

VERITAS Volume Manager 4.1 Release Notes

HP-UX 11i v2



i n v e n t

Manufacturing Part Number: 5991-1147

May 2005

Edition 4

© Copyright 2004 - 2005 Hewlett-Packard Development Company L.P.

Legal Notices

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

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

Trademark Notices

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

VERITAS™ is a registered trademark of VERITAS Software Corporation.

Publication History

The manual publication date and part number indicate its current edition. The publication date will change when a new edition is released. The manual part number will change when extensive changes are made.

To ensure that you receive the new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

- First Edition: December 2002, 5187-1878, HP-UX 11i Version 1(B.11.11)
CD-ROM (Software Pack 11i December 2002) and Web (<http://docs.hp.com>)
- Second Edition: August 2003, 5971-4774, HP-UX 11i Version 2 (B.11.23)
Web (<http://docs.hp.com>)
- Third Edition: December 2003, B3929-90015, HP-UX 11i Version 2 (B.11.23)
Web (<http://docs.hp.com>)
- Fourth Edition: June 2005, 5991-1147, HP-UX 11i v2

Getting Help

For license information (U.S. and Canadian Customers) contact:

- Phone: 650-960-5111
- Email: hplicense@mayfield.hp.com

For license information (Europe) contact:

- Phone: +33.(0)4.76.14.15.29
- Email: codeword_europe@hp-france-gen1.om.hp.com

For the latest information on patches visit:

- <http://itrc.hp.com>

For Technical Support visit:

- <http://welcome.hp.com/country/us/en/support.html>

1. VERITAS Volume Manager 4.1 Release Notes

Document Organization	1
Conventions	1
Overview of Base-VXVM 4.1 for HP-UX	2
Unsupported Features	3
Features of Base-VXVM 4.1 compared to LVM	3
New Features in VERITAS Volume Manager 4.1	3
Changed Features in VERITAS Volume Manager 4.1	6
Working With Base-VXVM 4.1 Root Disks	7
Adding Dump Volumes Using Volume Manager Disks	7
Removing a Mirrored Volume Manager Root Disk	7
VEA Graphical User Interface and SAM	8
Coexistence with HP Logical Volume Manager (LVM)	8
Documentation	9
Compatibility Information and Installation Notes	9
Supported Platforms	9
Software Requirements	10
Known Problems and Workarounds	10
Installation Issues	10
Utility Issues	11
DMP Issues	16

Contents

1 VERITAS Volume Manager 4.1 Release Notes

This document provides information on VERITAS Volume Manager 4.1 on HP-UX 11i v2. Before you install the packages, please read this entire document.

Document Organization

- “Conventions” on page 1
- “Overview of Base-VXVM 4.1 for HP-UX” on page 2
- “Compatibility Information and Installation Notes” on page 9
- “Known Problems and Workarounds” on page 10

Conventions

The following table describes the typographic conventions used in this guide.

Table 1-1

Typeface	Usage	Examples
monospace	Computer output, file contents, files, directories, software elements such as command options, function names, and parameters	Read tunables from the <code>/etc/vx/tunefstab</code> file. See the <code>ls (1)</code> manual page for more information.

Table 1-1 (Continued)

Typeface	Usage	Examples
<i>italic</i>	New terms, book titles, emphasis, variables to be replaced by a name or value	See the <i>VERITAS Volume Manager 4.1 Release Notes</i> for details.
<i>monospace</i> <i>(bold and italic)</i>	Variables to be replaced by a name or value in user input	# mount -F fstype mount_point
Symbol	Usage	Examples
%	C shell prompt	Not applicable
\$	Bourne/Korn/Bash shell prompt	Not applicable
#	Superuser prompt (all shells)	Not applicable
\	Continued input on the following line	# mount -F vxfs \ /h/filesys
[]	In a command synopsis, brackets indicates an optional argument	ls [-a]
	In a command synopsis, a vertical bar separates mutually exclusive arguments	mount [suid nosuid]

Overview of Base-VXVM 4.1 for HP-UX

This release of the VERITAS Volume Manager includes only Base-VXVM 4.1.

