

VERITAS Volume Manager 4.1 Release Notes

HP-UX 11i v2

Seventh Edition



Manufacturing Part Number: 5991-5288

June 2006

Printed in the United States

© Copyright 2006 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 Symantec Corporation.

1. VERITAS Volume Manager 4.1 Release Notes

| | |
|--|----|
| New Features in VERITAS Volume Manager 4.1 | 1 |
| Unsupported Features | 5 |
| Features requiring special licenses | 6 |
| Features of Base-VXVM compared to LVM | 6 |
| VERITAS Base-VXVM 4.1 Licenses | 7 |
| Feature Availability by Product | 7 |
| VEA Graphical User Interface and SAM | 8 |
| Coexistence with HP Logical Volume Manager (LVM) | 9 |
| Compatibility Information and Installation Notes | 9 |
| Supported Platforms | 9 |
| Software Requirements | 9 |
| Known Problems and Workarounds | 10 |
| Installation Issues | 10 |
| Utility Issues | 12 |
| Configuration Issues | 19 |
| DMP Issues | 20 |
| Device Issues | 24 |
| Hot-Relocation Issues | 26 |
| Cluster Functionality Issues | 27 |
| Data Integrity Issues | 28 |
| Snapshot and Snapback Issues | 28 |
| Intelligent Storage Provisioning Issues | 29 |
| Miscellaneous Issues | 29 |
| VEA Issues | 30 |
| Internationalization Issues | 33 |
| Patches and Fixes | 34 |

Contents

Preface

The *VERITAS Volume Manager 4.1 Release Notes* provides information on the September 2005 release of VERITAS Volume Manager 4.1.

Technical Support

For license information contact:

- Software License Manager
<http://licensing.hp.com/welcome.slm>
- HP Licensing Services contact
<http://licensing.hp.com/licenseAdmins.slm>
(Europe) contact:
Phone: +353.(0)91.75.40.06
Email: codeword_europe@hp.com
(U.S. and Canada)
Phone: +1 650.960.5111
Email: hplicense.na@hp.com
(Asia Pacific)
Phone: 0120.42.1231 or 0426-48-9310 (Inside Japan)
+81.426.48.9312 (Outside Japan)
Email: sw_codeword@hp.com

For latest information on available patches visit:

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

For technical support visit:

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

Typographic Conventions

Table 1 describes the typographic conventions used in this document.

Table 1 **Typographic Conventions**

| Typeface | Usage | Examples |
|---------------|--|--|
| monospace | Computer output, 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> manpage for more information. |
| <i>italic</i> | New terms, book titles, emphasis, variables replaced with a name or value | See the <i>VERITAS Volume Manager 4.1 Release Notes</i> for details. |
| % | C shell prompt | Not applicable |
| \$ | Bourne/Korn shell prompt | Not applicable |
| # | Superuser prompt (all shells) | Not applicable |
| \ | Continued input on the following line; you do not type this character | <code># mount -F vxfs \ /h/filesys</code> |
| [] | In command synopsis, brackets indicates an optional argument. | <code>ls [-a]</code> |
| | In command synopsis, a vertical bar separates mutually exclusive arguments. | <code>mount [suid nosuid]</code> |
| blue text | An active hypertext link | In PDF and HTML files, click on links to move to the specified location. |

Related Documentation

For more information about VERITAS 4.1 products refer to the following documents located in the `/usr/share/doc` directory:

- *VERITAS File System 4.1 Release Notes*
- *VERITAS File System 4.1 Administrator's Guide*

- *VERITAS Volume Manager 4.1 Hardware Notes*
- *VERITAS Volume Manager 4.1 Release Notes*
- *VERITAS Volume Manager 4.1 Troubleshooting Guide*
- *VERITAS Volume Manager 4.1 Migration Guide*
- *VERITAS Volume Manager 4.1 Administrator's Guide*
- *VERITAS Enterprise Administrator (VEA 500 Series) Getting Started*
- *VERITAS Storage Foundation 4.1 Cross-Platform Data Sharing Administrator's Guide*
- *VERITAS Flashsnap Point-In-Time-Copy Solutions Administrator Guide*

HP Encourages Your Comments

HP encourages your comments concerning this document. We are committed to providing documentation that meets your needs.

You can send comments to: netinfo_feedback@cup.hp.com

You must include document title, manufacturing part number, and any comment, error found, or suggestion for improvement you have concerning this document. Also, include what we did right so we can incorporate it into other documents.

