

# VERITAS Volume Manager™ 3.5

---

## Release Notes

Solaris

July 2002  
30-000744-011

  
VERITAS

---

## Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

## Copyright

Copyright © 2002 VERITAS Software Corporation. All rights reserved. VERITAS, VERITAS SOFTWARE, the VERITAS logo and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation in the USA and/or other countries. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

VERITAS Software Corporation  
350 Ellis Street  
Mountain View, CA 94043  
Phone 650-527-8000  
Fax 650-527-2908  
[www.veritas.com](http://www.veritas.com)



# Contents

---

Important Notice for VERITAS Volume Replicator Users .....	1
Organization .....	2
Getting Help .....	3
Using VRTSexplorer .....	5
Searching the VERITAS Knowledge Base .....	6
Conventions .....	8
Upgrading Disk Group Versions .....	9
Important Software Patch Information .....	10
Solaris Patches for VxVM .....	10
DMP Coexistence with Alternate Pathing .....	12
VxVM Point Patches .....	12
New Features .....	13
Changed Features .....	14
VxVM 3.5 Licenses .....	14
Photon Licensing .....	15
Termination of Support Statements .....	16
Termination of Support for VxVM Releases .....	16
Termination of Sun Support for Solaris 2.6 .....	16
Additional Changes in Hardware and Format Support .....	17
Software Fixes and Enhancements in VxVM 3.5 .....	18
Software Limitations and Problems in VxVM 3.5 .....	19
Disk Controller Firmware Upgrades .....	19
T3B Firmware Upgrade on Solaris 9 .....	20



---

Installation Issues . . . . .	22
Upgrade Issues . . . . .	23
Utility Issues . . . . .	27
Device Issues . . . . .	29
Hot-Relocation Issues . . . . .	32
DMP Issues . . . . .	33
Cluster Functionality Issues . . . . .	34
Data Integrity Issues . . . . .	35
Snapshot and Snapback Issues . . . . .	36
Miscellaneous Issues . . . . .	37
Solaris Issues . . . . .	39
Replacing Failed Disks . . . . .	41
VEA . . . . .	43
VEA Issues . . . . .	44
Adding Unsupported Disk Arrays to the JBOD Category . . . . .	45
Encapsulating and Mirroring the Root Disk . . . . .	46
Full Disk Mirroring Considerations . . . . .	46
Booting From DMP Devices . . . . .	48
VxVM and Multi-Host Failover Configurations . . . . .	48
Import Lock . . . . .	48
Failover . . . . .	49
Corruption of Disk Group Configuration . . . . .	49
Enclosure-based Naming . . . . .	50
Enclosure-based Naming on Persistent Simple or Nopriv Disks . . . . .	51
Working with VxVM Disks on a Sun StorEdge T3 Array . . . . .	51
Available Documentation . . . . .	52
Displaying Documentation Online . . . . .	53
English Versions of the Documentation . . . . .	53
Printing Documentation . . . . .	54



# VERITAS Volume Manager™ Release Notes

---

This document provides release information for the VERITAS Volume Manager™ (VxVM®) Release 3.5. This release also includes the new VERITAS Enterprise Administrator™ (VEA™) graphical user interface.

Release Notes are not installed along with any Volume Manager package. VERITAS recommends that you copy this document to the `/opt/VRTS/doc` directory so that the information is available for your future reference.

VxVM product enhancement is an ongoing process. To check for any additional information about this release, please use the steps shown in the section called “New Features” on page 13.

---

**Note** Before you install the packages, please review this entire document.

---

## Important Notice for VERITAS Volume Replicator Users

If you are using VERITAS Volume Replicator (VVR) and VRAS, *do not upgrade to VxVM 3.5* until you read the *VERITAS Volume Replicator Release Notes* for installation requirements.

## Organization

- u Getting Help
- u Conventions
- u Upgrading Disk Group Versions
- u Important Software Patch Information
- u New Features
- u Changed Features
- u VxVM 3.5 Licenses
- u Termination of Support Statements
- u Software Fixes and Enhancements in VxVM 3.5
- u Software Limitations and Problems in VxVM 3.5
- u VEA
- u Adding Unsupported Disk Arrays to the JBOD Category
- u Encapsulating and Mirroring the Root Disk
- u Booting From DMP Devices
- u VxVM and Multi-Host Failover Configurations
- u Enclosure-based Naming
- u Working with VxVM Disks on a Sun StorEdge T3 Array
- u Available Documentation



## Getting Help

If you have any comments for or problems with VERITAS products, contact VERITAS Technical Support:

- u U.S. and Canadian Customers: 1-800-342-0652
- u International Customers: +1 (650) 527-8555
- u E-mail: [support@veritas.com](mailto:support@veritas.com)

For license information:

- u Phone: 1-925-931-2464
- u Email: [license@veritas.com](mailto:license@veritas.com)
- u Fax: 1-925-931-2908

For software updates:

- u Email: [swupdate@veritas.com](mailto:swupdate@veritas.com)

For late-breaking news about this release, please see the section called "New Features" on page 13.

For additional technical support information, including

- TechNotes
- Product alerts
- Hardware compatibility lists

visit the VERITAS Technical Support Web site at:

<http://support.veritas.com>

For information about VERITAS products and services:

Phone 1-800-258-UNIX (1-800-258-8649) or 1-650-527-8000

Email [vx-sales@veritas.com](mailto:vx-sales@veritas.com)

For information about the Array Support Library

Before using any vendor supplied Array Support Library (ASL) provided on your product CD, you should visit the VERITAS Technical Support site for information on the latest ASL packages:

<http://support.veritas.com>

Click on the Knowledge Base Search section, and select Volume Manager for UNIX from the Search Product pull-down menu. Enter the search phrase "VERITAS enabled arrays" (quotes must be included in the search).

Any newly issued third party array drivers can later be found at the same site.



For additional information about VERITAS and VERITAS products, visit the Web site at:

<http://www.veritas.com>





## Using VRTSexplorer

The VRTSexplorer program can help VERITAS Technical Support engineers diagnose the cause of technical problems associated with VERITAS products. You can download this program from the VERITAS FTP site or install it from the VERITAS Installation CD. For more information, consult the README file in the `support` directory on the VERITAS Installation CD.

### Downloading from the Web

1. Use a web browser or the `ftp` program to download the VRTSexplorer program from the following URL:

```
ftp://ftp.veritas.com/pub/support/vxexplore.tar.Z
```

Save the file to a temporary directory such as `/tmp` as shown in the example session below.

2. Log in as `root` on the system that is experiencing the problem. Extract the contents of the downloaded file to the directory `/tmp/VRTSexplorer`:

```
# cd /tmp
# zcat vxexplore.tar.Z | tar xvf -
```

3. Run the VRTSexplorer program located in the VRTSexplorer directory:

```
# /tmp/VRTSexplorer/VRTSexplorer
```

4. When VRTSexplorer prompts you for a destination directory for the information that it collects, press Return to accept the default directory `/tmp`, or enter a path name of your own choice. VRTSexplorer writes the results of its analysis to a compressed tar file named `VRTSexplorer_casenumbe_r_hostname.tar.Z` in the specified directory.

### Installing from the CD

---

**Note** VRTSexplorer is part of VRTSspt package

---

VRTSspt is a group of support tools for troubleshooting a system and collecting information on its configuration. This package is optional; the tools are not required for operation of any VERITAS software. In general, the tools only gather configuration information.

Refer to the `README.VRTSspt` file in the `/support` directory for more information

To install VRTSspt package and VRTSexplorer:



1. Log in as superuser.
2. Place the VERITAS software CD into a CD-ROM drive connected to your system.
3. Load and mount the CD-ROM by starting the `volmgt` daemon. After starting the daemon, the path is `/cdrom/CD_name`.  

```
# /etc/init.d/volmgt start
```
4. Move to the `/volume_manager/support` directory and install the `VRTSspt` package:  

```
# cd /cdrom/volume_manager/support  
# pkgadd -d. VRTSspt
```
5. The program is installed in the `/opt/VRTSspt` directory. To run the program, enter:  

```
# /opt/VRTSspt/VRTSexplorer/VRTSexplorer
```
6. To remove the `VRTSspt` package, enter:  

```
# pkgrm VRTSspt
```

## Using the File Upload Facility

1. Use the file upload facility of your web browser, or the `ftp` program, to transfer the `VRTSexplorer` output file to the VERITAS Technical Support anonymous FTP site:  

```
ftp://ftp.veritas.com/incoming
```
2. Telephone VERITAS Technical Support at the number listed under “Getting Help” on page 3. Tell them that you have run `VRTSexplorer` and provide the name of the file that you transferred to the FTP site.

Alternatively, if you have already been assigned a call ID number by VERITAS Technical Support, send email to `support@veritas.com` and. Include your case ID number in the subject line.

## Searching the VERITAS Knowledge Base

The VERITAS Knowledge Base is an online document resource center that contains all publicly available information on VERITAS products. To access this information:

1. Go to the VERITAS Technical Support website at <http://support.veritas.com>.
2. Click on the **Knowledge Base Search** section.



3. Select **Volume Manager for UNIX** from the **Search Product** pull-down menu and enter the search phrase “late breaking information”.

This will show you any information that has been amended or added to the VxVM 3.5 product documentation.



## Conventions

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

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> manual page for more information.
<b>monospace (bold)</b>	User input	<b># mount -F vxfs /h/filesys</b>
<i>italic</i>	New terms, book titles, emphasis, variables replaced with a name or value	See the <i>User's Guide</i> for details. The variable <i>ncsize</i> determines the value of...