Unsupported Features

The following features are not supported in VERITAS Volume Manager 4.1 for HP-UX:

- The FlashSnap (FastResync and Dynamic Disk Group Split and Join and Storage Checkpoints).
- VERITAS Volume Replicator (VVR)
- VERITAS Cluster Server (VCS)
- The VERITAS Cluster File System(CFS)
- VERITAS Cluster Volume Manager (CVM)
- Quicklog

Features of Base-VXVM 4.1 compared to LVM

The VERITAS Volume Manager 4.1 for HP-UX includes many features that are not available with Logical Volume Manager (LVM) on HP-UX 11i version 2. Some of these features are:

- VERITAS Enterprise Administrator (VEA)—a Java-based administrative GUI
- RAID-5
- Dynamic multipathing for I/O load balancing
- Support for up to 32 mirrors
- Striped mirrors
- Online relayout of volumes
- Device Discovery Layer (DDL)
- Boot performance enhancements

New Features in VERITAS Volume Manager 4.1

New features in VxVM 4.1 include the following:

- Enhancements to Device Discovery Layer (DDL)

The following enhancements have been done to the DDL feature:

- Dynamic discovery of disks or their attributes
- Allowing DMP kernel extensions to support multipathing and/or load balancing in a configuration that is specific to a particular disk array.
- Previous releases of VxVM supported two naming schemes, OS-Native Scheme and Enclosure-based Naming. Although VxVM provided the means of changing between the two schemes dynamically, it was necessary to restart `vxconfigd` which is a time-consuming operation. With this release, the `vxconfigd` restart operation is no longer necessary.
- Base-VXVM now supports persistent device names. The disk names, once assigned, will remain constant across reconfiguration and rebooting.
- Device discovery is now multithreaded. Performance improvements have also been made to the DMP.

- Multi-Device Support

Multi-Device Support allows several volumes to be represented by a single logical object. All I/O to and from the underlying logical volume is directed via the I/O interfaces of the volume sets. This feature is designed primarily for the multi-device file system enhancement for VERITAS File System.

- Configuration Backup and Restore Support

Configuration Backup and Restore support allows you to backup and restore or replicate all Base-VXVM configuration data for disk groups and volumes. This applies only to Base-VXVM configuration data and does not apply to user and application data.

- Hitachi Extended Copy Facility

Using the Hitachi e-copy facility a host server can request data movement from a disk to a tape or to another disk without involving the server CPU time I/O during data movement.

- Storage Expert

System Administrators sometimes find that gathering and interpreting data about large and complex configurations can be a difficult task. Consisting of a set of simple commands that collect and analyse Base-VXVM configuration data, Storage Expert produces a summary report that recommends configuration improvements.

- Detecting and correcting the Serial Brain Split Condition in a Disk Group

Under certain circumstances a serial split brain condition can occur in a disk group where the actual serial numbers on disks are different from those listed in the configuration database copies on the disks. To diagnose and troubleshoot this problem, you can use the `vxsplitlines` command. This command also details the commands that you can use to correct the condition.

- Online Node Addition

Base-VXVM now allows you to add a node to a cluster dynamically.

- Dynamic LUN Expansion

Previous releases of Base-VXVM did not allow you to take advantage of a change in the size of a disk once the disk had been initialized; any attempt to do so would result in loss of data. However, in this release, Dynamic LUN capability allows you to resize the device while preserving the existing data.

- Unique Number for Messages

To enable easy identification of messages, each message contains a unique identifier. Using this ID number when reporting messages to HP support will enable them to respond more easily to customer requests.

- Support for `B_FAILFAST` flag

Dynamic Multipathing supports the `B_FAILFAST` flag if the underlying device driver exports the `ddi-failfast-supported` property. The `B_FAILFAST` flag allows the disk drivers to avoid heroic, time-consuming retry behavior when it is not required or is inappropriate. The `B_FAILFAST` flag is intended to allow modified error recovery behavior in the disk driver when no communication can be established with the device. However it does not dictate specific retry behavior, nor does it imply or guarantee deterministic failure time

When I/O is requested to a device associated with the `ddi-failfast` supported property, and the device has multiple paths available or I/O request is for a Read operation from a volume that has an alternate source (mirror), DMP will set the `B_FAILFAST` flag in the I/O request. While DMP sets the `B_FAILFAST` flag, DMP has no knowledge of whether the underlying disk driver's behavior is influenced by `B_FAILFAST`.

- Coexistence of DMP with Third Party Drivers

With this release, the generic framework within the Device Discovery Layer (DDL) enables DMP to coexist with other third party multi-pathing drivers such as EMC PowerPath.

Changed Features in VERITAS Volume Manager 4.1

Changed features in VxVM 4.1 include the following:

- Default Disk Group Behavior in Base-VXVM 4.1

In previous versions of VxVM, for VxVM to function, a disk group with the name `rootdg` (root disk group) was required. The `rootdg` disk group had to contain at least one disk. This requirement no longer exists, however you might find it convenient to create a system-wide default disk group.

