

VERITAS Volume Manager™ 3.2

Release Notes

Solaris

August 2001
30-000396-011


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 © 1998-2001 VERITAS Software Corporation. All rights reserved. VERITAS is a registered trademark of VERITAS Software Corporation in the U.S. and other countries. The VERITAS logo and VERITAS Volume Manager are trademarks of VERITAS Software Corporation. All other trademarks or registered trademarks are the property of their respective owners.

Printed in the USA/Ireland, August 2001.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
Phone 650-527-8000
Fax 650-527-2581
www.veritas.com



Contents

Important Notice for VERITAS Volume Replicator Users	1
Important Notice for SANPoint Foundation Suite Users	1
Organization	1
Getting Help	3
Conventions	4
Upgrading Disk Group Versions	5
Important Software Patch Information	6
Solaris Patches for VxVM	6
DMP Coexistence with Alternate Pathing	7
Getting Product Information Updates	8
New Features	9
Performance Enhancements in VxVM 3.2	11
Termination of Support Statements	12
Termination of Support for VxVM Releases	12
Additional Changes in Hardware and Format Support	12
Documentation Errata	13
Obtain a License Key	14
Software Fixes and Enhancements in VxVM 3.2	16
Software Limitations and Problems in VxVM 3.2	19
Firmware Upgrades	20
Installation Issues	22
Upgrade Issues	23
Utility Issues	23



Device Issues	26
Hot-Relocation Issues	27
DMP Issues	28
Cluster Functionality Issues	28
Data Integrity Issues	31
Miscellaneous Issues	31
Solaris Issues	33
VMSA	37
Software Fixes in VMSA 3.2	37
Software Limitations and Problems in VMSA 3.2	38
Adding Unsupported Disk Arrays to the JBOD Category	42
Encapsulating and Mirroring the Root Disk	42
Full Disk Mirroring Considerations	43
Booting From DMP Devices	44
VxVM and Multi-Host Failover Configurations	44
Import Lock	45
Failover	45
Corruption of Disk Group Configuration	46
Enclosure-based Naming	46
Using Enclosure-based Naming on Simple or Nopriv Disks	47
Working with VxVM Disks on a Sun StorEdge T3 Array	48
Available Documentation	49
Displaying Documentation Online	49
English Versions of the Documentation	50
Printing Documentation	52

VERITAS Volume Manager™ Release Notes

This document provides release information for the VERITAS Volume Manager™ (VxVM®) Release 3.2. This release includes the VERITAS Volume Manager Storage Administrator™ (VMSA) Release 3.2 graphical user interface.

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 “[Getting Product Information Updates](#)” on page 8.

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) 3.1.1 and VRAS 3.1.1, *do not install* VxVM 3.2 until you read the *VERITAS Volume Replicator Release Notes* for installation requirements.

Important Notice for SANPoint Foundation Suite Users

If you are using SANPoint Foundation Suite HA 2001-05 (which uses VxVM 3.1.1), *do not upgrade* VxVM 3.1.1 to VxVM 3.2.

A subsequent patch of VxVM 3.2 will be released to accommodate SANPoint Foundation Suite users.

Organization

- ◆ [Getting Help](#)
- ◆ [Conventions](#)
- ◆ [Upgrading Disk Group Versions](#)
- ◆ [Important Software Patch Information](#)



- ◆ [Getting Product Information Updates](#)
- ◆ [New Features](#)
- ◆ [Performance Enhancements in VxVM 3.2](#)
- ◆ [Termination of Support Statements](#)
- ◆ [Documentation Errata](#)
- ◆ [Software Fixes and Enhancements in VxVM 3.2](#)
- ◆ [Software Limitations and Problems in VxVM 3.2](#)
- ◆ [VMSA](#)
- ◆ [Software Limitations and Problems in VMSA 3.2](#)
- ◆ [Adding Unsupported Disk Arrays to the JBOD Category](#)
- ◆ [Encapsulating and Mirroring the Root Disk](#)
- ◆ [Booting From DMP Devices](#)
- ◆ [VxVM and Multi-Host Failover Configurations](#)
- ◆ [Enclosure-based Naming](#)
- ◆ [Working with VxVM Disks on a Sun StorEdge T3 Array](#)
- ◆ [Available Documentation](#)

Getting Help

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

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

For license information:

- ◆ Phone: 1-925-931-2464
- ◆ Email: license@veritas.com
- ◆ Fax: 1-925-931-2487

For software updates:

- ◆ Email: swupdate@veritas.com

For late-breaking news about this release, please see the section called “[Getting Product Information Updates](#)” on page 8.