Symbol	Usage	Examples
%	C shell prompt	
\$	Bourne/Korn shell prompt	
#	Superuser prompt (all shells)	
\	Continued input on the following line	<b># mount -F vxfs \ /h/filesys</b>
[ ]	In a command synopsis, brackets indicates an optional argument	<code>ls [-a]</code>
	In a command synopsis, a vertical bar separates mutually exclusive arguments	<code>mount [suid nosuid]</code>



## Upgrading Disk Group Versions

All disk groups have a version number associated with them. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version, so you need to upgrade existing disk groups before you can perform the tasks. The following table summarizes the disk group versions that correspond to each VxVM release from 2.0 onward:

VxVM Release	Cluster Protocol Versions	Disk Group Version	Supported Disk Group Versions
2.0	n/a	20	20
2.2	n/a	30	30
2.3	n/a	40	40
2.5	n/a	50	50
3.0	n/a	60	20-40, 60
3.1	n/a	70	20-70
3.1.1	10, 20	80	20-80
3.2	30	90	10-90
3.5	40	90	20-90

You can get a disk group version listing by specifying a disk group name with this command:

```
# vxdg list disk_group_name
```

You can also determine the disk group version by using the `vxprint(1M)` command with the `-l` format option.

To upgrade a disk group, use the following command:

```
# vxdg -T version upgrade disk_group_name
```

Unless a disk group version is specified, this command upgrades the disk group to the highest version supported by the VxVM version on your system.

For shared disk groups, the latest disk group version is only supported by the latest cluster protocol version. See *Upgrading Cluster Functionality* in the *VERITAS Volume Manager Administrator's Guide* for more information on changing the cluster protocol version.

To see the current cluster protocol version, type:

```
# vxdctl support
```



## Important Software Patch Information

---

**Caution** Please read this section completely *before you install* VERITAS Volume Manager (VxVM) 3.5. You must reyour system after installing the patches. If you have already installed VxVM, install the patches and then reboot your system.

---

---

**Note** To install VEA on Solaris 2.6, 7, or 8, you must download and install the appropriate JAVA 1.3.1 patches for Solaris 2.6, 7, or 8, refer to <http://seer.support.veritas.com/docs/248266.htm>

---

### Solaris Patches for VxVM

---

**Caution** Do not install VxVM until you have installed the following mandatory patches. The patches must be installed on your system for VxVM to work properly.

---

#### Check for Existing Patches

Some required system patches may already be present in your operating system. To display a list of patches that are already installed on your system, enter:

```
# showrev -p
```

For more information, see `showrev(1M)`.

#### Download Required Patches

If the patches shown in the required list are not already installed, go to <http://sunsolve.Sun.com> to download them. You need to install the appropriate patches and then reboot.

**DISCLAIMER:** Patch version and information is determined at the time of product release. For the most current patch version and information, please contact your vendor.

#### Required Solaris Patches

For Solaris Version	Sun Patch ID Number	Notes
Solaris 9	None	



## Required Solaris Patches

For Solaris Version	Sun Patch ID Number	Notes
Solaris 8	SUNWsan package (see your vendor to obtain) 109529 (-06 or higher) 111413 (-06 or higher)	Regardless of whether you have an A5K drive connected, you must install these A5K patches for VxVM to function. You must install the SUNWsan package and patch 109529 before installing patch 111413. You can then install VxVM. This order of steps is important.  <b>Note</b> To get the SUNWsan package, you need to go to <a href="http://www.sun.com/storage/san/">http://www.sun.com/storage/san/</a> then click on link: "Sun StorEdge SAN 2.0 release Software/Firmware Upgrades and Documentation." You have to be a registered user, and it is free.
	108528 (-14 or higher)	Patch 108528 requires many other patches, and each of these could have its own patch prerequisites. If these prerequisite patches are missing then additional solaris patches may be needed to be installed.
	108827 (-19 or higher)	
	110722 (-01 or higher)	If you have Sun's Alternate Pathing (AP), you must install Solaris patch 110722-01 for VxVM 3.5 to co-exist with the AP driver version 2.3.1. VxVM 3.5 works only with AP driver version 2.3.1. Please see "DMP Coexistence with Alternate Pathing" before installing VxVM.
Solaris 7	106541 (-19 or higher)	
	107473 (-07 or higher)	Regardless of whether you have an A5K drive connected, you must install an A5K patch for VxVM to function. You should install this patch before installing VxVM
	110722 (-01 or higher)	If you have Sun's Alternate Pathing (AP), you must install Solaris patch 110722-01 for VxVM 3.5 to co-exist with the AP driver version 2.3.1. VxVM 3.5 works only with AP driver version 2.3.1. Please see "DMP Coexistence with Alternate Pathing" before installing VxVM.



### Required Solaris Patches

For Solaris Version	Sun Patch ID Number	Notes
Solaris 2.6	105181 (-31 or higher)	
	105375 (-26 or higher)	Recommended if you have an A5K drive connected. Install this patch before installing VxVM.
	110722 (-01 or higher)	If you have Sun's Alternate Pathing (AP), you must install Solaris patch 110722-01 for VxVM 3.5 to co-exist with the AP driver version 2.3.1. VxVM works only with APdriver version 2.3.1. Please see "DMP Coexistence with Alternate Pathing" before installing VxVM.

## DMP Coexistence with Alternate Pathing

VxVM DMP co-exists with Sun's Alternate Pathing (AP) driver 2.3.1. For VxVM versions 3.1.1 and later to function, the DMP driver must always be present on the system. You need to upgrade AP to 2.3.1 and install Solaris patch 110722-01 before upgrading to this release of VxVM.

---

**Note** This patch is required for DMP to co-exist with AP 2.3.1. VxVM version 3.5 does not support earlier versions of AP.

---

To confirm which version of the AP driver is installed, run the following command from the prompt:

```
# pkginfo -l SUNwapdv | grep VERSION
```

### Upgrading from VxVM 3.1 and Earlier

If you were using VxVM 3.1 and earlier, upgrading to VxVM 3.5 automatically enables DMP if it was previously disabled.

### Upgrading from VxVM 3.1.1 or VxVM 3.2

The DMP configuration will remain the same after upgrading from 3.1.1 or 3.2.

## VxVM Point Patches

Please refer to the `Patches` directory for information about the point-patches included in this release. You can find the `Patches` directory at the same level as the `pkgs` and `scripts` directories on your CD-ROM: `/cdrom/cdrom0/volume_manager/Patches`.





## New Features

New features in VxVM 3.5 include the following:

- u GAB messaging

A cluster environment can now use GAB as the transport agent for messaging. This means that you do not need to assign IP addresses for the endpoints of your cluster's private network when you configure the clustering functionality of VxVM.

- u New licensing package

A new licensing package, `VRTSvlic`, is introduced. In time `VRTSvlic` will become the common licensing solution for all VERITAS products. Existing license keys are automatically translated into license keys compatible with the `VRTSvlic` package. The old and new licensing packages can co-exist on the same system to support older VERITAS products that do not understand the new licensing utilities. For more information, see the *VERITAS Volume Manager 3.5 Installation Guide*.

- u Support for Sun JumpStart and Live Upgrade

Support for the `-R root_path` option to the `pkgadd` and `pkgrm` commands is implemented in the VxVM product installation scripts. VxVM packages now meet Sun's remote installation package requirements as specified in the document *Level 0 Solaris Package Requirements*. VxVM packages also comply with the *Advanced Packaging Guidelines*, and are compatible in JumpStart and Live Upgrade environments.

- u Support for Solaris 9

VxVM 3.5 installs on Solaris 9 systems. For more information about upgrade scenarios, see the *VERITAS Volume Manager 3.5 Installation Guide*.

- u VxVM tunables

The values of various tunables are adjusted to match the performance expected from current systems. In addition, VxVM tunables are now adjusted by setting their values in the driver configuration files, `/kernel/drv/vxio.conf` and `/kernel/drv/vxdmp.conf`. For more information about configuring VxVM tunables, see the "Performance Monitoring and Tuning" chapter in the *VERITAS Volume Manager 3.5 Administrator's Guide*.

- u VERITAS Enterprise Administrator (VEA)

VEA replaces the VERITAS Volume Manager Storage Administrator (VMSA) as the graphical configuration tool for VxVM.

- u Simplified Upgrade Procedure

You now have the choice of upgrading to VM 3.5 using either `pkgadd`, or the `upgrade_start/upgrade_finish` method. See the *VERITAS Volume Manager 3.5 Installation Guide* for details.

- u DMP Support for SERVPoint SAN



The SERVPoint SAN product can now be connected to any server or platform (supported by VxVM 3.2 or higher) and appear as if that server or platform was a standard supported disk array.

---

**Note** For information about new features in VERITAS Volume Replicator (VVR), please see the VVR Release Notes.

---

- u The DMP feature of VxVM supports the following arrays:

Sun StorEdge 3900 and 6900 Series (see <http://docs.sun.com/db?p=coll/992.1> for documentation)

- u Restoring VERITAS Volume Manager 3.5 Configuration

You can restore VxVM 3.5 configuration information from the backup directory /VXVM-CFG-BAK after VxVM 3.5 has been removed by `pkgrm`. For more information on this feature, please refer to the *VERITAS Volume Manager 3.5 Installation Guide*.

## Changed Features

Changed features in VxVM 3.5 include the following:

- u Support for `vxprint`

`vxprint` is no longer a `setuid-root` program, so `pkgadd` no longer displays a message about `setuid-root` programs being installed.

## VxVM 3.5 Licenses

The following table shows the levels of licensing in VERITAS Volume Manager 3.5 and the features supported at each level:

VxVM License	Description of Supported Features
Full	Concatenation, spanning, rootability, volume resizing, multiple disk groups, co-existence 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, volume snapshots, multipath DMP, and VEA.
Add-on Licenses	Features that augment the Full VxVM license such as FlashSnap™ (FastResync and Dynamic Disk Group Split and Join), and clustering functionality (cluster-shareable disk groups and shared volumes).



---

**Note** You need a Full VxVM license to make effective use of Add-on licenses to VxVM. For example, FastResync (a FlashSnap license feature) reduces the time taken to resynchronize volume snapshots (a Full VxVM license feature).