For operations that require a disk group, the system wide default disk group will be used if the `-g` option is not specified with the `Base-VXVM` command. The main benefit of creating a default disk group is that the `Base-VXVM` commands default to the default disk group and you need not to use the `-g` option. To verify the default disk group after it has been created, enter the command:

```
# vxdg defaultdg
```

- Simplification of `vxinstall`

VxVM 4.1 no longer requires a `rootdg` to be configured. `vxinstall` is more simple now, and you can no longer use this command to configure disk groups on your system after installation. You must now use `vxdiskadm` or the VERITAS Enterprise Administrator.

- Temporary File Systems for Use by DMP

In this release, two temporary file systems are mounted for use by DMP. When they are mounted, output from the `df` command is similar to the following:

Filesystem	Kbytes	Used	avail	Capacity	Mounted on
...					
dmpfs	902488	0	902488	0%	/dev/vx/rdmp
dmpfs	902488	0	902488	0%	/dev/vx/dmp

- 32-Terabyte File support

The Base-VXFS 4.1 release supports a new disk layout (version 5) that enables the creation of file systems up to 32 terabytes in size. File systems created on Base-VXFS 4.1 will by default use the new disk layout, but files created on Version 5 can still only be a maximum of 2 terabytes. File systems larger than 1TB can be created only on 64-bit kernels and must be created on a VERITAS Volume Manager volume. An 8 terabyte file system requires a 2K block size, a 16 TB file system requires a 4K block size, and a 32 TB file system requires an 8K block size. An online conversion utility, `vxupgrade`, is provided to upgrade existing disk layouts to Version 5 on mounted file systems.

- A change in this release is that the default disk type is now `auto`. This means that disks that used to show up in the output from `vxdisk list` as `simple` will now show up as `auto:hpdisk`.

Working With Base-VXVM 4.1 Root Disks

This section describes the following:

- “Adding Dump Volumes Using Volume Manager Disks” on page 7
- “Removing a Mirrored Volume Manager Root Disk” on page 7

Adding Dump Volumes Using Volume Manager Disks

Complete the following procedure to use Volume Manager volumes for additional dump volumes in configurations with LVM or Base-VxVM root disks:

Step 1. Remove a disk from LVM control. Execute the following:

```
# pvcreate -f /dev/rdisk/c0t0d0
# pvremove /dev/rdisk/c0t0d0
```

Step 2. Put a disk under Volume Manager control and place in `rootdg`:

Use `vxdiskadm` option 1, or VEA.

Step 3. Add a volume to this disk:

Use `vxassist` or VEA.

Step 4. Use `crashconf` as you would under LVM.

For example;

```
# crashconf /dev/vx/dsk/rootdg/your_dump_volume
```

Removing a Mirrored Volume Manager Root Disk

Use the following commands to remove the mirrored Volume Manager Root Disk:

```
# daname=<c0t0d0>          /* of mirrored root disk */
# dmname=`vxprint -g rootdg -F %dmname -e sd_da_name~/ $daname/`|head -1`
```

```
# plxnames=`vxprint -g dg -F "%assoc" -e sd_dm_name=="$dmname"`  
# vxplex -o rm dis $plxnames  
# vxdg -g rootdg rmdisk $dmname
```

VEA Graphical User Interface and SAM

The VERITAS Enterprise Administrator (VEA) provides a Java-based graphical user interface for managing Base-VXVM. VEA has two parts:

- server
- client

The server must run on the system running Base-VXVM. The client can run on the server machine or the client software can be installed on a different HP-UX 11i v2 system to manage Base-VXVM remotely. Note that only HP-UX 11i clients are supported.

The HP-UX System Administration Manager (SAM), and VEA exist as independent entities. The VEA client should be launched from SAM. SAM is used to manage LVM objects and the VEA is used to manage VxVM objects. However, VEA recognizes and labels LVM volumes and disks, and similarly, SAM recognizes and labels VxVM volumes and disks. To manage VxVM disks graphically, you must use VEA. For information about VEA, see the *VERITAS Enterprise Administrator (VEA 500 Series) Getting Started guide*.

Coexistence with HP Logical Volume Manager (LVM)

The VERITAS Volume Manager for HP-UX coexists with HP Logical Volume Manager (LVM). Before VxVM 4.1, the VERITAS Volume Manager could not be used to control the root/boot disk.

Both LVM and VxVM utilities are interoperable and will not overwrite disks that are being managed by the other volume manager. As mentioned above, the administrative utilities (SAM and VEA) recognize and identify all disks on the system.

Although this release is targeted at new customer installations, a conversion utility, `vxvmconvert`, is provided for converting LVM volume groups to VxVM volume groups. Refer to the *VERITAS Volume Manager 4.1 Migration Guide* for details on using `vxvmconvert`.

Documentation

The following documents describe the VERITAS Volume Manager for HP-UX. They are available for download from the web at <http://www.software.hp.com>.