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

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

<http://www.veritas.com>



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.
monospace (bold)	User input	# mount -F vxfs /h/filesys
<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	# mount -F vxfs \ /h/filesys
[]	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 only work on disk groups with the current disk group version, so you need to upgrade existing disk groups before you can perform these tasks. The following table summarizes the disk group versions that correspond to each VxVM release:

VxVM Release	Cluster Protocol Versions	Disk Group Version	Supported Disk Group Versions
1.2	n/a	10	10
1.3	n/a	15	15
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	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 *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) 3.2. You must reboot your system after installing these patches. If you have already installed VxVM, install the patches and then reboot your system.

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.

Some required system patches may already be present in your operating system. You should check to see if your system already contains the patches needed. Use the command `showrev -p` to display the patches included in your system. For more information, see `showrev(1M)`.

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 8	SUNWsan package (see your vendor to obtain) 109529 (-06 or later) 111413 (-02 or later)	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 these steps is important.
Solaris 7	107473 (-07 or later)	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.
Solaris 2.6	105375 (-20 or later) You do not need to install patch #105357: Ignore this install script message typo.	



Required Solaris Patches

For Solaris Version	Sun Patch ID Number	Notes
Solaris 8, 7, 2.6	110722 (-01 or later)	If you have Sun's Alternate Pathing (AP), you must install Solaris patch 110722-01 for VxVM 3.2 to co-exist with the AP driver version 2.3.1. However, you must also upgrade AP. Please see DMP Coexistence with Alternate Pathing before installing VxVM.
Solaris 7	106541 (-16 or later)	If you are running Solaris 7, you must install Sun patch ID 106541.

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 3.2 does not support earlier versions of AP, but supports Solaris 2.6 and higher.

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 3.2 automatically enables DMP if it was previously disabled.

Upgrading from VxVM 3.1.1

If you were running VxVM 3.1.1, upgrading to VxVM 3.2 preserves your DMP settings. If you had DMP prevented in VxVM 3.1.1, it remains prevented in VxVM 3.2.



Getting Product Information Updates

There may be additional product information about this release which is available on the VERITAS Technical Support website. To check for any new or updated information, do the following:

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 3.2 product documentation.

New Features

New features with VxVM 3.2 include the following:

◆ Persistent FastResync

There are two methods which FastResync uses to track changes to a volume: persistent, which is on-disk, and non-persistent, which is host memory based. The Persistent FastResync maintains FastResync maps on disk which enables the maps to survive system reboots and other system and cluster downtime. VERITAS recommends using the Persistent FastResync feature for the initial release of VxVM 3.2. Both methods will be recommended in the first patch to VxVM 3.2.

This recommendation is due to unpredictable results when two plexes of a mirror are detached at different times and reattached simultaneously. An example of this would be if a three-way mirror loses two mirrors to hardware failures and then has both mirrors repaired and reattached simultaneously.

Note Persistent FastResync requires a full FastResync license. You do not need a new license if you have an existing FastResync license.

◆ Disk Group Split and Join

This feature provides support for general disk group reorganization and for moving volume snapshots to another host for off-host backup. Additional options to the `vxchg` command enable you to take advantage of the ability to remove all VxVM objects from an imported disk group, and move them to a newly created target disk group (`split`), and to remove all VxVM objects from an imported disk group and move them to an imported target disk group (`join`). The `move` operation enables you to move a self-contained set of VxVM objects between imported disk groups.

Note Disk Group Split and Join requires a full FastResync license. You do not need a new license if you have an existing FastResync license.

◆ Device Discovery Layer

The addition of the Device Discovery Layer (DDL) to VxVM simplifies the process of adding support for new disk arrays to VxVM. The DDL gives VxVM the ability to add a new type of disk array dynamically without the need to reboot the system. New disk arrays are discovered using a device discovery feature that is part of the `vxconfigd` daemon; this feature can be invoked manually by running `vxctl enable`.

◆ Enclosure-Based Naming

This feature allows disk devices to be named for enclosures rather than for the controllers through which they are accessed (that is, `c##t##s##` device names). In a Storage Area Network (SAN), information about disk location provided by the



operating system may not correctly indicate the physical location of the disks. For example, using `c#t#d#s#` naming, disks in separate enclosures that are connected to the same host controller are assigned controller-based device names. The new enclosure-based naming allows VxVM to name enclosures as separate physical entities.