---

## Photon Licensing

VxVM automatically generates a photon license when it detects an A5k array. In addition, VxVM 3.5 allows you to remove your Photon device for servicing for up to two weeks. To do this, when the current license is due to expire, VxVM automatically removes the old license, and re-issues a new Photon license with another two-week expiration date.

You can now remove your Photon device for servicing, without affecting other operations.

If you perform a license query during this time, you may see information displayed as follows:

```
License Key =
IGUE-WR66-HB9S-V7G3-NE3W-8PPP-PNP8-PZO6-OP
Product Name = VERITAS Volume Manager
License Type = DEMO_NODE_LOCK
OEM ID = 4090
Demo End Date = Fri Aug 30 01:00:00 2002 <=====
EXPIRATION DATE
(13.6 days from now).
Node Lock Type = (Hostid and Architecture ID)
Features :=
PHOTON                      = Enabled
CPU Count                    = Not Restricted
```

---

**Note** This message does not mean that your Photon array feature license will soon expire for good. You will NOT lose those VxVM features associated with this license in 13.6 days time!

---



## Termination of Support Statements

### Termination of Support for VxVM Releases

The following software is no longer supported by VERITAS:

- u VERITAS Volume Manager Storage Administrator (VMSA)
- u VERITAS Volume Manager Release 1.3.x
- u VERITAS Visual Administrator Release 1.3.x
- u VERITAS Volume Manager Release 2.0.x
- u VERITAS Volume Manager Release 2.1.x
- u VERITAS Volume Manager Release 2.2.x
- u VERITAS Volume Manager Release 2.3.x
- u VERITAS Volume Manager Release 2.4.x
- u VERITAS Volume Manager Release 2.5.x
- u Solaris 2.3 operating system
- u Solaris 2.4 operating system
- u Solaris 2.5 and 2.5.1 operating system
- u VxVM no longer supports the Sun-4c product line:
  - SPARCstation 1
  - SPARCstation 1+
  - SPARCstation 2
  - SPARCstation IPC
  - SPARCstation IPX
  - SPARCstation SLC

### Termination of Sun Support for Solaris 2.6

Sun Microsystems has announced the End of Support Life for Solaris 2.6. Please contact Sun customer support for more information.

---

**Note** This does not mean that you cannot install VxVM 3.5 on Solaris 2.6.

---



## **Additional Changes in Hardware and Format Support**

- u This release of VxVM provides documentation only in PDF format. The documentation is no longer available in Postscript format.
- u The Sun SparcStorage Array (SSA) is no longer supported. Please disregard any references to this array in all VxVM documentation.



## Software Fixes and Enhancements in VxVM 3.5

The following problems are fixed in the 3.5 Release. (Lists of incidents fixed in previous releases of VxVM can be found in the applicable Release Notes.)

VERITAS Incident Number	Description of fixes
29959	<code>vxdmpadm enable/disable ctrlr</code> displays an error and exits when an invalid controller name is specified.
32559	The <code>vxr5check</code> utility does not work on degraded volumes. This is documented in the <code>vxr5check(1M)</code> manual page.
45668	Resize of a mirrored volume with plexes of different types is currently unsupported. This is documented in the <code>vxassist(1M)</code> and <code>vxresize(1M)</code> manual pages. (Sun Bug I.D. 4339626)
55735	Sun patch 110342-02 (and later) fixed the problem with multiple entries for the same AP metanode occasionally being visible to VxVM.
63854	<code>vxconfigd</code> no longer dumps core when it tries to import an unsupported DG version. Instead the following error is reported: <code>vxvm:vxdg: ERROR: Disk group cdg: import failed: Disk group has no valid configuration copies</code>
68454	When renaming enclosures, the enclosure name cannot contain more than 25 characters. The disk access name can be up to 31 characters, which includes 25 characters for the enclosure name plus 6 characters for an underbar and numbering ( <code>_nnnnn</code> ).
72148	The problem with the commands <code>vxddladm excludearray pid=&lt;PID&gt;</code> and <code>vxddladm excludearray vid=&lt;VID&gt;</code> excluding more than intended are fixed.
72734	The manual page for <code>vxddladm(1M)</code> and the help text from the command now match.
73503, 75214	While performing a <code>relayout</code> on a mirrored volume, the <code>vxassist</code> command no longer keeps the volume as mirrored even if the <code>layout</code> attribute is specified as <code>stripe</code> or <code>nomirror</code> .
73583	The installation script no longer asks you to install patch 105357.
75530	When using <code>vxdisksetup</code> with T3 LUNs, the command no longer fails with the error "Device does not match the kernel configuration." (Sun Bug I.D. 4491753)
76980	The problem that was reported with <code>vxdg free</code> not reporting the same amount of free space as <code>vxprint -ht</code> could not be reproduced.
80615	If the last active volume is disabled during a resync, I/O no longer continues to the stale plex and no longer corrupts the data.



## Software Limitations and Problems in VxVM 3.5

See the following sections for information about known problems and issues in this release of VxVM:

- u “Disk Controller Firmware Upgrades” on page 19
- u “T3B Firmware Upgrade on Solaris 9” on page 20
- u “Installation Issues” on page 22
- u “Upgrade Issues” on page 23
- u “Device Issues” on page 29
- u “Hot-Relocation Issues” on page 32
- u “DMP Issues” on page 33
- u “Cluster Functionality Issues” on page 34
- u “Data Integrity Issues” on page 35
- u “Miscellaneous Issues” on page 37
- u “Solaris Issues” on page 39

### Disk Controller Firmware Upgrades

The following procedure is a workaround for SUN Bug I.D. 4164338.

Obtain firmware upgrades as appropriate from your disk drive vendor. Download the appropriate files and documentation from the vendor’s support website.

To upgrade disk controller firmware, you do not need to reboot the system or unload the VxVM in-kernel drivers `vxdump`, `vxio` and `vxspec` (to guarantee data availability during the firmware upgrade procedure to a disk participating in a RAID mirror configuration).

With DMP enabled on the system with a volume mirrored across 2 controllers on one HBA, set up the configuration as follows:

1. Disable the plex associated with the disk device:

```
# /opt/VRTS/bin/vxplex -g diskgroup det plex
```

2. Stop I/O to all disks through one controller of the HBA by executing the following command:

```
# /opt/VRTS/bin/vxdmpadm disable ctlr=first_cntlr
```

For the other controller on the HBA, enter:

```
# /opt/VRTS/bin/vxdmpadm -f disable ctlr=second_cntlr
```



3. Upgrade the firmware on those disks for which the controllers have been disabled using the procedures that you obtained from the disk drive vendor.

4. After doing the upgrade, enable all the controllers by executing:

```
# /opt/VRTS/bin/vxdmpadm enable ctlr=second_cntlr
```

5. Enable the plex associated with the device:

```
# /opt/VRTS/bin/vxplex -g diskgroup att volume plex
```

This command takes some time depending upon the size of the mirror set.

6. Upgrade the disk group:

```
# vxdg -T version upgrade diskgroup
```

This command takes some time depending upon the size of the mirror set.

## T3B Firmware Upgrade on Solaris 9

On Solaris 9 only, a T3B upgrade to firmware version 2.1 must follow the procedure below. Not using the procedure leads to disabled disk groups or an inability to mount file systems. The procedure is a result of VERITAS incident number 95877.

1. Use the `umount` command to unmount related filesystems

2. Stop all VxVM volumes:

```
# vxvol stop vol_name
```

3. Stop VxVM:

```
# vxdctl stop  
# vxiod -f set 0
```

4. Upgrade T3B firmware to version 2.1.

5. Start VxVM:

6. # vxiod set 10  
# vxconfigd -m disable  
# vxdctl enable

7. Start the VxVM volumes:

```
# vxvol -g dg_name start vol_name
```





8. Use the `mount` command to remount the file system.



## Installation Issues

VERITAS Incident Number	Description
88367	The <code>pkgadd</code> command is not finding patches correctly. To check what patches are installed on your system, run <code>showrev -p</code> . If you have the correct patches installed, you can ignore any messages about missing patches that you may see during installation.
102081	Using <code>vxunroot</code> to unencapsulate and remove the rootdisk from VxVM will currently fail if the rootdisk is the only disk in <code>rootdg</code> . The rootdisk is not removed from the configuration, and re-encapsulation of the rootdisk fails.  Workaround (1): Remove the <code>VXVM-CFG-BAK</code> directory before installing the VxVM package. This allows the disk to be re-encapsulated and the file <code>/etc/vx/reconfig.d/state.d/install-db</code> to be created. The disk can then be re-encapsulated during the <code>vxinstall</code> .  Workaround (2): Run the following commands to re-encapsulate the rootdisk: <pre># vxdctl disable # vxencap rootdisk device_name</pre> (You can ignore any error messages displayed by <code>vxdmpadm</code> .)
104091	If you exclude all disks (using <code>/etc/vx/disks.exclude</code> ), then run <code>vxinstall</code> using the quick install option, nothing changes. However, at the end of the install process, <code>vxinstall</code> asks you if you want to reboot. The reboot option is not needed, and should be ignored.  <b>Workaround:</b> Reply <b>no</b> to the reboot option, and <code>vxinstall</code> will exit gracefully.
none	Remove a VM disk from a disk group after evacuating any data on the disk. You can permanently remove the disk from VxVM control by removing the VM metadata partition from the disk. Use the low-level command, <code>vxdiskunsetup</code> , as follows: <pre># /usr/lib/vxvm/bin/vxdiskunsetup c#t#d#</pre> <b>Note</b> This command permanently removes a disk from VxVM control. It should be used with caution only by a system administrator who is trained and knowledgeable about VxVM.