- *VERITAS Volume Manager 4.1 Administrator's Guide*
- *VERITAS Volume Manager 4.1 Troubleshooting Guide*
- *VERITAS Enterprise Administrator (VEA 500 Series) Getting Started*
- *VERITAS Volume Manager 4.1 Migration Guide*
- *VERITAS Volume Manager 4.1 Hardware Notes*
- *VERITAS Storage Foundation Cross-Platform Data Sharing Guide*
- *VERITAS 4.1 Upgrade and Installation Guide*

Manpages are also available on your system via the `man(1)` command.

Compatibility Information and Installation Notes

Supported Platforms

The following platforms are supported for VERITAS Volume Manager 4.1 for HP-UX:

IPF	cx2600, rx1600, rx2600, rx4640, rx5600, rx7600, rx8600 series, zx2000, Integrity Superdome and zx6000 workstations
PA-RISC	rp2400, rp3340, rp4400, rp5400, rp7400 and rp8400 series 64-bit A, L, and N classes

Software Requirements

This release of VERITAS Volume Manager operates only on HP-UX 11i 64-bit operating systems with the September 2004 HP-UX 11i Version 2 release or later (for Itanium-based and PA-RISC systems).

Known Problems and Workarounds

Base-VXVM 4.1 has the following known problems and workarounds:

- “Installation Issues” on page 10
- “Utility Issues” on page 11
- “DMP Issues” on page 16
- “Device Issues” on page 20
- “VEA Issues” on page 21

Installation Issues

Cautionary Note When Using HP-UX Maintenance Mode Boot (MMB)

HP-UX Maintenance Mode Boot (MMB) is meant to be used in the recovery from catastrophic failures that have prevented the target machine from booting. If a mirrored root is configured, then when booting in MMB mode, only one mirror is activated. Therefore, any writes to the root filesystem in this mode could cause root filesystem corruption later when both mirrors are configured.

The `vx_emerg_start` script is provided to be used when starting the Volume Manager in MMB mode. This script will avoid writing to the root file system, unless absolutely necessary. If it needs to update the `volboot` file, it tells you to invoke the `vx_emerg_start` script, using the `-f` option, to perform the write.

HP recommends that after the `vx_emerg_start` script has been run to start the Volume Manager while in MMB mode on a mirrored root, half of the mirror not booted from is removed. This can be carried out as follows:

1. Determine which disk you booted from.
2. Use the `vxdisk list` command to find your boot disk in the `DEVICE` column on the far left.

3. Find the Disk Media (DM) name of your boot device by looking up your boot device in the DISK column. This will be a name such as `rootdisk01`, `rootdisk02` and so on. Also note the name of the mirror disk DM.
4. Use the `vxprint -g rootdg rootvol` command. If you have a mirrored root volume, you will see two lines with `p1` on the far left side. Look at each plex (or `p1`) entry and immediately below it will be the subdisk associated with the plex. It will start with an `sd` in the far left column.
5. Look at the NAME field immediately to the right of the `sd` column. This will show the subdisk name, which is made up of the DM name followed by `-nn`, where `nn` is a number such as `03`, `04`, and so on. This should allow you to identify the DM name of the disk that is not your boot disk. You can remove the plex and its associated subdisk by executing the `vxplex` command as follows:

```
vxplex -o rm dis plex name
```

For example, to remove the `rootvol` plex associated with `rootdisk02`:

```
# vxprint -g rootdg rootvol
TY  NAME          ASSOC      KSTATE  LENGTH  PLOFFS   STATE
v   rootvol       root       ENABLED 524288  -        ACTIVE
p1  rootvol-01     rootvol    ENABLED 524288  -        ACTIVE
sd  rootdisk01-03 rootvol-01 ENABLED 524288  0        -
p1  rootvol-02     rootvol    ENABLED 524288  -        ACTIVE
sd  rootdisk02-03 rootvol-02 ENABLED 524288  0        -
# vxplex -o rm dis rootvol-02
```

NOTE The `TUTIL0` and `PUTIL0` fields have been removed in the above `vxprint` output for readability.

6. Once the system has been repaired and is up in normal mode, the root volume can be mirrored using the command:

```
# vxassist -g rootdg mirror rootvol dm:rootdisk02
```

Utility Issues

vxassist relayout Considerations

- Problem: The `vxassist` relayout operation requires all mirrors in the volume to have the same layout (ref. incident 90840).

- **Workaround:** If the volume contains mirrors with different layouts, then you need to relayout the mirror plexes to the same layout before performing the volume relayout operation.

vxassist Command Does not Add a Mirror and a Log

- **Problem:** The `vxassist` command does not add a mirror and a log when processing a command such as the following:

```
# vxassist mirror volume layout=log ...
```

The mirror is added, and the log is omitted.