In many advanced disk arrays, you can use hardware-based storage management to represent several physical disks as one logical disk device to the operating system. In such cases, VxVM also sees a single logical disk device rather than its component disks. For this reason, when reference is made to a *disk* within an enclosure, this disk may be either a physical or a logical device.

For information about using `vxprint` to view `c#t#d#s#` device names, see [“Using vxprint with Enclosure-Based Disk Names”](#) on page 53 of the *VERITAS Volume Manager Administrator’s Guide*.

- ◆ Specifying Ordered Allocation of Storage to Volumes

A new *useopt* argument, *ordered*, to the `vxassist -o` option, allows you to specify that storage is allocated in the order: firstly concatenate disks, then form columns, and finally form mirrors.

- ◆ MPXIO: VxVM 3.2 provides support for coexistence with Sun’s multipathing `mpxio` driver.

- ◆ New location for man pages

The new default location for the VxVM man pages is `/opt/VRTS/man`. You should make the appropriate change to the `MANPATH` variable.

- ◆ Cluster support for SmartSync™ Recovery Accelerator for Mirrored Oracle® Databases

SmartSync works with Version 7.3.2 or later of the Oracle Universal Database™. This feature is already implemented in VxVM and is now supported in a cluster environment.

- ◆ Clusterization of Layered Volumes is implemented in VxVM 3.2

- ◆ 8-node Support

The cluster functionality of VxVM now allows up to 16 nodes in a cluster to simultaneously access and manage a set of disks under VxVM control.

- ◆ New Licensing Package: VRTSlic

In VxVM 3.2, licensing binaries have been separated out from the `VRTSvxvm` package into their own `VRTSlic` package. This package must be installed before the `VRTSvxvm` package; the Installation Guide has been updated to take account of this change.

- ◆ Starting with 3.2, VxVM supports SunFire 3800, 4800, and 6800 systems.

- ◆ VMSA has new icons.

- ◆ The private region size has been increased to provide larger default sizes and maximum database and kernel log sizes. This can affect root disk mirroring. See [Full Disk Mirroring Considerations](#) for more information.

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

Performance Enhancements in VxVM 3.2

VxVM 3.2 offers significant enhancements which speed transactions, reduce processing time, and improve bandwidth usage. These improvements do not require special tuning and are particularly noticeable in large configurations.

- ◆ The maximum size of a private region has been doubled; this allows large sites to create more volume manager objects in a disk group, and is particularly useful for users who use layered volumes.
- ◆ Object creation time (especially with `vxassist`) has been reduced dramatically for configurations over approximately 5,000 objects. Volume creation time remains nearly constant up to and beyond 50,000 objects.
- ◆ Disk group import times are faster, due to improvements in object indexing and new code which improves I/O bandwidth usage while validating the on-disk configuration copies. All configurations notice faster import times, but large configurations experience the most dramatic improvements. If existing large disk groups are combined, the results are even more impressive.
- ◆ System management commands like `vxprint` and `vxstat` run in a fraction of the time, and have less impact on other operations.
- ◆ Modifying existing objects with commands such as `vxedit`, `vxmend` and `vxdisk` takes less time. In large configurations, operations are almost instantaneous instead of taking up to half a minute.
- ◆ Now any system configuration can reach the maximum I/O throughput of the system without approaching any VxVM limitations. Large numbers of parallel resynchronous operations no longer have a significant effect on other administrative commands.



Termination of Support Statements

Termination of Support for VxVM Releases

The following software is no longer supported by VERITAS:

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

Additional Changes in Hardware and Format Support

- ◆ The Volume Manager Visual Administrator (VxVA) software is no longer available with VxVM.
- ◆ Currently, documentation on the CD-ROM is shipped in Postscript and PDF formats. The next release of VxVM will provide documentation only in PDF format and will not be available in Postscript format.

- ◆ The Sun Sparc Storage Array (SSA) is no longer supported. Please disregard any references to this array in all VxVM documentation.
- ◆ A license for RAID-5 is no longer needed to use this feature or VxVM in general.

Documentation Errata

Please note the following changes to the documentation for VxVM 3.2.

VERITAS Volume Manager Installation Guide

1. On page 50, under the section title *Configuring New Disks*, please replace the command shown in Step 3. Instead of using `vxdisk scandisks` to see the newly initialized disks, use the command `vxctl enable`. The `vxdisk scandisks` command is invalid in VxVM 3.2.
2. VERITAS has a new corporate fax number that replaces the number shown in VxVM 3.2 documentation. Please make a note of the correct number:
650/527-2581
3. Please replace the one page section called “*Obtain a License Key*” of the *VERITAS Volume Manager Installation Guide*. Please insert the following page in place of the original page 7.