1 VERITAS Volume Manager 4.1 Release Notes

This document provides information on VERITAS Volume Manager Release 4.1. The VERITAS Volume Manager 4.1 is supported on HP-UX 11i v2.

The following topics are discussed in this document:

- “New Features in VERITAS Volume Manager 4.1” on page 1
- “Unsupported Features” on page 5
- “Features of Base-VXVM compared to LVM” on page 6
- “VERITAS Base-VXVM 4.1 Licenses” on page 7
- “Compatibility Information and Installation Notes” on page 9
- “Known Problems and Workarounds” on page 10

New Features in VERITAS Volume Manager 4.1

New features in Base-VXVM 4.1 include the following:

- **Intelligent Storage Provisioning (ISP)**

The ISP Service is an alternative to the traditional method of creating and managing volumes. ISP creates volumes from available storage with capabilities that you specify by consulting the externally defined rule base for creating volumes, and comparing it to the properties of the storage that is available. Please refer VERITAS Storage Foundation 4.1 Intelligent Storage Provisioning Administrator’s Guide for more information.

- **Enhancements to Device Discovery Layer (DDL)**

The Device Discovery Layer was introduced in an earlier release of VxVM. This release enhances the DDL by adding the following functionality:

- 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 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.

- **Cross-Platform Data Sharing Feature (CDS)**

CDS allows the sharing of data between heterogeneous systems where each system has direct access to the physical devices that are used to hold the data. Please refer *VERITAS Storage Foundation Cross-Platform Data Sharing Administrator's Guide for more information*. CDS does not support VERITAS Volume Replicator (VVR).

- **Hitachi Extended Copy Facility**

The Hitachi E-Copy facility allows a host server to request data movement from disk to tape or to another disk without the overhead of server CPU time I/O during the actual movement of data.

- **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 recover from such a situation, you can use the `vxsplitlines` command to diagnose the problem. 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.

- **Dynamic Multipathing Enhancements**

- Coexistence of DMP with Third Party Drivers
- Failback is now a cluster-wide operation that provides significant performance improvements.
- A/PF type arrays are now supported by DMP in a clustered environment.
- Support for `B_FAILFAST` flag

DMP 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 times.

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.

- Performance improvements have been made to DMP.

- **Instant Snapshots**

Instant (copy-on-write) snapshots are easier to configure and offer more flexibility of use than traditional third-mirror break-off snapshots. Once created, an instant snapshot is immediately available for use. For preference, new volumes should be configured to use instant snapshots, but this also requires rewriting of administration scripts that assume the traditional third-mirror snapshot model. Legacy volumes can also be reconfigured to use instant snapshots.

- **Disability Requirements Conformance**

VxVM is compliant with the US Government's section 508 article.

- **Menu-driven Installation**

With this release, VxVM provides support for installation through a menu-driven front end. This enhancement, which will be common across all VERITAS products, is intended to provide a common look and feel to the installation experience. The traditional installation process using `pkgadd` is still supported, but users are encouraged to move to the menu-driven interface.

- **Enhancement of Local Detach Policy**

Local Detach policy now works uniformly on both master and slave nodes in a cluster. If the master node cannot handle the local detach policy, it leaves the cluster rather than disabling disk groups.

- **Default Disk Group Behavior in Base-VXVM 4.1**

In previous versions of VxVM, a disk group with the name `rootdg` (root disk group) was required. For VxVM to function, the `rootdg` disk group had to exist, and it 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 VxVM command is not specified with the `-g` option. The main benefit of creating a default disk group is that the Base-VXVM commands default to the default disk group and you will not need to use the `-g` option. To verify the default disk group after it has been created, enter the command:

```
# vxdg defaultdg
```

- **Simplification of `vxinstall`**

Because Base-VXVM no longer requires a `rootdg` to be configured, `vxinstall` has been much simplified, and you can no longer use this command to configure disk groups on your system after installation. You should 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:

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

- **Default Disk Format**

With the introduction of the Cross-Platform Data Sharing (CDS) feature, the format for a disk defaults to CDS rather than to simple, which was the format in earlier releases. In addition, the simple disk format has been renamed as `hpdisk`. For instructions on converting from a CDS format disk to a `hpdisk` format disk, see the *Creating and Administering Disk Groups* chapter in the *VERITAS Volume Manager Administrator's Guide*. The section, *Creating a disk Group*, includes information on disk formats and disk group formats.