VERITAS Incident Number	Description
none	If you are planning to initialize disks, check to see if any of the disks were previously under VxVM control. If so, and if they were used on the same host system, the disk groups they represent are imported automatically during the installation process if the proper deinstallation procedures were <i>not</i> followed. An attempt during installation to initialize or encapsulate disks that were previously under VxVM control fails. After installation, if you no longer want to use those disk groups, use the destroy option of the <code>vxchg (1M)</code> command to remove those disk groups. Alternately, you can use <code>vxdiskunsetup (1M)</code> to remove the disks from VxVM control. Be aware that these options can result in data loss if used incorrectly.
none	There is an error in the instructions for Loading the Software from CD-ROM in the <i>VERITAS Volume Manager 3.5 Installation Guide</i> . In the sub-section "Removing Previous Licensing Packages," the command to de-install the old <code>VRTSvlic</code> package, in Step 1, should read:  # <code>pkgrm VRTSvlic</code> (not, # <code>pkgrm VRTSlic</code> )

## Upgrade Issues

VERITAS Incident Number	Description
94896	You may find that the <code>upgrade_finish</code> command fails with the message: "can't unload the module: Device busy".  Workaround: Bring down the <code>vxconfigd</code> daemon and re-run <code>upgrade_finish</code> . For example:  # <code>kill -9 &lt;pid of vxconfigd&gt;</code> # <code>upgrade_finish</code>
102047	After upgrading your system to Volume Manager 3.5 and starting the reboot, you may see messages of the form:  <code>vxvm: Volume vol1 for mountpoint /vol1 is not in rootdg disk group; Cannot start the volume early</code> <code>vxvm: Volume vol2 for mountpoint /vol2 is not in rootdg disk group; Cannot start the volume early</code> <code>VxVM starting special volumes ( swapvol rootvol usr var tmp )</code> <code>...</code>  These messages are informational only, and can safely be ignored.



<b>VERITAS Incident Number</b>	<b>Description</b>
102638	<p>You must have VRTSlic version 1.0.1 installed before you can upgrade to VxVM 3.5. That is, upgrade is not supported with VRTSlic version 1.0.</p> <p>You can use the <code>pkginfo</code> command to determine the current version of VRTSlic:</p> <pre># pkginfo -l VRTSlic PKGINST: VRTSlic NAME: VERITAS Licensing Utilities CATEGORY: system,utilities ARCH: sparc VERSION: 1.0.1,REV=08.15.2001.22.01       (VERSION needs to be "1.0.1" or higher) BASEDIR: / VENDOR: VERITAS Software DESC: Licensing commands PSTAMP: VERITAS-LIC1.0-2049.0 INSTDATE: Oct 15 2001 18:08 HOTLINE: (800) 342-0652 EMAIL: support@veritas.com STATUS: completely installed FILES: 23 installed pathnames 13 shared pathnames 15 directories 2 executables 756 blocks used (approx)</pre>



VERITAS Incident Number	Description
103454	<p>There are three situations in which an attempt to replace a failed disk may be unsuccessful:</p> <ol style="list-style-type: none"> <li>1. If you attempt to use <code>vxdiskadm option 5</code> to replace a failed disk that was initialized using a pre-3.5 version of Volume Manager, then the replacement may fail if the disk was full and the new replacement disk has been initialized using Volume Manager version 3.5. For more information and workarounds, please refer to Situation 1, “Replacing Failed Disks” on page 41.</li> <li>2. If you attempt to use <code>vxdbg -k adddisk</code> to replace a failed disk, you may encounter a problem if any of the following hold true: <ul style="list-style-type: none"> <li>u The failed disk was initialized using a pre-3.5 version of Volume Manager.</li> <li>u The replacement disk was initialized in the same manner as the failed disk, but was initialized using version 3.5 of Volume Manager</li> <li>u The amount of space used by subdisks on the failed disk exceeds the size of the public region of the replacement disk.</li> </ul> <p>For more information and workarounds, please refer to Situation 2, “Replacing Failed Disks” on page 41.</p> </li> <li>3. When replacing a failed disk using the VEA GUI, the following error message may appear: <pre>vxvm:vxdbg: ERROR: associating disk-media disk01 with clt10d0s2: Disk public region is too small</pre> <p>For more information and workarounds, please refer to Situation 3 in “Replacing Failed Disks” on page 41.</p> </li> </ol>



VERITAS Incident Number	Description
103919	<p>Under certain conditions, <code>upgrade_start</code> may fail to find all volumes in rootdisk. It will then quit prematurely and output the following error message:</p> <pre>Unexpected inconsistency in &lt;filename&gt;</pre> <p>The reason for this is that <code>upgrade_start</code> uses <code>/etc/vx/reconfig.d/disk.d/rootdisk&gt;vtoc</code> to detect all the volumes in the rootdisk. This error will occur if the Volume Table Of Contents (VTOC) of the rootdisk is different from the <code>/etc/vx/reconfig.d/disk.d/rootdisk&gt;vtoc</code>.</p> <p>VTOC mismatch is most commonly introduced when the encapsulated rootdisk fails and is then replaced with a new disk and is then re-sync'd from the mirror disk (commonly done using <code>vxdiskadm</code> option #4), followed by replacing the rootdisk (using <code>vxdiskadm</code> option #5). This operation may cause the VTOC of the new rootdisk to be different from the original VTOC of the rootdisk.</p> <p><b>Workaround:</b></p> <p>Move the <code>/etc/vx/reconfig.d/disk.d/rootdisk&gt;vtoc</code> file to a different place before using <code>upgrade_start</code>. This enables <code>upgrade_start</code> to continue, but it will only un-encapsulate file system volumes <code>/</code>, <code>swap</code>, <code>/opt</code>, <code>/usr</code> and <code>/var</code>. All other volumes (if they exist) will remain encapsulated and will be inaccessible between <code>upgrade_start</code> and <code>upgrade_finish</code>. If you wish to access data from these other volumes, then you need to manually unencapsulate the volumes before running <code>upgrade_start</code>.</p>
none	<p>If a swap volume specified in <code>/etc/vfstab</code> is mirrored at the time that <code>upgrade_start</code> is run, the <code>upgrade_finish</code> script starts a resynchronization of the volume. This can cause a message similar to the following to be printed when the command to reboot the system is issued:</p> <pre>xvm:vxvol: tutil0 field for plex plex_name changed unexpectedly</pre> <p>This message can be ignored.</p>



VERITAS Incident Number	Description
none	<p>For a system on which the root file system is contained on a mirrored volume, the <code>upgrade_start</code> script can choose a mirror on a disk other than the normal boot disk to perform the upgrade. If this occurs, the reboot after running <code>upgrade_finish</code> can initially fail, claiming that the mirror on the boot disk is stale, as follows:</p> <pre>vxvm:vxconfigd: Error: System boot disk does not have a valid rootvol plex.Please boot from one of the following disks:  disk: *diskname*Device: *device* ... vxvm:vxconfigd:Error: System startup failed The system is down.</pre> <p>Boot the system from one of the disks named. If the <code>eeeprom</code> option <code>use-nvramrc?</code> is set to <code>true</code>, boot the system by specifying <code>vx-diskname</code>.</p>

## Utility Issues

VERITAS Incident Number	Description
6154	<p><code>vxassist</code> 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.</p>
11286	<p>Using <code>vxdg free</code> with a non-existent <code>disk-media-name</code> does not print an appropriate error message; it prints only a header.</p>
13418	<p>The <code>vxdisksetup</code> utility allows the administrator to specify some region/partition configurations that are invalid. Specifically, overlapping private and public regions can be specified, but can cause failures or data corruption when the disk is actually used. The administrator must check that the partitioning of the disk does not cause overlapping public and private partitions when the default partitioning is overridden from the command line.</p>



VERITAS Incident Number	Description
13488	<p>The <code>vxassist</code> command does not add a mirror and a log when processing a command such as the following:</p> <pre># vxassist mirror volume layout=log ...</pre> <p>The mirror is added, but the log is silently omitted. To add a log and a mirror, add the mirror and the log in two separate <code>vxassist</code> invocations, as follows:</p> <pre># vxassist mirror volume ... # vxassist addlog volume ...</pre>
56184	<p>The <code>vxassist</code> <code>relayout</code> process can change the volume layout to a layered volume. For example, you might end up with a stripe-mirror volume instead of a mirror-stripe volume.</p> <p>If the final layout is not what you intended, then there are two solutions:</p> <ol style="list-style-type: none"> <li>1. If the task is not complete, stop the relayout and reverse it by: <pre>/usr/sbin/vxrelayout -g dg_name -t task_id reverse vol_name</pre> </li> <li>2. If the task has completed, you can relayout the volume back to the original layout.</li> </ol>
57225	<p>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 unless you specify the attribute <code>layout=nodiskalign</code> to <code>vxassist</code>, as shown here:</p> <pre># vxassist growby volume length layout=nodiskalign</pre> <p><b>Note</b> Specifying <code>layout=nodiskalign</code> permanently enforces this layout policy on the volume.</p> <p>[Sun Bug I.D. 4554411]</p>
93840	<p><b>Stopping and Starting Layered Volumes</b></p> <p>If a layered volume is in use when a <code>vxvol stopall</code> command is issued, then only the sub-volumes are disabled. The layered volume remains enabled.</p> <p>When a <code>vxvol stop layered-volume</code> command is issued, then only the top layered volume is stopped. The sub-volumes remain enabled.</p> <p>When a <code>vxvol start layered-volume</code> command is issued, then only the top layered volume is started. The sub-volumes remain disabled.</p>
102381	<p><code>vxdiskadm</code> requires two attempts to replace a failed disk.</p> <p>The first attempt can fail with a message of the form:</p> <pre>/usr/lib/vxvm/voladm.d/bin/disk.repl: test: argument expected</pre> <p>The command is not completed and the disk is not replaced.</p> <p>If you now rerun the command, using Option 5, the replacement successfully completes.</p>