Obtain a License Key

During installation, you will be prompted for a license key, which is issued by VERITAS Customer Support.

Before contacting VERITAS, you need the host ID and the machine type.

1. Obtain the host ID. Enter:

```
# hostid
```

2. Obtain the host machine type. Enter:

```
# uname -i
```

The host type is listed in the first line of output that follows a blank line. For example, the host type of a SPARC server 1000 is `SUNW,SPARCserver-1000`.

3. A License Key Request Form is included in your product package. Fill in the form and fax it to VERITAS Customer Support at 1-925-931-2487. (See [“Getting Help”](#) on page ix for contact information.)

Upgrades

If you are upgrading VxVM, you can view the currently installed licenses with the following command:

```
# vxlicense -p
```

and

```
# /usr/lib/vxvm/bin/vxliccheck -pv
```

You may need to add new licenses to support features in VxVM 3.2. Add new licenses with the following command:

```
# vxlicense -c
```



Software Fixes and Enhancements in VxVM 3.2

The following problems have been fixed in the 3.2 Release. (Lists of incidents fixed in previous releases of VxVM can be found in the applicable Release Notes.)

VERITAS Incident Number	Description
25660	The issue that <code>pkgm</code> does not remove <code>/dev/vx</code> has been fixed.
25836	VxVM now breaks up I/O requests that are greater than pool size, and also reduces stabilization of mirrored volumes to improve performance. (Sun Bug I.D. 4379956)
29900	The issue that <code>vxprint -m specify_writecopy</code> can be turned on but not off has been fixed.
30427	In a clustered environment where the master node is running <code>vxvol stopall</code> and <code>vxldg deport</code> , the <code>vxconfigd</code> daemon on a slave node no longer core dumps with a SEGV violation when a <code>vxdisksetup</code> of local disks is attempted. (Sun Bug I.D. 4416683)
31074	The minimum, default and maximum private region lengths have been increased to 160, 2048, and 524288 blocks respectively to accommodate layered volumes. (Sun Bug I.D. 4402252)
31139	VxVM now closes layered opens on shutdown.
36454	The system shutdown script for VxVM has been modified to check and automatically update the versions of libraries that <code>vxconfigd</code> uses at system startup before <code>/usr</code> is mounted.
40024	Names for disk names, plex names and volume names can be up to 31 characters.
41428	The <code>vxbootsetup</code> utility now creates a swap partition on a mirrored root disk regardless of where the primary swap partition appears in the <code>/etc/vfstab</code> file. (Sun Bug I.D. 4291994)
43074	The start-up time required for the import of large disk groups has been reduced, and the performance of configuration transactions on large disk groups has been improved.
44662	The <code>vxresize</code> command now allows resizing of volumes other than those of type <code>fsген</code> or <code>raid5</code> (such as <code>ген</code>) provided that the <code>-f</code> option is specified.
44744	VxVM now supports live savecore.
48303	The <code>vxassist convert</code> command no longer behaves inconsistently when attempting to convert to a target layout format that is the same as the source layout format.

VERITAS Incident Number	Description
50434	The VRTSvmmman package no longer creates a spurious /opt/opt/VRTSvxvm directory during an upgrade of VxVM. (Sun Bug I.D. 4354976)
50501	Installation of VxVM is no longer allowed on an unsupported version of Solaris. (Sun Bug I.D. 4344709)
50659	The vxresize command no longer fails if largefiles is not enabled on a VxFS file system.
51121	The layout=diskalign attribute to vxassist relayout now correctly aligns subdisks on cylinder boundaries.
51621	Hot-Relocation is now possible on version 60 disk groups even if the highest disk group version supported by VxVM is greater than 60.
52861	The vxassist relayout command no longer fails when applied to an A5200 (photon) array with only a 98 license present. (Sun Bug I.D. 4371487)
53042	In a cluster, the vxconfigd daemon on a slave no longer hangs due to certain logging failures on the master node.
54632, 62149	A system panic that could occur when the ssd driver is used in conjunction with underlying third-party SCSI drivers has been fixed. (Sun Bug I.D. 4363977)
54856, 54721, 54717, 54718, 54719, 54539, 55342	Significant performance enhancements have been made to VxVM 3.2 For details, please see the section called " Performance Enhancements in VxVM 3.2 " on page 11.
54875	Disk groups with version less than 40 are now correctly supported for certain transactions when a disk fails. (Sun Bug I.D. 4390757)
55244	The data change maps on a layered volume no longer fail if one or more log plexes are added to a layered replicated volume.
55575	The vxassist make command no longer creates unstartable stripe-mirror volumes if the specified volume size is 1 sector greater than a integral multiple of the stripe width.
55638	The attribute nmirror=src_nmirror may be specified to vxassist to preserve the number of mirrors during a volume relayout. The recommended way of changing the number of columns during a relayout is to use the ncol attribute. (Sun Bug I.D. 4398953)
56510	The Clariion ATF 3.1.0 is supported for use with DMP in VxVM.
57467	The upgrade_finish script now checks correctly for a Solaris 8 environment. (Sun Bug I.D. 4403373)
57670	The DESC line for pkginfo -l VRTSvmmdev is now complete.