CDS does not normally require an *Add-on* license, however if you want to import a disk group from a platform that is different from your own platform, you will need an *Add-on* license.

CDS disk groups have certain physical disk alignment restrictions for subdisk. For more information on this feature and CDS disk groups in general, see the *VERITAS Storage Foundation 4.1 Cross-platform Data Sharing Administrator's Guide*

A further 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`.

Unsupported Features

The following features are not supported with the OEM version of VERITAS Volume Manager 4.1 for HP-UX:

- VERITAS Volume Replicator (VVR)
- VERITAS Cluster Server (VCS)

Features requiring special licenses

- The FlashSnap (FastResync and Dynamic Disk Group Split and Join and Storage Checkpoints).
 - The VERITAS Cluster File System(CFS)
 - Quicklog
-

Features of Base-VXVM compared to LVM

The VERITAS Volume Manager 4.1 for HP-UX (Base-VXVM) includes many features that are not available with LVM on HP-UX. 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

VERITAS Base-VXVM 4.1 Licenses

The following table shows the supported features available with VERITAS Volume Manager 4.1 licensing

Table 1-1 Veritas Volume Manager Supported Features

| VxVM License | Description of Supported Features |
|---------------------|---|
| Base | Concatenation, spanning, rootability and root disk mirroring, multiple disk groups and VEA. |
| Full | Base features plus volume resizing, coexistence with native volume manager, striping, mirroring, DRL logging for mirrors, striping plus mirroring, mirroring plus striping, RAID-5, RAID-5 logging, Smartsync, hot sparing, hot-relocation, online data migration, online relayout, Intelligent Storage Provisioning, Storage Expert, Device Discovery Layer, dynamic multipathing (DMP). |

Unless you have installed the required license, you may not be able to use certain features of VERITAS Volume Manager 4.1. For example, you require a VERITAS Volume Manager 4.1 license to be able to create mirrored volumes other than the root disk.

Feature Availability by Product

The following table lists the features available with the VERITAS Volume Manager for HP-UX.

Table 1-2 VERITAS Volume Manager 4.1 for HP-UX Feature Availability

| Feature | VERITAS Volume Manager 4.1 for HP-UX | Required VxVM License |
|---------------------------------|---|------------------------------|
| Concatenation | Supported | Base |
| Java-based admin GUI | Supported | Base |
| Hot-relocation and unrelocation | Supported | Full |
| Mirroring (RAID-1) | Supported | Full |

Table 1-2 VERITAS Volume Manager 4.1 for HP-UX Feature Availability

| Feature | VERITAS Volume Manager 4.1 for HP-UX | Required VxVM License |
|--|---|------------------------------|
| Number of mirrors supported | Supported | Full |
| Online migration | Supported | Full |
| Online relayout | Supported | Full |
| Online resizing of volumes | Supported | Full |
| Path failover support (active/passive peripherals) | Supported | Full |
| RAID-5 | Supported | Full |
| Striped Mirrors (RAID 1+0) | Supported | Full |
| Striping (RAID 0) | Supported | Full |
| Task monitor | Supported | Full |

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: a server and a 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 version 2 system to manage Base-VXVM remotely. Note that only HP-UX 11i clients are supported.

SAM, the HP-UX system administration manager, and VEA exist as independent entities. The VEA client can and 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*.

WARNING **System Administration Manager (SAM) does not recognize a VxFS 4.1 file system with Disk Layout Version 6.**

Coexistence with HP Logical Volume Manager (LVM)

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

Both LVM and Base-VXVM utilities are aware of the other volume manager, 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, `vxxmconvert`, 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 `vxxmconvert`.

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, and 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 (for Itanium-based and PA-RISC systems) installed.

Known Problems and Workarounds

The following known problems and workarounds have been identified:

- “Installation Issues” on page 10
- “Utility Issues” on page 12
- “DMP Issues” on page 20
- “Device Issues” on page 24
- “Hot-Relocation Issues” on page 26
- “Cluster Functionality Issues” on page 27
- “Data Integrity Issues” on page 28
- “Snapshot and Snapback Issues” on page 28
- “Intelligent Storage Provisioning Issues” on page 29
- “Miscellaneous Issues” on page 29
- “VEA Issues” on page 30
- “Internationalization Issues” on page 33

Installation Issues

Upgrading Systems Running VxVM 3.5 Prior to Command Cumulative Patch 06

Before upgrading a system that is running under VxVM 3.5 at a patch level prior to Command Cumulative Patch 06 (PHCO_30834), it is strongly recommended that you download and apply this patch, and then run the `ckpublen.sh` utility script. If the script reports that any disks need to be re-initialized, back up the file systems and data residing on the volumes on those disks, and restore them after re-initializing the disks and recreating the volumes. You can then proceed to upgrade the system with the VERITAS Storage Foundation 4.1 software.