VERITAS Incident Number	Description
104205	<p>When resizing concat-pro and stripe-pro volumes, you may see an error message of the form:</p> <pre>vxvm:vxresize: ERROR: vxvm:vxplex: WARNING: Volume vol2-L04: Dirty region logging will not be enabled vxvm:vxresize: ERROR: until the volume is fully resync'ed.</pre> <p>The error message can be safely ignored. The volumes are correctly resized. (This error message can also appear when using VEA to resize the volumes.)</p>
none	<p>Ensure that an entry for <code>/usr/sbin</code> precedes <code>/sbin</code> in the setting of the <code>PATH</code> environment variable. This ensures that the correct <code>vxchg</code> binary is run if the full pathname is not specified. (Two versions of the <code>vxchg</code> binary exist. The version of <code>vxchg</code> in <code>/sbin</code> is only used at system startup. Administrators should use the version in <code>/usr/sbin</code>.)</p>
none	<p>Due to the current implementation of a resize of layered volumes, it is recommended that you do not <code>grow</code> or <code>shrink</code> layered volumes (stripe-mirror, concat-mirror, and so on) while resynchronization is ongoing.</p> <p>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.</p> <p>If the system reboots before the <code>grow</code> or <code>shrink</code> of a layered volume completes, the volume is left with an intermediate layout. In this case, you have to use <code>vxassist convert</code> to restore the volume to its original layout.</p> <p>After a layered volume is resized, the volume names, the plex names and the subdisk names associated with the subvolumes, are changed.</p>
none	<p>The <code>vxprint</code> command no longer has <code>setuid</code> to <code>root</code> permission when installed.</p>

## Device Issues

VERITAS Incident Number	Description
5316	<p>VxVM tracks disks using long unique identifiers that it stores on each disk. VxVM expects each disk to have a unique identifier, and does not effectively guard against the situation where two disks have the same identifier. Duplicate identifiers should only occur as a result of the administrator using <code>dd</code> or some other utility to perform physical copies of the contents of an entire disk.</p> <p><b>Workaround:</b> Avoid performing exact physical disk copying.</p>



<b>VERITAS Incident Number</b>	<b>Description</b>
8818	It is possible to prevent any access of a disk by VxVM. For example, startup of VxVM can be severely impacted by a disk with errors that result in I/O operations that take a long time to fail. However, when VxVM starts up, it accesses every disk on the system by reading its VTOC and possibly a few blocks from one partition. There is currently no mechanism to prevent this. A disk can be offlined persistently, but the offline state is only recognized <i>after</i> the probe of all disks.
68530	<p>If the system has a disk array connected to a switch in the name server zone in fabric mode, you may see some or all of the following messages during boot-up:</p> <pre>vxvm:vxconfigd: WARNING: Disk xxxxx in group yyy: disk device not found vxvm:vxconfigd: ERROR: enable failed: Error in disk group configuration copies. No valid disk found containing disk group; transactions are disabled. vxvm:vxconfigd: FATAL ERROR: Rootdg cannot be imported during boot.</pre> <p>The messages occur early in the boot sequence. As boot-up progresses, VxVM finds the items if they are there. You can ignore the messages. [Sun Bug I.D. 4462372]</p>
73154	When Hitachi DF400 and DF500 arrays are configured as Active/Active mode, performance is degraded.



VERITAS Incident Number	Description
103454	<p>There are three situations in which an attempt to replace a failed disk may be unsuccessful:</p> <ol style="list-style-type: none"> <li>1. If you attempt to use <code>vxdiskadm option 5</code> to replace a failed disk that was initialized using a pre-3.5 version of Volume Manager, then the replacement may fail if the disk was full and the new replacement disk has been initialized using Volume Manager version 3.5. For more information and workarounds, please refer to Situation 1, “Replacing Failed Disks” on page 41.</li> <li>2. If you attempt to use <code>vxdbg -k adddisk</code> to replace a failed disk, you may encounter a problem if any of the following hold true: <ul style="list-style-type: none"> <li>u The failed disk was initialized using a pre-3.5 version of Volume Manager.</li> <li>u The replacement disk was initialized in the same manner as the failed disk, but was initialized using version 3.5 of Volume Manager</li> <li>u The amount of space used by subdisks on the failed disk exceeds the size of the public region of the replacement disk.</li> </ul> <p>For more information and workarounds, please refer to Situation 2, “Replacing Failed Disks” on page 41.</p> </li> <li>3. When replacing a failed disk using the VEA GUI, the following error message may appear: <pre>vxvm:vxdbg: ERROR: associating disk-media disk01 with clt10d0s2: Disk public region is too small</pre> <p>For more information and workarounds, please refer to Situation 3 in “Replacing Failed Disks” on page 41.</p> </li> </ol>
103991	<p>Do not run the <code>vxrelayout</code> and <code>vxassist</code> commands to relayout a volume that is part of root disk. This action may ruin the layout of the root disk; so that you will not be able to boot from the disk. On an encapsulated root disk, a relayout would cause the upgrade to fail.</p>
none	<p>Disks with insufficient space (less than 2048 disk blocks) for the allocation of an on-disk database copy cannot be encapsulated. The database requires at least the same space as is allocated for other disks in the same disk group. The default size is 2048 blocks. To work around this, relocate the data on the last partition of the disk to a volume on a different disk, and free the space by reducing the partition size to 0.</p> <p>The space for the database must be allocated from the beginning or the end of the disk, with the exception of the root disk. The root disk can be encapsulated by carving out space from the <code>swap</code> partition if there is no space at the beginning or at the end of the disk. This is done by creating a subdisk for the private partition in the space obtained from the <code>swap</code> partition.</p> <p><b>Workaround:</b> The problem of insufficient space on a disk to store private VxVM information has no workaround. VxVM requires at least a small region of private storage (2048 blocks) for proper disk identification.</p>



VERITAS Incident Number	Description
none	Booting off a fabric device is not supported. Also, encapsulating a fabric boot device is not supported, nor is mirroring a locally encapsulated boot device to a fabric device.
105173	On T3 array, Volume Manager may get the following failure when trying to add a disk (typically from vxinstall or vxdisksetup): <pre>vxvm:vxdisk: ERROR: Device XXXX: online failed Device path not valid</pre> This can happen in cases where the T3 disk was re-partitioned (or re-formatted) prior to disk(s) being added.

## Hot-Relocation Issues

VERITAS Incident Number	Description
14894	Except for <code>rootvol</code> and <code>swapvol</code> , 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).
14895	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.



## DMP Issues

VERITAS Incident Number	Description
102097	<p>When VxVM detects disks with duplicate disk IDs (unique internal identification), VxVM attempts to select the appropriate disk (through array vendor-specific logic). If a disk can not be selected, VxVM does not import any of the duplicated disks into a disk group. In the rare case when VxVM cannot make the choice, you must choose which duplicate disk to use.</p> <p>In previous releases, VxVM would select the first disk found if the selection process failed. The change in this release is to avoid the selection of the wrong disk, which could lead to data corruption. An array with hardware mirroring capability is particularly susceptible to such data corruption.</p> <p>The following examples describe where user intervention is needed.</p> <p>Example 1:</p> <p>When DMP is disabled to an array that has multiple paths, then each path to the array is claimed as a unique disk. VxVM detects duplicate disks whenever an array is configured as a unique disk for each accessible path and gives the following message:</p> <pre>vxvm:vxconfigd: NOTICE: Unable to resolve duplicate diskid.</pre> <p>When DMP is suppressed, VxVM can not know which path to select as the true path. You must decide which path to use. Decide which path to exclude and then either edit the file <code>/etc/vx/vxvm.exclude</code> or, if <code>vxconfigd</code> is running, use the <code>vxdiskadm</code> option 17 selection 1 (suppress all paths through a controller from VxVM's view) or selection 2 (suppress a path from VxVM's view).</p> <p>The following example shows a <code>vxvm.exclude</code> file with paths <code>c6t0d0s2</code>, <code>c6t0d1s2</code>, and <code>c6t0d2s2</code> excluded from VxVM:</p> <pre>exclude_all 0 paths c6t0d0s2 /pci@1f,4000/SUNW,ifp@2/ssd@w50060e8003275705,0 c6t0d1s2 /pci@1f,4000/SUNW,ifp@2/ssd@w50060e8003275705,1 c6t0d2s2 /pci@1f,4000/SUNW,ifp@2/ssd@w50060e8003275705,2 # controllers # product # pathgroups</pre>



VERITAS Incident Number	Description
	<p>Example 2:</p> <p>Some arrays such as EMC, HDS, and so on provide hardware mirroring. When a lun pair is split, depending on how the process is performed, you may get two disks with the same diskid. With duplicate disk IDs, VxVM gives the following message:</p> <pre>vxvm:vxconfigd: NOTICE: Unable to resolve duplicate diskid.</pre> <p>Check with your array vendor to make sure that your site uses the proper split procedure. If you know which luns to use, decide which path to exclude and then either edit the file <code>/etc/vx/vxvm.exclude</code> or, if <code>vxconfigd</code> is running, use the <code>vxdiskadm</code> option 17 selection 1 (suppress all paths through a controller from VxVM's view) or selection 2 (suppress a path from VxVM's view).</p> <p>Example 3:</p> <p>When you have disks duplicated using <code>dd</code> or any other disk copying utility, VxVM gives the following message:</p> <pre>vxvm:vxconfigd: NOTICE: Unable to resolve duplicate diskid.</pre> <p>Choose which set of duplicated disks that you want to exclude, and either edit the file <code>/etc/vx/vxvm.exclude</code> or, if <code>vxconfigd</code> is running, use the <code>vxdiskadm</code> option 17 selection 1 (suppress all paths through a controller from VxVM's view) or selection 2 (suppress a path from VxVM's view).</p>
none	Messages from the <code>Format</code> command can be ignored while you are running DMP. Refer to the <i>VERITAS Volume Manager Administrator's Guide</i> .
none	DMP is a required module in releases of VxVM 3.1.1 and later.
none	You can no longer use the steps to disable/re-enable DMP that were supported in VxVM 3.1 and earlier.