VERITAS Incident Number	Description
58257	A VxVM tunable has been added to shorten delays in reporting an I/O error to VxVM. Previously, the Solaris disk driver could take a long time to report the error, due to DMP and error correction retries. Now the system checks the time taken by the driver to report errors. If this time is more than the value of <code>dmp_failed_io_threshold</code> , retries are not done and the error is reported to the application. This tunable has a default value of 10 minutes.
58304	The <code>vxdiskadm</code> menu display has been improved.
58623	The <code>vxconfigd</code> daemon now does not crash on reboot after encapsulating the root disk on a FCAL (diamond/crystal) or A5200 (photon) disk array. (Sun Bug I.D. 4412150)
59331	The <code>vxunroot</code> command no longer incorrectly programs the OBP boot-device string for fibre-channel devices. (Sun Bug I.D. 4416609)
60126, 60184	VxVM now automatically generates a license key on installation if you have an A5200 (photon) array with a Sun StorEdge PCI Single Fibre Channel network adapter card (X6799A). (Sun Bug I.D. 4313519, 4413604 and 4413609)
61236	Before VxVM 3.2, a disk was not marked as failed and a plex was not detached when a disk containing a simple volume failed. Now, when I/O errors occur on the last plex of a volume, <code>vxconfigd</code> gets notified and performs a validate to determine whether the disk should be disabled and the plex should be detached.
63746	Volume resize using <code>vxassist</code> on a layered volume no longer fails.
66168	Enabling FastResync on Existing Volumes with Associated Snapshots If a volume was created before release 3.2 of VxVM, and it has any attached snapshot plexes or it is associated with any snapshot volumes, follow the procedure given in the section "Enabling Persistent FastResync on Existing Volumes with Associated Snapshots" in the chapter "Administering Volumes" of the Administrator's Guide. If you do not perform the reconfiguration described in this section, full resynchronization is required every time that <code>snapback</code> is used to reattach a snapshot to its original volume.
67068	Before VxVM 3.2, <code>dumpadm</code> content type customization (<code>dumpadm -c</code>) was lost upon reboot with encapsulated root. This has been fixed.
67686	<code>vxdiskrm</code> does not return appropriate help or usage message. (Sun Bug I.D. 4462419) The <code>vxdiskrm</code> utility is obsolete. Use the <code>vxdiskunsetup</code> utility instead.
67833	The problem with <code>vxresize</code> (<code>fsadm</code>) failing to recognize invalid length specification has been fixed.

VERITAS Incident Number	Description
67867	SAN-aware vxassist works with SAL authentication turned on.
68004	vxdkmpadm doiocctl command is no longer required.
68553	vxdiskadm's "Suppress devices from VxVM's view" no longer lists erroneous devices/paths. (Sun Bug I.D. 4464764)
68906	vxconfigd no longer reports erroneous/misleading SCSI INQUIRY messages. (Sun Bug I.D. 4464389)
68909	The Activity cursor for device discovery during vxinstall is now terminated when a control-c is issued. (Sun Bug I.D. 4463865)
69189	The problem with a Raid5 volume in rootdg (constructed with T3) not starting on reboot has been fixed. (Sun Bug I.D. 4464739)
69315	VxVM 2.6 installation no longer refers to an obsolete patch (105223). The correct patch is 105181. (Sun Bug I.D. 4465903)
69318	The error in which vxdiskadd adds a disk twice during encapsulation has been fixed. (Sun Bug I.D. 4465371)
69687	The error in which no valid log copies are found when trying to import a shared disk group has been fixed.
69939	The problem in which vxvg join produced strange errors when rootdg was the source has been fixed.
72756	The problem of being unable to initialize T3 disks in partner-pair configuration has been fixed. (Sun Bug I.D. 4480992)
73206	The vxassist error when creating RAID5 volume on MPXIO-enabled photon disks no longer occurs.
75174	The error seen with the command ddl_final_claim_device: No such file or directory has been fixed. (Sun Bug I.D. 4489407)