- **Workaround:** If a log and a mirror are to be added, add the mirror and the log in two separate `vxassist` invocations, as follows:

```
# vxassist mirror volume ...
```

```
# vxassist addlog volume ...
```

Base-VXVM and Multi-Host Failover Configurations

Outside the context of clustering functionality, Base-VXVM disk groups can be imported (made available) from only one host at any given time. When a host imports a disk group as private, the volumes and configuration of that disk group becomes accessible to the host. If the administrator or system software wants to privately use the same disk group from another host, the host that already has the disk group imported (importing host) must deport (give up access to) the disk group. Once deported, the disk group can be imported by another host.

If two hosts are allowed to access a disk group concurrently without proper synchronization, such as that provided by the Oracle Parallel Server, the configuration of the disk group, and possibly the contents of volumes, can be corrupted. Similar corruption can also occur if a file system or database on a raw disk partition is accessed concurrently by two hosts. This is not a problem limited to Base-VXVM.

Import Lock

When a host in a non-clustered environment imports a disk group, an import lock is written on all disks in that disk group. The import lock is cleared when the host deports the disk group. The presence of the import lock prevents other hosts from importing the disk group until the importing host has deported the disk group.

Specifically, when a host imports a disk group, the import normally fails if any disks within the disk group appear to be locked by another host. This allows automatic re-importing of disk groups after a reboot (autoimporting) and prevents imports by another host, even while the first host is shut down. If the importing host is shut down without deporting the disk group, the disk group can only be imported by another host by clearing the host ID lock first (discussed later).

The import lock contains a host ID (in Base-VXVM, this is the host name) reference to identify the importing host and enforce the lock. Problems can therefore arise if two hosts have the same host ID.

NOTE Since Base-VXVM uses the host name as the host ID (by default), it is advisable to change the host name of one machine if another machine shares its host name. To change the host name, use the `vxdctl hostid new_hostname` command.

Failover

The import locking scheme works well in an environment where disk groups are not normally shifted from one system to another. However, consider a setup where two hosts, Node A and Node B, can access the drives of a disk group. The disk group is first imported by Node A, but the administrator wants to access the disk group from Node B if Node A crashes. This kind of scenario (failover) can be used to provide manual high availability to data, where the failure of one node does not prevent access to data. Failover can be combined with a *high availability* monitor to provide automatic high availability to data, when Node B detects that Node A has crashed or shut down, Node B imports (fails over) the disk group to provide access to the volumes.

Base-VXVM supports failover, but it relies on the administrator or on an external high-availability monitor to ensure that the first system is shut down or unavailable before the disk group is imported to another system. For details on how to clear locks and force an import, see the `vxdg(1M)` manual page and the section on moving disk groups between systems in the VERITAS Volume Manager 4.1 Administrator's Guide.

Corruption of Disk Group Configuration

If `vxdgimport` is used with `-C` (clears locks) and/or `-f` (forces import) to import a disk group that is still in use from another host, disk group configuration corruption is likely to occur. Volume content corruption is also likely if a file system or database is started on the imported volumes before the other host crashes or shuts down.

If this kind of corruption occurs, you can use the `/usr/sbin/dgcfgrestore` command to restore the disk group configuration from a backup. The disk group configuration daemon, `dgcfgdaemon`, automatically makes a backup of a disk group configuration whenever it is changed. Though each disk group has numerous configuration copies, redundancy does not help in this case because corruption often affects all the configuration copies.

Disk group configuration corruption usually appears as missing or duplicate records in the configuration databases. This can result in a variety of `vxconfig d` error messages, including errors such as:

```
Association not resolved
```

```
Association count is incorrect
Duplicate record in configuration
Configuration records are inconsistent
```

Though some errors are reported only for specific group configuration copies, they are usually applicable to all copies. The following message appears along with the error:

Disk group has no valid configuration copies

See the VERITAS Volume Manager 4.1 Troubleshooting Guide for more information on VxVM error messages.

If you use the VERITAS VCS product, all disk group failover issues can be managed correctly. VCS includes a high availability monitor and includes failover scripts for Base-VXVM, Base-VXFS, and for several popular databases.

The `-t` option to `vxldg` prevents automatic re-imports on reboot, and is necessary when used with a host monitor (such as VCS) that controls imports itself, rather than relying on automatic imports by VxVM.

vxrecover Needs at Least One ACTIVE or CLEAN Plex to Start a Volume

- **Problem:** The `vxrecover` command starts a volume only if it has at least one plex that is in the ACTIVE or CLEAN state and is not marked STALE, IOFAIL, REMOVED, or NODAREC. If such a plex is not found, VxVM assumes that the volume no longer contains valid up-to-date data, so the volume is not started automatically. A plex can be marked STALE or IOFAIL as a result of a disk failure or an I/O failure.
- **Workaround:** In such cases use the following command:

```
# vxvol -f start volume
```
- However please determine what caused the problem before you run this command. It is likely that the volume needs to be restored from backup, and it is also possible that the disk needs to be replaced.