## Cluster Functionality Issues

VERITAS Incident Number	Description
20448	<p>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</p> <pre># vxvol -f resync volume</pre> <p>[Sun Bug I.D. 4087612]</p>



VERITAS Incident Number	Description
none	VxVM does not currently support RAID-5 volumes in cluster-shareable disk groups. Creating and using RAID-5 volumes on shared disks may cause a system panic.
none	The use of file systems other than VERITAS Cluster File System™ (CFS) on volumes in cluster-shareable disk groups can cause system deadlocks.
none	If the <code>vxconfigd</code> 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 <code>vxconfigd</code> 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 <code>vxconfigd</code> program must therefore be started on the master first.
none	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: <pre>cannot assign minor #</pre> This message is accompanied by the console message: <pre>WARNING:minor number ### disk group <i>group</i> in use</pre>

## Data Integrity Issues

VERITAS Incident Number	Description
88090	To avoid data corruption, detaching a plex from a volume whose data integrity is being recovered at system startup has the following effects: <ul style="list-style-type: none"> <li>u Snapback of any existing snapshot of the volume requires a full resynchronization.</li> <li>u Reattachment of any existing plex that is currently detached from the volume requires a full resynchronization.</li> </ul>
none	Disk drive 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

VERITAS Incident Number	Description
14909	If a disk fails after a snapshot is complete, the snapshot plex does not detect the failure or detach from the volume. This is because a snapshot plex is a write-only plex, so it notices I/O errors, but does not detach.
87206	<p>When you use the following command:</p> <pre># vxassist -o resyncfromoriginal snapback snapvo</pre> <p>to execute a default snapback operation (resynchronizing from the original volume), and the system crashes while the operation is still in progress, you will find, when the system comes back up and the volumes are restarted, that the snapshot plexes are not associated with any volume.</p> <p>To recover from this situation, you should use the following procedure</p> <ol style="list-style-type: none"> <li>1. Use the following commands to discover the original volume name of each snapshot plex: <pre># volrid=`vxprint -g diskgroup -p -F "%snap_riid" \ plexname`</pre> <pre># volprint -g diskgroup -n -v -e v_riid=\$volrid</pre> </li> <li>2. Using the information discovered in step 1, identify the volumes to which the snapshot plexes originally belonged, and reattach them to those original volume as in the following command: <pre># volplex att original_volume plex1 [plex2 ...]</pre> </li> </ol>
None	<p>It is recommended that you do not use Persistent FastResync or snapshot disks as root disks.</p> <p>(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.)</p>
None	<p>When taking a snapshot of a cluster mounted vxfs filesystem, the following warning message might appear:</p> <pre>vxio: WARNING: vxvm:vxio: Plex plex&gt; detached from volume vol</pre> <p><b>Workaround:</b> None. This is normal and is <i>NOT</i> the result of any error conditions.</p>
None	<p>Normally, a filesystem would have no work to do. However, if none of the filesystems for a cluster mounted vxfs filesystem are mounted, then it is likely that the <code>fsck</code> of the snapshot would take longer than is necessary, depending on the I/O activity at the time of the snapshot.</p> <p><b>Workaround:</b> When taking a snapshot of a cluster mounted vxfs filesystem, you should ensure that at least one of the volumes defined in the command line is mounted on the CVM master.</p>





## Miscellaneous Issues

VERITAS Incident Number	Description
9936	RAID-5 volumes cannot be mirrored.
13741	<p>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.</p> <p>If it is important that a disk group <i>not</i> be auto-imported when the system is rebooted, the disk group should be imported temporarily when the intention is to deport the diskgroup (for example, in HA configurations). Use the <code>-t</code> flag to <code>vxvg import</code>.</p>
14450	<p>During very fast boots on a system with many volumes, <code>vxconfigd</code> may not be able to auto-import all of the disk groups by the time <code>vxrecover -s</code> is run to start the volumes. As a result, some volumes may not be started when an application starts after reboot.</p> <p><b>Workaround:</b> Check the state of the volumes before starting the application, or place a sleep (<code>sleep sec</code>) before the last invocation of <code>vxrecover</code>.</p>
14915	<p>The <code>vxrecover</code> 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:</p> <pre># vxvol -f start volume</pre> <p>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.</p>
25644	<p>The way that VxVM handles minor numbers for volume devices can cause problems when upgrading from VERITAS Volume Manager 2.x to 3.x. The problems are unneeded disk group remappings and the <code>vxconfigd</code> command can hang. The problem does not affect disk group versioning.</p>



VERITAS Incident Number	Description
96801	<p>The following informational message is currently described in the <i>Volume Manager 3.5 Troubleshooting Guide</i>.</p> <pre>vxvm:vxconfigd: WARNING: library and vxconfigd disagree on existence of client number</pre> <p>However, the message may be displayed in the form:</p> <pre>vxvm:vxconfigd::library and vxconfigd disagree on existence of client 9188</pre> <p>(that is, the word: "WARNING" (or "INFO") may not appear in the actual message.)</p> <p>You can safely ignore this message.</p>
101652	<p>The <code>vxlicrep</code> command output may currently show an installed license for the "Storage Expert" feature. However, this (separately licensable) feature is not supported in VxVM 3.5 release (it will be supported in future VxVM releases).</p>
102598	<p>If you are using VERITAS Volume Replicator (VVR) 3.5, and plan to share disk groups, you should refer to incident 102598 in the <i>VERITAS Volume Replicator Release Notes</i>.</p>
none	<p>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.</p>
none	<p>The Sun Online:Backup™ facility does not accept the long device path names for volumes. A limitation of Online: Backup is that it does not accept device paths longer than 24 characters.</p> <p><b>Workaround:</b> Use symbolic links to the longer <code>/dev/vx/dsk/volname</code> paths from a shorter pathname.</p>
none	<p>The following messages may get displayed on the console during a system reboot or during VxVM initialization when you are running <code>vxinstall</code>:</p> <pre>No VVR license installed on the system; vradmind not started No VVR license installed on the system; in.vxrsyncd not started</pre> <p>These messages are informational only, and can be safely ignored if you are not a VERITAS Volume Replicator (VVR) user.</p>



## Solaris Issues

VERITAS Incident Number	Description
6211	Driver close calls should take priority on SVR4-derived operating systems. If this is not done, it is possible to receive a simultaneous last-close and first-open operation, swapping the ordering of the arrival of operations to VxVM, and as a result, leave the volume device closed. (This can happen if VxVM detects an additional open of the device followed a last close).
8948	A Solaris SCSI disk driver offlines a disk when the disk is not available and notifies the console. When the disk is connected back to the system, the disk is not automatically opened. All processes in which the disk must be open cannot read or write from or to the disk until the disk is opened by some other utility, causing the disk to come online. To address this situation, issue a command that opens a partition of the disk. For example:  # : < /dev/rdsk/c1t5d0s2



VERITAS Incident Number	Description
13312	<p>The versions of the kernel drivers for VxVM are incompatible with some versions of the Solaris operating system. Multiple kernel modules are installed and properly maintained by the installation and upgrade software. It is possible for a mismatch to occur (for example, if the administrator moves the kernel driver files). If a mismatch occurs, the VxVM kernel prints a warning message on the console similar to the following message:</p> <pre>WARNING: vxio: incompatible kernel version (5.X), expecting 5.X</pre> <p>If this message is displayed, the system must be booted for recovery (as explained in the <i>VERITAS Volume Manager Troubleshooting Guide</i>) and the correct kernel modules installed. To install the correct kernel module versions, <code>cd</code> to the <code>kernel/drv</code> directory of the mounted root file system. To list the VxVM kernel modules, use the following command:</p> <pre># ls -l vxio* vxspec* vxdmp*</pre> <p>The release-specific versions of the kernel modules are stored as <i>module.OS_release</i>, where <i>OS</i> and <i>release</i> are the result of running the <code>uname -s</code> and <code>uname -r</code> commands on the system, respectively.</p> <p>For example, on a misconfigured system running Solaris 2.6, the listing for <code>vxio*</code> may be similar to the following:</p> <pre>-rw-r--r-- 1 root other 1682424 ... vxio -rw-r--r-- 1 root sys 1647664 ... vxio.SunOS_5.6 -rw-r--r-- 1 root sys 1661340 ... vxio.SunOS_5.7 -rw-r--r-- 1 root sys 1682424 ... vxio.SunOS_5.8</pre> <p>The size of the <code>vxio</code> kernel module that is in use matches the <code>vxio.SunOS_5.8</code> version. To correct the problem, copy the <code>SunOS_5.6</code> versions to the in-use module name:</p> <pre># cp vxio.SunOS_5.6 vxio</pre> <p>Finally reboot the system.</p>
13388	<p>During encapsulation, VxVM does not consider a partition to be a swap partition unless its partition tag (as shown by <code>prtvtoc</code>) is <code>swap</code> or <code>3</code>. Any partition used as a swap partition but not tagged as such is encapsulated as a file system. In the <code>vfstab</code> file, a note is made that the partition has been encapsulated, but the <code>vfstab</code> entry <i>is not</i> translated, and thus, the partition is not added as a swap area as part of the boot process.</p> <p>All partitions used as swap must be marked with the <code>swap</code> tag to be properly encapsulated.</p>
none	<p>Since the disk label is stored in block 0 of the disk, block 0 must not be used (that is, no application should write any information in block 0). Special protection has been built into VxVM to protect block 0 from being overwritten.</p>



VERITAS Incident Number	Description
none	The <code>dd</code> command in Solaris uses only <code>lseek()</code> to seek to a particular offset in a file. It does not use <code>llseek()</code> . This causes <code>dd</code> to fail on volumes greater than 2 gigabytes.
none	On Solaris, slice 2 of a disk is the full disk by default. When finding connected disks, VxVM checks slice 2 of a disk. Slice 2 on a disk must always be defined as the full disk slice with a tag of 0x05.
none	If the PROM version is not at least version 2, the system is not suitable for the configuration of a bootable root volume; the following message appears on the console each time the machine is booted after adding the VxVM package:  <pre>VxVM: Root volumes are not supported on your PROM version.</pre> Any attempt to encapsulate the root disk fails on these machines.
none	If multiple swap partitions are encapsulated on your disks, VxVM names them as <code>swapvol</code> , <code>swapvol1</code> , <code>swapvol2</code> , and so on. When the system is rebooted, the following error message is displayed:  <pre>/dev/vx/dsk/swapvol2 : Overlapping swap files are not allowed</pre> However, the swap devices are correctly added with no ill effects on the system. To avoid seeing this message, shorten the names of swap volumes (other than <code>swapvol</code> ) from <code>swapvoln</code> to <code>swapn</code> .