Software Limitations and Problems in VxVM 3.2

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

[“Firmware Upgrades”](#) on page 20

[“Installation Issues”](#) on page 22

[“Upgrade Issues”](#) on page 23

[“Utility Issues”](#) on page 23

[“Device Issues”](#) on page 26



[“Hot-Relocation Issues”](#) on page 27

[“DMP Issues”](#) on page 28

[“Cluster Functionality Issues”](#) on page 28

[“Data Integrity Issues”](#) on page 31

[“Miscellaneous Issues”](#) on page 31

[“Solaris Issues”](#) on page 33

Firmware Upgrades

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

You need to get firmware upgrades as appropriate from your disk drive vendor

To upgrade the firmware for disk controllers, it is not necessary to reboot the system and to unload the VxVM in-kernel drivers `vxvmp`, `vxio` and `vxspec` drivers (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, set up the configuration as follows:

1. Disable the plex associated with the disk device:

```
# /usr/sbin/vxplex -g diskgroup det plex
```
2. Stop I/O to all disks through the specified HBA by executing the following command:

```
# /usr/sbin/vxdmpadm disable ctlr=first_ctlr
```

For last controller on this device, enter:

```
# /usr/sbin/vxdmpadm -f disable ctlr=second_ctlr
```
3. Upgrade the firmware on those disks for which the controllers have been disabled using the procedures from the disk drivers' website.
4. After doing the upgrade, enable all the controllers by executing:

```
# /usr/sbin/vxdmpadm enable ctlr=second_ctlr
```
5. Enable the plex associated with the device:

```
# /usr/sbin/vxplex -g diskgroup att volume plex
```

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

6. Upgrade the disk group:

```
# vxvg -T version upgrade diskgroup
```

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



Installation Issues

VERITAS Incident Number	Description
none	For internal SEAGATE disks, there is a danger of their being overwritten by the VxVM initialization routines when you run <code>vxinstall</code> . To avoid this, put the addresses of the disks (for example, <code>c1t1d0</code>) in the <code>/etc/vx/disks.exclude</code> file, or put the addresses of their controllers (for example, <code>c1</code>) in the <code>/etc/vx/cntrlr.exclude</code> file.
none	<p>Remove a VM disk from a disk group after evacuating any data on the disk. You can permanently remove this disk from VxVM control by removing the VM metadata partition from that disk. Use the low-level command, <code>vxdiskunsetup</code>, as follows:</p> <pre># /usr/lib/vxvm/bin/vxdiskunsetup c#t#d#</pre> <p>Note 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.</p>
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.
73583	<p>Installation script output asks you to install an incorrect patch #105357. This patch does not need to be installed.</p> <p>Workaround: Ignore the request to install patch 105357.</p>



Upgrade Issues

VERITAS Incident Number	Description
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>
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
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>



VERITAS Incident Number	Description
6154	There is no protection built into <code>vxassist</code> to prevent the user from shrinking the <code>swap</code> volume without first shrinking what the system sees as available swap space. If it is necessary to shrink the <code>swap</code> volume, this 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.
11286	Using <code>vxdbg free</code> with a non-existent <code>disk-media-name</code> does not print an appropriate error message; it only prints a header.
13418	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. This problem will be addressed in a future release.
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. If a log and a mirror are to be added, add the mirror and the log in two separate <code>vxassist</code> invocations, as follows:</p> <pre># vxassist mirror volume ... # vxassist addlog volume ...</pre> <p>This problem will be addressed in a future release.</p>
32559	The <code>vxr5check</code> utility does not work on degraded volumes.
45668	Resize of a mirrored volume with plexes of different types is currently unsupported. (Sun Bug I.D. 4339626)
none	<p>Due to the current implementation to handle the 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 shrinks or grows their sizes. This causes any ongoing operation, such as the 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, the user has 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>