Free Space Reported Differently

- **Problem:** `vxldg free` reports free space differently than `vxprint-ht`.
- **Workaround:** None.

Subdisks are not Aligned on Cylinder Boundaries After a Relayout

- When `relayout` is performed on a volume, VxVM does not grow subdisks such that they end on cylinder boundaries. If you subsequently increase the size of the volume, its subdisks are not grown using contiguous disk space.

- To ensure that a volume's subdisks are grown using contiguous disk space, specify the attribute `layout=nodiskalign` to `vxassist`, as shown here:

NOTE Specifying `layout=nodiskalign` permanently enforces this layout policy on the volume.

Resizing Layered Volumes Fails while Resynchronization is Ongoing

- **Problem:** Due to the current implementation to handle the resize of layered volumes, it is recommended that you do not grow or shrink layered volumes (`stripe-mirror`, `concat-mirror`, and so on) while resynchronization is ongoing.

Internally, Base-VXVM converts the layout of layered volumes and updates the configuration database before it shrinks or grows their sizes. This causes any ongoing operation, such as the resynchronization, to fail.

- **Workaround:** If the system reboots before a layered volume completes growing or shrinking, the volume is left with an intermediate layout. In this case, the user has to use `vxassist` to restore the volume to its original layout.

After a layered volume is resized, the volume names, the plex names, and the subdisk names associated with the subvolumes, are changed.

Stopping and Starting Layered Volumes

If a layered volume is in use when a `vxvol stopall` command is issued, then only the sub-volumes are disabled. The layered volume remains enabled. When a `vxvol stop layered-volume` command is issued, then only the top layered volume is stopped. The sub-volumes remain enabled. When a `vxvol start layered-volume` command is issued, the top layered volume is started. The sub-volumes remain disabled.

VxVM Commands Do Not Always Show Current Status of VxVM Disks

- **Problem:** The VEA and frequently used Base-VXVM commands, such as `vxdisk` and `vxprint`, do not necessarily show the current status of disks managed by Base-VXVM. Base-VXVM builds and maintains a configuration database in system memory. This configuration database also includes Disk Access (DA) records with information about the disk devices obtained by the `vxconfigd` scan pass. Base-VXVM relies on the operating system kernel to notify it of disk status changes. The HP-UX kernel does not currently notify Base-VXVM of disk status changes.
- **Workaround:** Use either of the following commands to force an update of the Base-VXVM configuration database:

```
# vxdisk online diskname
```

This command updates the status of the diskname disk.

```
# vxdctl enable
```

This command updates the status of all the VxVM disks.

NOTE The `vxdctl` command initiates an entire disk device scan. Therefore the length of time Base-VXVM takes to scan all of the devices in the environment of that particular host will increase as the number of devices increases. If you know that a particular disk state has changed, it is faster to use `vxdiskonline diskname` to update that disk only.

Run `vxdctl enable` to Show Status Changes for LVM Disks

- Problem: Base-VXVM output will not reflect status changes for LVM disks until `vxdctl(1M)` is run. For example, if you clear an LVM disk with `pvremove(1M)`, the `vxdisk` command will still list the status of that disk as LVM, until you run the `vxdctlenable` command. This is also true for VEA output and the output from other Base-VXVM commands.
- Workaround: Run `vxdctlenable` after making any changes to LVM disks to update Base-VXVM's database.

Possible Incorrect I/O Counts On Objects

- Problem: The Base-VXVM configuration daemon may hang inside the kernel while processing a configuration change, because of incorrect object I/O counts. The I/O count could become incorrect due to a race condition in the Base-VXVM kernel. Since `vxconfig` is hanging inside the kernel, Base-VXVM utilities do not work.
- Workaround: None. Contact VERITAS Technical Support for help in analyzing this problem.

DMP Issues

Action Required After Swapping Fibre Channel Paths to an Array

If the Fibre Channel paths to an array become swapped, either by accident or design, this changes the `nportid` setting for the disks in that array. You must use the `fcmsutil` command (see the `fcmsutil(1M)` manual page) to reconfigure the disks before attempting to rediscover the VxVM configuration.

If the disks are not reconfigured using the `fcmsutil` command, the `ioscan`, `insf -e` and `vxctlenable` commands may either take a very long time to complete or they may hang. Messages should also be visible in `/var/adm/syslog/syslog.log` to indicate that the `fcmsutil` command should be used to correct the problem.