## Replacing Failed Disks

---

**Note** The failed disk and the replacement disk should have the same disk geometry; that is, they should have the same bytes per sector, sectors per track, tracks per cylinder and sectors per cylinder, same number of cylinders, and the same number of accessible cylinders. (You can use the `prtvtoc` command to get disk information.)

---

There are three situations in which an attempt to replace a failed disk may be unsuccessful. These situations, and the appropriate workarounds, are described below.

---

**Note** Any of these situations could happen if the disk geometry is different between the failed and replacement disks.

---

### Situation 1

If you attempt to use `vxdiskadm` option 5 to replace a failed disk that was initialized using a pre-3.5 version of Volume Manager, then the replacement may fail if the disk was full and the new replacement disk has been initialized using Volume Manager version 3.5.



Two possible scenarios are:

**Scenario A:**

The disk to be used as the replacement does not show up in the list of available replacements when using `vxdiskadm` option 5, even though the disk has already been initialized.

**Scenario B:**

The disk to be used as the replacement does not show up in the list of available replacements when using `vxdiskadm` option 5, because it has not been initialized.

Two workarounds are available for replacing disks previously initialized using VxVM 3.2 or earlier. You should first use workaround 1, then, if that fails, use workaround 2.

**Workaround 1:**

- u Redo `vxdiskadm` option 5.
- u When selecting disk for replacement, choose “none.”
- u When prompted about whether to “Use a default private region length for the disk?”, respond “y.”

For example:

```
.... The following devices are available as replacements:
c0t1d0s2 c1t9d0s2 c1t10d0s2
You can choose one of these disks to replace disk09. Choose “none”
to initialize another disk to replace disk09.
Choose a device, or select “none” [<device>,none,q,?] (default:
c0t1d0s2) none
```

---

**Note** You *must always* select the default private region length.

---

```
Use a default private region length for the disk?
[y,n,q,?] (default: y) y
```

---

**Note** `vxdiskadm` does not always provide a complete list of disk replacements. In the above example, “c0t1d0s2 c1t9d0s2 c1t10d0s2” may not be a complete list.

---

**Workaround 2:**

- u Exit out of `vxdiskadm`.
- u Use `vxdbg -g diskgroup rmdisk medianame` to remove the disk from the disk group.
- u Use the `vxdisksetup` program to initialize the target disk:

```
# /usr/lib/vxvm/bin/vxdisksetup -i c##t##d## privlen=XXXX
old_layout
```

where XXXX is the size of the private region of the source disk.

If you don't know the previous configuration of your disk such as the private region length, you can use the default private region length:

```
# /usr/lib/vxvm/bin/vxdisksetup -i c#t#d# old_layout
```

- u Add the newly initialized target disk to the source disk group:

```
# vxdg -g diskgroup adddisk medianame=c#t#d#
```

- u Go back into vxdiskadm option 5 again to replace your failed disk.

## Situation 2

If you attempt to use `vxdg -k adddisk` to replace a failed disk, you may encounter a problem if any of the following hold true:

- u The failed disk was initialized using a pre-3.5 version of Volume Manager.
- u The replacement disk was initialized in the same manner as the failed disk, but was initialized using version 3.5 of Volume Manager
- u The amount of space used by subdisks on the failed disk exceeds the size of the public region of the replacement disk.

For example:

```
vxdg -k adddisk disk01=clt10d0s2
vxvm:vxdg: ERROR: associating disk-media disk01 with clt10d0s2:
Disk public region is too small
```

**Workaround:** Please refer to the Workarounds in Situation 1.

## Situation 3

When using the VEA GUI to replace a failed disk, if you see the following error message:

```
vxvm:vxdg: ERROR: associating disk-media disk01 with clt10d0s2:
Disk public region is too small
```

VEA GUI is invoking the `vxdg -k adddisk` command to replace a failed disk.

**Workaround:** There is no workaround using the VEA GUI. Please refer to the Workarounds in Situation 1.

# VEA

Refer to the *VERITAS Volume Manager Installation Guide* for information on how to set up and start the VEA server and client.



## VEA Issues

VERITAS Incident Number	Description
82995	VEA may core dump when closing the Volume View display window when using the CDE desktop. If this occurs, restart the VEA application.
78319	If, after installing and starting VEA, you obtain a new license, you will not have access to the newly licensed features until you restart the VEA service: <ol style="list-style-type: none"> <li>1. Stop the VEA backend service: # /opt/VRTSob/bin/vxsvc -k</li> <li>2. Start the VEA backend service: # /opt/VRTSob/bin/vxsvc</li> </ol>
97076, 76683	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.
101471	Object updates in the Volume View may be incorrect. <b>Workaround:</b> Close then re-open the Volume View.
102306	VEA fails to start and goes into a continuous loop if the <code>vea.preference</code> file is corrupt. Refer to <a href="http://seer.support.veritas.com/docs/248359.htm">http://seer.support.veritas.com/docs/248359.htm</a> for more information.
102368	Object updates in the Disk/Volume Map View may be incorrect. <b>Workaround:</b> Close then reopen the Disk/Volume MapView.
102970	Strings in the search dialog could become corrupted when running VEA in non-English locales. <b>Workaround:</b> Run VEA after setting the <code>LC_ALL</code> environment variable to <code>C</code> , using the command: # export LC_ALL=C
103454	The following error may appear when replacing a failed disk: vxvm:vxdg: ERROR: associating disk-media disk01 with clt10d0s2: Disk public region is too small For more information and workarounds, please refer to Situation 3 in “Replacing Failed Disks” on page 41.
104205	When resizing concat-pro and stripe-pro volumes, a spurious error message may appear; this can be safely ignored. See “Utility Issues” for more information.
104526	The <b>Actions &gt; Disk Mirror</b> menu is incorrectly disabled if you do not have a full VxVM license. <b>Workaround:</b> Use the <code>vxmirror</code> command line to mirror the disk.





VERITAS Incident Number	Description
None	<p>The following X Window System error may occur when starting VEA:</p> <pre>Xlib: connection to "hostname:0.0" refused by server Xlib: Client is not authorized to connect to Server</pre> <p><b>Workaround:</b> Allow X server access by typing:  <code>xhost + [hostname]</code></p>
None	<p>VEA currently does not allow you to split/join shared disk groups.</p> <p><b>Workaround:</b> Use the command line interface to effect the split/join.</p> <p><b>Note</b> After using the command line to split the disk group, VEA will not allow you to split/join the same disk group. However, if you right-click on the split disk group, the source and target disk groups are displayed, allowing you to rejoin the split.</p>

## Adding Unsupported Disk Arrays to the JBOD Category

Use this procedure when adding new array support after VxVM is installed. For more information, see `vxddladm help addjbod`.

1. Enter the following command to add the new JBOD category:

```
# vxddladm addjbod vid=vendorid
```

where *vendorid* is FUJITSU, IBM, or SEAGATE as appropriate.

---

**Note** In VxVM 3.5, SEAGATE is added as JBOD device by default.

---

2. Enter the following command to bring the arrays under VxVM control:

```
# vxctl enable
```

This procedure ensures that dynamic multipathing is set up correctly on the array. Otherwise, VxVM treats the independent paths to the disks as separate devices, which can result in data corruption.



## Encapsulating and Mirroring the Root Disk

If you plan to mirror the root disk (which contains the root file system) so that an alternate root disk exists for booting purposes, you should place the root disk under VxVM control through encapsulation. The root disk can be encapsulated either during the `vxinstall` process (while installing VxVM), from the `vxdiskadm` menu (after VxVM is installed), or from VEA.

---

**Note** If the root disk is full, you may not be able to use the `vxdiskadm` menu. See “Full Disk Mirroring Considerations” on page 46 for the workaround if needed.

---

Once encapsulated, the root disk can be mirrored using `vxdiskadm`.

---

**Caution** If the root disk is encapsulated and the dump device is covered by the swap volume, the `savecore -L` operation is unsafe because it will overwrite the swap. As a workaround, you can configure some other partition as a dedicated dump device instead of the default swap.

---

You cannot grow or shrink any volume associated with an encapsulated boot disk (`rootvol`, `usr`, `var`, `opt`, `swapvol`, etc.) because these map to a physical underlying partition on the disk and must be contiguous. Refer to the section on rootability in the *VERITAS Volume Manager Administrator's Guide*.

## Full Disk Mirroring Considerations

The larger private region size that was introduced in VxVM 3.2 may create mirroring scenarios which require manual instructions in order to work. Three such scenarios are:

- u Mirroring a full root disk to a target disk which is exactly the same size as the source disk. A full disk means a disk that has no free cylinders.
- u Mirroring a disk created using an earlier version of VxVM to a target disk which is exactly the same size as the source disk. You only need to use this step if mirroring using `vxdiskadm` fails. For details, see `vxdiskadm(1M)`.
- u Mirroring a full VxVM disk (not a root disk) that was encapsulated in 3.5 to a target disk that is the same size as the source disk. You only need to use this step if mirroring using `vxdiskadm` fails. For details, see `vxdiskadm(1M)`.