VERITAS Incident Number	Description
63854	vxconfigd coredumps trying to import a unsupported DG version.
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 (_nnnnn).
72148	<p>Executing the command <code>vxddladm excludearray pid =<PID></code> and <code>vxddladm excludearray vid=<VID></code> will exclude more than intended. For example, entering <code>vxddladm excludearray vid=NEC</code> prevents you from using <code>vxddladm excludearray pid=<123></code>. The command incorrectly claims that <code>pid=123</code> has already been excluded.</p> <p>Workaround:</p> <p>To avoid problems, specify both the vid and pid in the same command. You can also specify only one ID, either vid or pid. For example, once you specify vid, you need to continue using vids in subsequent commands. If you specify a vid and then in a subsequent command, enter a pid, you will encounter an error. To clear exclusions, enter <code>vxddladm includearray all</code>.</p>
72734	The manual page for <code>vxddladm(1M)</code> and the help text from the command do not match. The attribute <code>lib</code> in the manual page should read <code>libname</code> .
73503	<p>While doing relayout on a mirrored volume, the <code>vxassist</code> command keeps the volume as mirrored even if the layout attribute is specified as <code>stripe</code> or <code>nomirror</code>. For example, see the following commands:</p> <pre># vxassist make vol lg layout=mirror-stripe ncol=3 # vxassist relayout vol layout=stripe ncol=2</pre> <p>The volume <code>vol</code> is converted to a 2-column volume, but it is still mirrored even if the layout attribute is specified as <code>stripe</code> and <code>nomirror</code>.</p>
76980	<code>vx dg free</code> reports free space differently than <code>vxprint -ht</code> .



Device Issues

VERITAS Incident Number	Description
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. This size defaults to 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 this 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>Workaround: There is no workaround to the problem of insufficient space on a disk to store private VxVM information. VxVM requires at least a small region of private storage (2048 blocks) for proper disk identification.</p>
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>Workaround: The only effective workaround is for the administrator to avoid performing exact physical disk copying.</p>
8818	<p>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.</p>



VERITAS Incident Number	Description
none	<p>If the system has a disk array connected to a switch in the Name server zone in fabric mode, you may see some of all of the following messages during boot-up:</p> <pre>vxvm:vxconfigd: WARNING: Disk xxxxx in group yyy: disk device not found</pre> <pre>vxvm:vxconfigd: ERROR: enable failed: Error in disk group configuration copies. No valid disk found containing disk group; transactions are disabled.</pre> <pre>vxvm:vxconfigd: FATAL ERROR: Rootdg cannot be imported during boot.</pre> <p>This messages occur early in the boot sequence. As boot-up progresses, VxVM finds these items if they are there. You can ignore these messages.</p>
none	<p>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.</p>
73154	<p>When Hitachi DF400 and DF500 arrays are configured as Active/Active mode, performance is degraded.</p>
75530	<p>T3 LUNs fails vxdisksetup with "Device does not match the kernel configuration" error. (Sun Bug I.D. 4491753)</p>

Hot-Relocation Issues

VERITAS Incident Number	Description
14894	<p>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).</p>
14895	<p>When a disk failure occurs, the hot-relocation feature notifies the system administrator of the failure and any relocation attempts through electronic mail messages. These 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.</p>



DMP Issues

VERITAS Incident Number	Description
29959	<code>vxddmpadm enable/disable ctrlr</code> succeeds even when an invalid controller name is specified.
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	In VxVM 3.1.1 and newer, DMP is a required module.
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	If a node leaves the cluster while a plex is being attached to a volume, the volume can remain in the SYNC state indefinitely. To avoid this, after the plex attach completes, resynchronize the volume manually with the following command # vxvol -f resync volume [Sun Bug I.D. 4087612]
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.



VERITAS Incident Number	Description
none	<p>It is possible to create a disk group containing disks that are attached to different nodes. This should be avoided. All disks in a given disk group should be attached to the same set of nodes.</p> <p>If CVM has deported a disk group because the disk group has lost access to one or more of its disks (due to a node leaving the cluster), the only way to try to regain access to the deported disks that are still attached to nodes in the cluster is to force-import the deported disk group. However, forcing an import in this situation is dangerous because it can cause mirrors to become unsynchronized in such a way that it cannot be determined which mirror has correct data.</p> <p>It is possible to have a private (non-shared) disk group consisting of disks that are attached to one node and owned by another node. This should be avoided.</p> <p>After a cluster reconfiguration, I/O via NETDisk may hang and cause vxconfigd to hang. Removing the hung node(s) from the cluster may fix this problem.</p> <p>It is possible to have private (non-shared) disk groups on physically shared disks. If these disks are on controllers that have been designated for fencing (i.e., reserved by Sun Cluster), the owner of the private disk group may not be able to access it when it is not in the cluster. For this reason, creating private disk groups on shared disks is not recommended unless the system administrator is fully aware of the consequences.</p>
none	<p>The use of file systems on volumes in cluster-shareable disk groups can cause system deadlocks. In particular, file systems on any type of volume in a shared disk group may lead to deadlocks during cluster reconfiguration events. <code>fs_{gen}</code> volumes are not supported in shared disk groups; only <code>gen</code> volume types are supported.</p>



