2 mirror copies in a volume.
Mirror Write Cache (MWC)	A MirrorDisk/UX mechanism whose use is optional, that tracks outstanding mirror write requests and provides a basis for the resynchronization of data blocks after a system crash or power failure.
physical volume	A disk that has been initialized by LVM becomes known as a physical volume.
public region	A region of a physical disk managed by VxVM that contains available space and is used for allocating subdisks.

private region
A region of a physical disk used to store private, structured VxVM information. The private region contains a disk header, a table of contents, and a configuration database. The table of contents maps the contents of the disk.

physical extent
A set of physical disk blocks on a physical volume that forms a basic unit of access in LVM. This also forms the allocation unit for logical volumes.

snapshot
A temporary extra copy (plex/mirror) created in a volume. A separate volume is created once its contents are brought in sync with the original volume.

set of continuous physical extents [PEs]
Set of physical sectors (blocks) contained within a single physical volume. A physical extent is a specific, contiguous region of the disk where data resides. This is of a constant size usually of 4 MB and has no partitions.

VGRA
The Volume Group Reserved Area (VGRA) is a region on an LVM disk that holds LVM configuration information and is at a fixed location. The location of this fixed location depends upon whether the disk is a boot disk or a non-boot disk. For boot disks, the VGRA starts at sector 2144. For non-boot disks, the VGRA starts at sector 128.

volume group
The collective identity of a set of physical volumes, which provide disk storage for the logical volumes. A set of physical volumes whose space can be combined and logically divided up into logical volumes. Only logical volumes and physical volumes that are a part of a volume group can map together, a physical volume can belong to only one volume group.

Index