Take the following steps if the paths have swapped:

1. Use the following command to obtain the `N_Port_id` settings for the affected paths:

```
# fcmsutil TD_device_file get remote all | grep N_Port_id
```

This is illustrated by the following example:

```
# fcmsutil /dev/td0 get remote all | grep N_Port_id
```

```
Target N_Port_id is = 0x00006b
```

```
Target N_Port_id is = 0x00006c
```

```
Target N_Port_id is = 0x00006d
```

```
Target N_Port_id is = 0x00006e
```

```
Target N_Port_id is = 0x000071
```

```
Target N_Port_id is = 0x000072
```

```
Target N_Port_id is = 0x000073
```

```
Target N_Port_id is = 0x000074
```

```
Target N_Port_id is = 0x000075
```

```
Target N_Port_id is = 0x000076
```

2. Run the following command on each disk found in step 1, specified by its `N_Port_id` setting:

```
# fcmsutil TD_device_file replace_dsk nport_ID
```

This is illustrated by the following example:

```
# for nport in 6b 6c 6d 6e 71 72 73 74 75 76
```

```
> do
```

```
> fcmsutil /dev/td0 replace_dsk 0x0000$nport
```

```
> done
```

3. Repeat steps 1 and 2 for all other paths that are affected, such as `/dev/td1`.

DMP Lists Disabled Paths That Have Been Reused

- **Problem:** When one of the multiple paths or cables to a disk array is disconnected, fails, or is swapped with another path, and then that same path or cable is reconnected or replaced, it is possible that HP-UX will recognize the recovered path as a new path, not as the same path that has recovered. In this case, DMP will list twice as many paths: the new ones in the `ENABLED` state and the old ones (that is, from before the paths were swapped, removed or replaced) in the `DISABLED` state. I/O continues to be routed correctly.
- **Workaround:** None necessary. Base-VXVM DMP will not automatically clean up the paths that are no longer in use, or that are in the `DISABLED` state. When the host is rebooted, the DMP database will be rebuilt without the `DISABLED` path definitions.

Disabling DMP

The disabling of DMP is not supported. The `vxdmptdis` and `vxdmpten` utilities that were provided prior to Base-VXVM 3.2 to disable and enable DMP are no longer provided.

Removing DMP Disks

- **Problem:** The `vxdiskrm` command only logically removes a disk. It does not remove the disk from DMP. Therefore, the disk may re-appear after a `vxctlenable` even if the disk has been physically removed.
- **Workaround:** Restart using `vxconfigd`. This causes the DMP database to be updated, and the physically removed device will no longer appear.

DMP and HP Surestore Disk Array xp256

HP Surestore Disk Arrays xp256, xp512, and xp48 are supported with Dynamic Multipathing (DMP) in active/active mode. If Base-VXVM encounters devices with identical Base-VXVM Disk IDs, the VxVM software will give a preference to selecting the P-Vol as opposed to the S-Vol to resolve the conflict. This disk selection policy will work only for xp256 microcode revisions 5244 and above.

DMP and EMC Symmetrix Disk Array Support

The Dynamic Multipathing feature (DMP) of the Volume Manager works with the EMC Symmetrix disk array only if the disk array is configured in the Common Serial Number Mode. This is because only the Common Serial Number Mode configuration provides unique device identification on different paths for DMP to configure its database. Without this mode, DMP cannot provide the multipathing functionality. DMP cannot access EMC Symmetrix disks as individual disks with a single path it found by scanning the device information tree in the kernel.

For the DMP functionality to work with the EMC Symmetrix disk array, the following conditions must be met:

- EMC Symmetrix users must program their disk arrays in the Common Serial Number Mode before installing the Volume Manager package.
- The minimum level of microcode on the EMC Symmetrix should be:
 - 5265 or 5266 for Sym-4 on models 3330, 3430, and 3700
 - 5265 or 5266 for Sym-4.8 on models 3630, 3830, and 3930
 - 5566 for Sym-5 on models 8130, 8430, and 8730
- The Common Serial Number Mode must be enabled for all channels that participate in the DMP.
- Ultra SCSI and fast-wide differential (FWD) SCSI support has been tested on the EMC OSD (Open System Director) controller.

DMP and HP FC10 Disk System Support

Because of the way LUNs are distributed across controllers by HP-UX, the enable/disable controller feature of Base-VXVM does not function as expected for the following configurations of FC10 disk systems:

- Multiple daisy chained FC10 disk systems connected to a single host
- Multiple FC10 disk systems connected to a single host through a hub

If the `vxddmpadm disable ctlr=ctlr_name` command is used for the above configurations to disable I/O through a particular controller, I/O can be stopped through paths that were not intended by the user.