In these scenarios, you may need to follow these steps to successfully create a mirror.

---

**Note** In the following examples, the traditional `c##d##` naming scheme is used. If you are using enclosure-based naming scheme, use the enclosure name instead.

---

1. Determine the size of the source disk's private region, using either of the following methods:
  - u If the source disk is a root disk, get the source disk's private region length using:



```
# vxprint -l rootdisk
```

where **rootdisk** is the disk media name of the root disk. This is typically *rootdisk*.

In the output, find the `privlen` value. In this example, the value is 3071:

```
devinfo: publen=39846240 privlen=3071
```

- u If the source disk is not a root disk, get the source disk's private region length using:

```
# vxdisk list diskname
```

where **diskname** is the disk media name of the source disk.

In the displayed output, note the `len` (length) value for the `private` field. In this example, the value of this field is 3071:

```
private: slice=4 offset=1 len=3071
```

2. Use the `vxdisksetup` program to initialize the target disk:

```
# /usr/lib/vxvm/bin/vxdisksetup -i c#t#d# privoffset=0 \
privlen=XXXX publen=XXXX
```

where `XXXX` is the size of the source disk's private region.

3. Add the newly initialized target disk to the source disk group:

```
# vxdg -g diskgroup adddisk medianame=c#t#d#
```

4. Use the `vxdiskadm` command and choose item 6 (Mirror volumes on a disk) to create the mirror:

- a. Use disk media name from step 1 as the source disk.

- b. Use **medianame** as the target disk.



## Booting From DMP Devices

When the root disk is placed under VxVM control, it is automatically accessed as a Dynamic Multipathing (DMP) device with one path if it is a single disk, or with more paths if the disk is part of a multiported disk array. By encapsulating the root disk, the system reliability is enhanced against loss of one or more of the existing physical paths to a disk. For more information, see the *VERITAS Volume Manager Administrator's Guide*.

## 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 become 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 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 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 `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.

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 `vxdbg(1M)` manual page and the section on moving disk groups between systems in the *VERITAS Volume Manager Administrator's Guide*.

## Corruption of Disk Group Configuration

If `vxdbg import` 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 must probably rebuild your configuration from scratch and reload all volumes in the disk group from a backup. To backup and rebuild the configuration, if nothing has changed, use `vxprint -mispvd` and store the output which can be fed to `vxmake` to restore the layouts. 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 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, VxFS, and for several popular databases.

The `-t` option to `vxchg` 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.

## Enclosure-based Naming

Starting with the previous 3.2 release, VxVM allows you to use either `c##d##s#` disk naming or enclosure-based disk access naming. Prior to VxVM 3.2, all disks were named according to the `c##d##s#` format and fabric mode disks were not supported by VxVM.

Disk devices controlled by MPXIO are always in fabric mode (irrespective of its hardware configuration), and are therefore named in the *enclosure name* format. This is true for both naming schemes.

The new enclosure-based naming scheme operates as follows:

- u All fabric or non-fabric disks in supported disk arrays are named using the *enclosure\_name\_#* format. For example, disks in the supported disk array, `enggdept` are named `enggdept_0`, `enggdept_1`, `enggdept_2` and so on. (You can use the `vxddmpadm` command to administer enclosure names.)
- u Disks in the DISKS category (JBOD disks) are named using the `Disk_#` format.
- u Disks in the OTHER\_DISKS category are named as follows:
  - Non-fabric disks are named using the `c##d##s#` format.
  - Fabric disks are named using the `fabric_#` format.
- u Enclosure names are limited to 25 characters. Disk access names include these 25 characters plus an underbar and numbering (`_nnnnn`). The maximum length of a disk access name is 31 characters.
- u You cannot mix naming schemes. Disks names must be in either `c##d##s#` format or enclosure-based name format.

See “Issues Regarding Simple/Nopriv Disks with Enclosure-Based Naming” in the “Administering Disks” chapter of the *VERITAS Volume Manager Administrator’s Guide* for details of how to switch between the two naming schemes.

To display the native OS device names of a VM disk (such as `disk01`), use the following command:

```
# vxdisk list diskname
```



## Enclosure-based Naming on Persistent Simple or Nopriv Disks

If you have persistent simple or nopriv disks, enclosure-based naming can cause errors in two situations:

1. After Upgrading to VxVM 3.5.
2. After Changing Your Disk Naming Scheme.

(For more information, refer to “Issues Regarding Persistent Simple/Nopriv Disks with Enclosure-Based Naming” in the *VERITAS Volume Manager Administrator’s Guide*. See also `vxdarestore (1M)` manual page.)

### After Upgrading to VxVM 3.5

If you upgrade your system to VxVM 3.5 and you have persistent simple or nopriv disks and you select to change from `c#t#d#` naming to enclosure-based naming, these disks may be in an error state after upgrading. To recover these disks, run `vxdarestore` after you have run the `upgrade_finish` script.

### After Changing Your Disk Naming Scheme

If you change your VxVM 3.5 system to enclosure-based naming using `vxdiskadm`, you need to check any persistent simple or nopriv disks, as these disks may be in an error state after the change. To recover these disks, run `vxdarestore`.

## Working with VxVM Disks on a Sun StorEdge T3 Array

On a VERITAS ServPoint SAN Appliance, or a T3 or T3+ array with the `c#t#d#` naming scheme, use the following procedures to initialize, remove, uninitialized, or replace VxVM disks.

---

**Note** `vxconfigd` is temporarily stopped and restarted each time the mode of the array is changed. Data I/O is always operating during this time.

---

### Initializing a Disk on a T3 or T3+ Array in `c#t#d#` Mode

1. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming mode to enclosure based format.
2. Issue `vxdisksetup` to add the new disk.



3. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming back to `c##t##d##` mode.

## Removing or Uninitializing a Disk on a T3 or T3+ Array in `c##t##d##` Mode

1. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming mode to enclosure based format.
2. Issue `vxdbg` to remove the disk from the disk group.
3. Issue `vxdiskunsetup` to remove the disk from VxVM control.
4. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming back to `c##t##d##` mode.

## Replacing a Disk on a T3 or T3+ Array in `c##t##d##` Mode

1. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming mode to enclosure based format.
2. Issue `vxdbg` to remove the disk from the disk group.
3. Issue `vxdiskunsetup` to remove the disk from VxVM control.
4. Replace the disk.
5. Issue `vxdisksetup` to add the new disk.
6. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming back to `c##t##d##` mode.

## Available Documentation

The following documents accompany this VxVM release:

- u *VERITAS Volume Manager Release Notes* (this document)
- u *VERITAS Volume Manager Hardware Notes*
- u *VERITAS Volume Manager Installation Guide*
- u *VERITAS Volume Manager Administrator's Guide*
- u *VERITAS Volume Manager User's Guide—VEA*





- u *VERITAS Volume Manager Troubleshooting Guide*
- u *VERITAS Volume Replicator Release Notes*
- u *VERITAS Volume Replicator Configuration Notes*
- u *VERITAS Volume Replicator Administrator's Guide*
- u *VERITAS Volume Replicator Administrative Services Administrator's Guide*
- u Online manual pages

VEA provides online help files for both VxVM and VVR. To access the online help files, select the appropriate item from the Help menu or click Help in a dialog box.

## Displaying Documentation Online

This product includes online documentation in Adobe Portable Document Format (PDF) formats. You can view the documents online in this format.

To view PDF documents, you must use the Adobe Acrobat Reader. You can use Acrobat reader as a stand-alone application, or as a plug-in to your web browser. However, VERITAS Software assumes no responsibility for the correct installation or use of Acrobat Reader. For more information on the latest versions of Acrobat Reader, or for help with installation problems, visit the Adobe web site at: <http://www.adobe.com>.

## English Versions of the Documentation

The manuals are provided on the CD-ROM under the `/CD_path/volume_manager/docs` directory. If you have installed the VRTSvmdoc package, the documents are available at `/opt/VRTSvxvm/docs`:

- u *VERITAS Volume Manager Installation Guide*  
`/opt/VRTSvxvm/docs/vxvm_ig.pdf`
- u *VERITAS Volume Manager Administrator's Guide*  
`/opt/VRTSvxvm/docs/vxvm_ag.pdf`
- u *VERITAS Volume Manager Troubleshooting Guide*  
`/opt/VRTSvxvm/docs/vxvm_tshoot.pdf`
- u *VERITAS Volume Manager Hardware Notes*  
`/opt/VRTSvxvm/docs/vxvm_hwnotes.pdf`
- u *VERITAS Volume Manager User's Guide — VEA*  
`/opt/VRTSvxvm/docs/vxvm_ug.pdf`



- u *VERITAS FlashSnap Point-In-Time Copy Solutions Administrator's Guide*  
/opt/VRTSvxvm/docs/pitc.pdf

Unformatted manual pages related to the VxVM are located in the `volume_manager/pkgs/VRTSvmman/reloc/VRTS/man` directory on the CD-ROM. If you have installed the VRTSvmman package, the manual pages can be found in the `/opt/VRTS/man` directory. If you add this directory to your MANPATH environment variable, you can view these man pages with the `man(1)` command.

The Volume Replicator manuals are installed under the following directories:

- u *VERITAS Volume Replicator Administrator's Guide*  
/opt/VRTSvxvm/docs/vr\_ag.pdf
- u *VERITAS Volume Replicator Configuration Notes*  
/opt/VRTSvxvm/docs/vr\_config.pdf

## Printing Documentation

To print the documentation, you must have access to a PostScript printer. If you are not sure how to do this, or whether or not you have this functionality, consult your system administrator.

You can print the documents using the print options in your Acrobat Reader viewer.