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, then it will request that you reinvoke the `vx_emerg_start` script, using the `-f` option, to perform the write.

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

- Determine which disk you booted from.
- Use the `vxdisk list` command to find your boot disk in the DEVICE column on the far left.
- 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.
- Use the `vxprint -g rootdg rootvol` command. If you have a mirrored root volume, you will see two lines with `pl` on the far left side. Look at each `pl` (or `plex`) entry and immediately below it will be the subdisk associated with the plex. It will start with an `sd` in the far left column.
- 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 |
| pl | rootvol-01 | rootvol | ENABLED | 524288 | - | ACTIVE |
| sd | rootdisk01-03 | rootvol-01 | ENABLED | 524288 | 0 | - |
| pl | 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.

- 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

Warning Message at Boot Time

A message such as the following is displayed if an attempt is made to open a volume at boot time before any disk group has been imported.

```
WARNING: VxVM vxio V-5-0-23 Open on an spurious volume device (hex_id) encountered. This device may be valid, but has not yet been configured in the kernel.
```

This message may be ignored. Once the disk group has been imported successfully, there should be no problem in accessing its volumes.

Shrinking a Swap Volume

`vxassist` has no built-in protection to prevent you from shrinking the swap volume without first shrinking what the system sees as available swap space. If it is necessary to shrink the swap volume, the operation must be done in single user mode and the system must be rebooted immediately. Failing to take these precautions can result in unknown system behavior or lock-up.

Adding a Log and Mirror to a Volume

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, but the log is silently omitted. To add a log and a mirror, add them in two separate `vxassist` invocations, as follows:

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

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

Using `vxdiskadm` to Replace a Failed Disk

The `vxdiskadm` command requires two attempts to replace a failed disk. The first attempt can fail with a message of the form:

```
/usr/lib/vxvm/voladm.d/bin/disk.repl: test: argument expected
```

The command is not completed and the disk is not replaced. If you now rerun the command, using Option 5, the replacement successfully completes.

Replacement of the old_layout Attribute

The `vxdisksetup` command gives the error message `Attribute unrecognized` when the `old_layout` attribute is used to make a disk into a VxVM controlled disk. The `old_layout` attribute is no longer supported. Use the `-noreserve` attribute instead .

Using vxvol and vxmend with Layered Volumes

The `vxvol` and `vxmend` commands do not handle layered volumes very well. When `vxmend` is executed on the top level volume to change the state of a volume, it is executed only on the top level volume; the change is not propagated to the lower level volumes. As a result, the volume states can become inconsistent and a subsequent `vxvol` `init` command might fail.

The `vxvol` command also exhibits the same problem. When a `vxvol` `init` command is executed on the top level volume, the change is not propagated to the volumes corresponding to its subvolumes.

Workaround:

When executing the `vxvol` or `vxmend` command on a layered volume, first issue the command to the lower level volumes in a bottom-up fashion; then execute the command on the top-level volume.

In this example, a volume, `vol`, has two subvolumes, `vol-L01` and `vol-L02`. The state of the volumes is first set to empty, and then the initialization commands are executed:

```
# vxmend -o force -g mydg fix empty vol
# vxmend -o force -g mydg fix empty vol-L01
# vxmend -o force -g mydg fix empty vol-L02
# vxvol -g mydg init zero vol
# vxvol -g mydg init zero vol-L01
# vxvol -g mydg init zero vol-L02
```

Growing or Shrinking Layered Volumes

Due to the current implementation of a resize of layered volumes, it is recommended that you do not grow or shrink layered volumes (for example; `stripe-mirror`, `concat-mirror`) during resynchronization. This limitation does not apply to ISP layered volumes.

Internally, VxVM converts the layout of layered volumes and updates the configuration database before it does the actual resize. This causes any ongoing operation, such as a resynchronization, to fail.

If the system reboots before the grow or shrink of a layered volume completes, the volume is left with an intermediate layout. In this case, you have to use `vxassist convert` to restore the volume to its original layout.