If each FC10 disk system is individually connected to the host (using separate HBAs), the above mentioned feature functions correctly and as expected.

DMP and JBOD Support

The Dynamic Multipathing feature on JBOD drives works only if the individual disks provide unique serial number information. If the individual disks do not provide unique serial numbers, even if the disks are configured for multipathing in hardware, the Base-VXVM does not provide multipathing support.

Device Issues

Rootability Cloning Script May Yield Too Many Processes Error

- Problem: Running the Base-VXVM rootability cloning script, `vxcp_lvmroot`, may yield the following error message if your `nproc` tunable is set to the default of 276:

```
/etc/vx/bin/vxcp_lvmroot[47]: The fork function failed. Too many  
already exist.
```

```
sh: The fork function failed. Too many processes already exist.
```

- Workaround: Increase the `nproc` static tunable to 1024 using `sam(1M)`, rebuild a new kernel, reboot your machine, and try the `vxcp_lvmroot` again.

Business Copy (BC) Limitation on XP Disk Arrays

- Problem: Base-VXVM will allow an import of the volumes from only one BC on a node where volumes have the same disk and group identification. If the Primary (P-Vol) is accessible, then this volume will be used. The Secondary (S-Vol) is only imported if the P-Vol is inaccessible at the time of import. There is currently no ability to change the disk group identification (group id) on XP BC volumes.
- Workaround: None

NOTE You cannot import the split BC to a machine on which P-Vol is being used if Base-VXVM and XP512 BCs are running together. The reason for this is that the split BC and the P-Vol have the same private region and Base-VXVM allows you to import only one of them by default it chooses P-Vol. If you want, however, you can import the split BC (S-Vol) onto a secondary host.

VxVM and Older Quantum Disk Drives

- Problem: The VxVM makes use of the kernel-to-kernel pass through `ioctl` SCSI command feature in HP-UX. Base-VXVM issues SCSI inquiry commands to devices on the system to recognize individual disks and sort out host to device connection pathways. Some older Quantum disks (models PD210S and PD425S) do not respond properly to SCSI inquiry command when the device is in certain states. As a consequence, these devices are not recognized by VxVM and, as such, cannot be used as disks for Base-VXVM. A `vxdisk`

Adding Swap Space Using VxVM Volumes The HP System Administration Manager (SAM) currently does not have the capability to add swap space using Base-VXVM volumes. Please refer to the VERITAS File System 4.1 and VERITAS Volume Manager 4.1 Installation Guide for more information and workarounds for this problem.

Enclosure-based Naming on Persistent Simple or Nopriv Disks On a system that has persistent simple or nopriv disks, these disks may be put into an error state if you change to enclosure-based naming. To recover these disks, run the `vxdarestore` command. For more information, refer to Issues Regarding Persistent Simple/Nopriv Disks with Enclosure-Based Naming in the *VERITAS Volume Manager 4.1 Administrator's Guide* and the `vxdarestore(1M)` manual page.

Ensuring that no SCSI-3 PGR Keys Exist on the Disk Array Before physically connecting an array to the host that is running Base-VXVM, you should ensure that there are no SCSI-3 PGR keys on the disk array LUNs. You can use the `vxfenadm` command to determine if any keys exist.

VEA Issues

The following issues have been identified as VEA problems:

Connecting to X Windows The following X Window System error may occur when starting VEA:

```
Xlib: connection to "hostname:0.0" refused by server Xlib: Client is not authorized to connect to Server
```

Workaround: Allow X server access by typing:

Restarting VEA after Obtaining New Licenses If, after installing and starting VEA, you obtain a new license (either explicitly, using `vxlicinst`, or implicitly, by installing bundles B9116AA, B9117AA, or B9118AA), you will not have access to the new licensed features until you restart the VEA service:

1. Stop the VEA backend service:

```
# /opt/VRTSob/bin/vxsvcctl stop
```

2. Start the VEA backend service:

```
# /opt/VRTSob/bin/vxsvcctl start
```

Updating Objects in Volume View

- Problem: Object updates in the Volume View may be incorrect.
- Workaround: Close then re-open the Volume View.

Updating Objects in Disk/Volume Map View

- Problem: Object updates in the Disk/Volume Map View may be incorrect.
- Workaround: Close then re-open the Disk/Volume Map View.

Mirroring Disks

- Problem: The Actions > Disk Mirror menu is incorrectly disabled if you do not have a full VxVM license.
- Workaround: Use the `vxmirror` command line to mirror the disk.

Using SAM to Launch the VEA Client

- Problem: It is currently not possible to launch the VEA client from SAM.
- Workaround: Until a patch to overcome this problem becomes available, you should use the command line interface to launch VEA.