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

Maximum Size of a VxVM Volume

VxVM supports volume lengths up to $2^{63}-1$ disk sectors when using VERITAS-specific `ioctl` calls. However, system calls such as `seek`, `lseek`, `read` and `write` are limited to a maximum offset that is determined by the operating system. For a system that supports large files, this is usually $2^{63}-1$ bytes. Otherwise, the maximum offset value is usually $2^{31}-1$ bytes (1 byte less than 2 terabytes).

vxconfigd Hangs Due to a Faulty Disk

If I/O hangs for some reason such as a disk failing while the VxVM configuration daemon, `vxconfigd`, is performing I/O from/to the disks, there is no way to communicate with `vxconfigd` via signals or native interprocess communication. This can potentially cause two kinds of problem:

- The node becomes unavailable for VxVM administrative commands.
- In a clustered or HA environment where VERITAS Cluster Server agents need to communicate with `vxconfigd` to determine the health of VxVM components, service groups start timing out and failing.

Base-VXVM and Multi-Host Failover Configurations

Outside the context of clustering functionality, 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, so this is not a problem limited to 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 departs the disk group. The presence of the import lock prevents other hosts from importing the disk group until the importing host has departed 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 departing 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 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 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 `vxctl 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.

VxVM can support 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.) There are typically numerous configuration copies for each disk group, but corruption nearly always affects all configuration copies, so redundancy does not help in this case.

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

```
Association not resolved
Association count is incorrect
Duplicate record in configuration
Configuration records are inconsistent
```

These errors are typically reported in association with specific disk group configuration copies, but usually apply to all copies. The following is usually displayed 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 VxVM and VxFS and for several popular databases.

The `-t` option to `vxdbg` 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.

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, but the log is silently 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 ...
```

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, to force the volume to start, use the following command:

```
# vxvol -f start volume
```
- However, try to 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:** `vxdbg 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:

```
# vxassist growby volume length layout=nodiskalign
```

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, 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 the grow or shrink of a layered volume completes, the volume is left with an intermediate layout. In this case, the user has to use `vxassist convert` 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, then only 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 VxVM commands, such as `vxdisk` and `vxprint`, do not necessarily show the current status of disks managed by VxVM. 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. VxVM relies on the operating system kernel to notify it of disk status changes. The HP-UX kernel does not currently notify VxVM of disk status changes.
- **Workaround:** Use either of the following commands to force an update of the 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 `vxdtlenable` command initiates an entire disk device scan. Therefore the length of time 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 which disk's state has been changed, it is faster to use `vxdiskonline diskname` to update that disk only.

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

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

Possible Incorrect I/O Counts On Objects

- **Problem:** The 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 `vxconfigd` is hanging inside the kernel, VxVM utilities will not work.
- **Workaround:** None. Contact VERITAS Technical Support for help in analyzing this problem.

Configuration Issues

`vxconfigd` daemon hangs while configuring VxVM

If a disk fails while configuring VxVM with the `vxconfigd` daemon, the following problems occur:

- The VxVM administrative commands cannot be run on the node on which the disk has failed.
- Service groups timeout and fail in a clustered or HA environment as the VERITAS Cluster Server agents cannot communicate with `vxconfigd` daemon to determine the health of VxVM components.

Also, there is no way to communicate with the `vxconfigd` daemon through signals or native interprocess communication.

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 `vxdctlenable` 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 become 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 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 simply 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. 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 `vxdumpdis` and `vxdumpen` utilities that were provided prior to 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,

Known Problems and Workarounds

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 should be enabled for all channels that participate in the DMP. This is done at installation time in the Symmetrix configuration by the EMC Customer Engineer.
- Ultra SCSI and fast-wide differential (FWD) SCSI support has been tested on the EMC OSD (Open System Director) controller.

See your sales representative for further updates on hardware support.

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 `vxdmadm 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 VxVM does not provide multipathing support.

Usage of `dmp_failed_io_threshold` Parameter

Data loss can occur after the failure of a single path in a multipathed environment. This might happen because of an interaction between an Application or file system and Volume Manager. This issue applies to configurations in which:

- VxVM volumes are not mirrored.
- There are multiple paths to the disk.
- A portion of the I/O path below DMP fails in such a way that the error is not returned to DMP before 10 minutes have elapsed. Examples of failures that have been seen to cause this condition, include bad hardware (HBA, FCOT, GBIC, Switch Port, Array Controller) and HBA reconfigured to retry endlessly.

NOTE A portion of the I/O path below DMP fails in such a way that the error is not returned to DMP before 10 minutes have elapsed. Examples of failures that have been seen to cause this condition, include bad hardware (HBA, FCOT, GBIC, Switch Port, Array Controller) and HBA reconfigured to retry endlessly.

If the delay in returning the I/O is caused by a problem in the I/O path to the device rather than the device itself, DMP will incorrectly return the error to the VxVM layer rather than retrying the I/O on another path. If the volume is mirrored, VxVM will satisfy the I/O from the other plex, and detach the plex that failed and prevented the volume from becoming hung.

If the volume is not mirrored, the error will be passed to the File System or application layer. This can result in the File System marking inodes for deletion when they are still valid. If raw volumes are in use, the application might believe that the data on the disk is corrupted when it is actually clean.

To prevent this possibility in situations where mirrored volumes are not used, the threshold should be tuned to a sufficiently high value that is unlikely to be reached. For example, to change the value of `dmp_failed_io_threshold` to 16 hours (57600 seconds), modify the value defined in `/kernel/drv/vxdmp.conf` as shown here:

```
dmp_failed_io_threshold=57600
```

After changing the value, reboot the system.

In situations in which mirrored volumes are in use, and an application time-out is being hit when there is still a valid plex with the data, the value of `dmp_failed_io_threshold` can be tuned to a smaller value so that the I/O can succeed on the mirror without triggering an application failure.

Device Issues

Rootability Cloning Script May Yield Too Many Processes Error

- Problem: Running the 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 In the case where VxVM and XP512 BCs are being used together; if the BCs are split, you cannot then import the split BC to the machine on which P-Vol is currently being used. The reason for this is that the split BC and the P-Vol will have the same private region, and 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. 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 VxVM. A `vxdisk` command may list the device in error state or may not list at all.

- Workaround: Do not attempt to define the device for VxVM.

Adding Swap Space Using VxVM Volumes

The HP System Administration Manager (SAM) currently does not have the capability to add swap space using VxVM volumes. Please refer to the VERITAS File System (HP OnlineJFS/JFS) 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.

Unsupported Disk Arrays

To ensure that DMP is set up correctly on a multiported JBOD or other disk array that is not supported by VxVM, use the procedure given in "Adding Unsupported Disk Arrays to the DISKS Category" in the Administering Disks chapter of the VERITAS Volume Manager Administrator's Guide. Otherwise, VxVM treats the independent paths to the disks as separate devices, which can result in data corruption.

Hitachi Arrays in Active/Active Mode

When Hitachi DF400 and DF500 arrays are configured in Active/Active mode, performance is degraded.

HDS9500V Array Support Library

In VxVM 3.5, support for the Hitachi Data Systems HDS9500V array was provided by the `VRTShds9500v` Array Support Library (ASL). Release 4.1 of VxVM includes built-in support for the HDS9500V array. If the `VRTShds9500v` ASL is currently installed on your VxVM 3.5 system, use the following command to remove this ASL before installing VxVM 4.1:

```
# swremove VRTShds9500v
```

Known Problems and Workarounds

If you have already upgraded to VxVM 4.1 with the `VRTShds9500v` ASL installed, remove this ASL and reboot the system. This will ensure that the ASL provided with VxVM 4.1 is the one that is claiming your HDS9500V devices.

Adding HP-EVA Disks

When HP-EVA disks are added to VxVM 4.1, debug messages such as the following are displayed:

```
# vxdctl enable
Printing Name-Value Pair
CAB_SERIAL_NO : 50001FE100270DF0Printing Name-Value Pair
CAB_SERIAL_NO : Printing Name-Value Pair
Printing Name-Value Pair
Printing Name-Value Pair
CAB_SERIAL_NO : Printing Name-Value Pair 50001FE100270DF0
LUN_SERIAL_NO : 50001FE100270DF0 600508B40010293D00006000012A0000Printing
Name-Value Pair
Printing Name-Value Pair
Printing Name-Value Pair
.
.
.
```

These messages are harmless and can be ignored.

Hot-Relocation Issues

Impact of Hot-Relocation on Performance

Except for `rootvol` and `swapvol`, hot-relocation does not guarantee the same layout of data or performance after relocation. It is therefore possible that a single subdisk that existed before relocation may be split into two or more subdisks on separate disks after relocation (if there is not enough contiguous space on a single disk to accommodate that subdisk).

Disk Information in Notification Messages

When a disk failure occurs, the hot-relocation feature notifies the system administrator of the failure and any relocation attempts through electronic mail messages. The messages typically include information about the device offset and disk access name affected by the failure. However, if a disk fails completely or a disk is turned off, the disk access name and device offset information is not included in the mail messages. This is because VxVM no longer has access to this information.

Cluster Functionality Issues

Volume in SYNC state

If a node leaves the cluster while a plex is being attached to a volume, the volume can remain in the SYNC state indefinitely. To avoid this, after the plex attach completes, resynchronize the volume manually with the following command:

```
# vxvol -f resync volume
```

RAID-5 Volumes

VxVM does not currently support RAID-5 volumes in cluster-shareable disk groups.

File Systems Supported in Cluster-Shareable Disk Groups

The use of file systems other than VERITAS Storage Foundation Cluster File System (SFCFS) on volumes in cluster-shareable disk groups can cause system deadlocks.

Reliability of Information About Cluster-Shareable Disk Groups

If the `vxconfigd` program is stopped on both the master and slave nodes and then restarted on the slaves first, VxVM output and VEA displays are not reliable until the `vxconfigd` program is started on the master and the slave is reconnected (which can take about 30 seconds). In particular, shared disk groups are marked disabled and no information about them is available during this time. The `vxconfigd` program must therefore be started on the master first.

Messages Caused by Open Volume Devices

When a node aborts from the cluster, open volume devices in shared disk groups on which I/O is not active are not removed until the volumes are closed. If this node later joins the cluster as the master while these volumes are still open, the presence of these volumes does not cause a problem. However, if the node tries to rejoin the cluster as a slave, this can fail with the following error message:

```
cannot assign minor #
```

This message is accompanied by the console message:

```
WARNING:minor number ### disk group group in use
```

Data Integrity Issues

Disks with Write-Back Caches

Disk drives configured to use a write-back cache, or disk arrays configured with volatile write-back cache, exhibit data integrity problems. The problems occur after a power failure, SCSI bus reset, or other event in which the disk has cached data, but has not yet written it to non-volatile storage. Contact your disk drive or disk array manufacturer to determine whether your system disk drives use a write-back cache, and if the configuration can be changed to disable write-back-caching.

Snapshot and Snapback Issues

Using Snapshots as Root Disks

It is recommended that you do not use snapshots of the root volume as a bootable volume. A snapshot can be taken to preserve the data of the root volume, but the snapshot will not be bootable. The data from the snapshot would have to be restored to the original root volume before the machine could be booted with the preserved data.

Warning Message when Taking a Snapshot of a SFCFS File System

When taking a snapshot of a SFCFS file system, the following warning message might appear:

```
VxVM vxio WARNING V-5-0-4 Plex plex detached from volume volume
```

Workaround: No action is required. This behavior is normal and is not the result of an error condition.

File System Check of a Snapshot

Normally, a file system would have no work to do when a snapshot is taken. However, if an SFCFS file system is not mounted, it is likely that the `fsck` of the snapshot will take longer than is usually necessary, depending on the I/O activity at the time of the snapshot.

Workaround: When taking a snapshot of an SFCFS file system, you should ensure that at least one of the volumes defined in the command line is mounted on the cluster master.

Mount Operation Can Cause Inconsistencies in Snapshots

Inconsistencies can arise in point-in-time copies if any of the following snapshot operations are performed on a volume while a file system in the volume is being mounted: `vxassist snapshot`, `vxplex snapshot`, `vxsnap make`, `vxsnap refresh`, or `vxsnap restore`.

Intelligent Storage Provisioning Issues

Intelligent Storage Provisioning Issues

To create application volumes successfully, the appropriate licenses must be present on your system. For example, you need a full VERITAS Volume Manager and a VERITAS FlashSnap license to use the instant snapshot feature. Vendors of disk arrays may also provide capabilities that require special licenses for certain features of their hardware.

Miscellaneous Issues

Auto-import of Disk Groups

If a disk that failed while a disk group was imported returns to life after the group has been deported, the disk group is auto-imported the next time the system boots. This contradicts the normal rule that only disk groups that are (non-temporarily) imported at the time of a crash are auto-imported.

If it is important that a disk group not be auto-imported when the system is rebooted, the disk group should be imported temporarily when the intention is to deport the disk group (for example, in HA configurations). Use the `-t` flag to `vxchg import`

Volumes Not Started Following a Reboot

During very fast boots on a system with many volumes, `vxconfigd` may not be able to auto-import all of the disk groups by the time `vxrecover -s` is run to start the volumes. As a result, some volumes may not be started when an application starts after reboot.

Workaround: Check the state of the volumes before starting the application, or place a sleep (`sleep sec`) before the last invocation of `vxrecover`.

Forcibly Starting a Volume

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. In such cases, to force the volume to start, use the following command:

```
# vxvol -f start volume
```

However, try to 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.

Failure of Memory Allocation

On machines with very small amounts of memory (32 megabytes or less), under heavy I/O stress conditions against high memory usage volumes (such as RAID-5 volumes), a situation occurs where the system cannot allocate physical memory pages any more.

Messages About VVR Licenses

The following messages may get displayed on the console during a system reboot or during VxVM initialization when you are running vxinstall:

```
No VVR license installed on the system; vradmind not started
```

```
No VVR license installed on the system; in.vxrsyncd not started
```

These messages are informational only, and can be safely ignored if you are not a VERITAS Volume Replicator (VVR) user.

Number of Columns in a RAID-5 ISP Volume

If an ISP volume is created with the RAID-5 capability, the parameters `ncols` and `nmaxcols` refer only to the number of data columns, and do not include the parity column. For this reason, the actual number of columns that are created in such a volume is always one more than the number specified.

VEA Issues

The following issues have been identified as VEA problems:

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 newly 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/vxsvcctrl 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.

Name Service Switch Configuration File

For VEA to operate successfully, the name service switch configuration file, `/etc/nsswitch.conf`, must be present on the system. See the `nsswitch.conf(4)` manual page for more information.

Accessing the Task Log

The task log accessed from the Log tree is not supported. At this time, entries are written to the log file in `/var/vx/isis/command.log`.

Setting a Comment on an ISP Volume

If you create a new ISP volume by right-clicking on a user template and selecting the New Volume menu item, a comment that you specify to the Create Volume Dialog is not set on the volume. To specify a comment for the newly created volume, select the volume, choose Properties from the pop-up menu, enter a comment in the Comment field and then click OK.

Administering a Cache Volume Created on an ISP Volume

It may not be possible to use the VEA GUI to add or remove mirrors to or from a cache volume (used by space-optimized instant snapshots) that is created on an ISP volume, or to delete a cache volume. The cache object, but not the cache volume, is visible in the graphical interface.

Workaround: Stop and restart the VEA server.

Permitting Remote Access to the X Windows Server

The following X Windows 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 access to the local X server by using the following command:

```
# xhost +[hostname]
```

Disk Group Creation Failure with Duplicate Disk ID

VEA fails to create a disk group with a duplicate disk ID, and gives no other options.

Creating non-CDS disks

By default, VEA always creates CDS disk groups, and will therefore initialize disks with the CDS format. In order to create non-CDS disk groups through VEA, you need to create the following two files (or, edit the files accordingly if they already exist). This changes the default format setting for disk creation from `format=cdsdisk` in the `/etc/default/vxdisk` file, and the default CDS attribute setting from `cds=on` in the `/etc/default/vxdg` file.

To create non-CDS disk groups through VEA:

1. Create the file `/etc/default/vxdisk`, and add a line describing the desired non-CDS disk format as:

```
format=simple
```

```
or
```

```
format=hpdisk
```

2. Create the file `/etc/default/vxdg`, and add a line describing the CDS attribute as:

```
cds=off
```

To create CDS disk groups at a later stage, you will need to change the settings back again to `format=cdsdisk` and `cds=on`. For more information, see the VERITAS Storage Foundation Cross-Platform Data Sharing Administrator's Guide.

Incorrect vxpool command

The VEA GUI may incorrectly show the `-p` option as an argument to the `vxpool list` command, although the command is not actually invoked.

Internationalization Issues

Comments in Japanese on a Snapshot Volume Are Not Saved or Displayed Correctly

Comments that are entered in the Japanese character set in the Snapshot Options dialog of the Create Instant Snapshot screen of the VEA GUI are not saved or displayed correctly.

Inaccuracies in ISP Attribute Fields

The ISP User Template Wizard shows two "attribute value" fields rather than one "attribute value" and one "attribute name" field.

Patches and Fixes

The VEA GUI patches PHCO_33080 and PHCO_33081 are not mandatory for customers using VEA on Itanium® 2 montecito platform.