VERITAS Incident Number	Description
none	<p>When a node leaves the cluster due to clean shutdown or abort, the surviving nodes perform a cluster reconfiguration. If the leaving node attempts to rejoin before the cluster reconfiguration is complete, the outcome depends on whether the leaving node is a slave or master.</p> <p>If the leaving node is a slave, the attempt will fail with the error messages:</p> <pre>Resource temporarily unavailable and one of the following: {vxclust} return from cluster_establish is configuration daemon error -1 master has disconnected</pre> <p>Retry at a later time.</p> <p>If the leaving node is a master, the attempt generates disk-related error messages on both nodes and the remaining node aborts. The joining node eventually joins and may become master.</p>
none	<p>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 VMSA 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. The <code>vxconfigd</code> program must therefore be started on the master first.</p>
none	<p>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:</p> <pre>cannot assign minor #</pre> <p>This message is accompanied by the console message:</p> <pre>WARNING:minor number ### disk group <i>group</i> in use</pre>
none	<p>To use the VxVM cluster functionality with a SPARCStorage Array, you must use firmware level 3.4 or higher.</p>

Data Integrity Issues

VERITAS Incident Number	Description
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.

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. It should be imported temporarily when the intention is to deport a 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 autoimport 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>Workaround: Check the volumes before starting the application or place a sleep (<code>sleep sec</code>) before the last invocation of <code>vxrecover</code>.</p>
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.



VERITAS Incident Number	Description
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>
32576	<p>Using an A5x00 array on a PCI bus as an encapsulated root (boot) disk or an alternate root disk is not supported. If you encapsulate or mirror to an external root disk on an A5x00 array with a PCI bus, it may not be possible to boot from that disk.</p> <p>Workaround: Try booting the system using the <code>boot -r</code> command. [Sun Bug I.D. 4259045]</p>
none	<p>On machines with low 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. For example, such a situation can result during heavy I/O stress exercised against RAID-5 volumes for 24 hours on a 32-megabyte machine.</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>Workaround: Use symbolic links to the longer <code>/dev/vx/dsk/volname</code> paths from a shorter pathname.</p>

VERITAS Incident Number	Description
	<p>Per vxassist man page, the convert operation is used to change the volume layout from mirror-stripe to stripe-mirror or vise versa. It can also be used to change the volume layout from mirror-concat to concat-mirror or vise versa. In other words, any mirrored volumes can be converted between the flat layout and the layered layout with the convert option.</p> <p>However, in 3.1.1, if the source volume is not mirrored, the convert option will succeed without checking if the source layout is mirrored or not. Such behavior has been corrected in 3.2. If the user wishes to enhance the data availability by adding mirrors to a volume, the correct way to achieve the purpose is to use "vxassist mirror" to add a mirror to a volume. If the volume size is very large and the user wishes to take advantage of less recover time, then "vxassist convert" can be used to convert the volume to a layered layout.</p>

Solaris Issues

VERITAS Incident Number	Description
6211	<p>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).</p>
8948	<p>A Solaris SCSI disk driver offlines a disk when it is not available and notifies the console. When the disk is connected back to the system, it 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:</p> <pre># : < /dev/rdisk/c1t5d0s2</pre>
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>
none	<p>The UNIX dd command uses only lseek() to seek to a particular offset in a file. It does not use llseek(). This causes dd to fail on volumes greater than 2 gigabytes.</p>



VERITAS Incident Number	Description
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	<p>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:</p> <pre>VxVM: Root volumes are not supported on your PROM version.</pre> <p>Any attempt to encapsulate the root disk fails on these machines.</p>
none	<p>If multiple swap partitions are encapsulated on your disks, VxVM names them as swapvol, swapvol1, swapvol2, and so on. When the system is rebooted, the following error message is displayed:</p> <pre>/dev/vx/dsk/swapvol2 : Overlapping swap files are not allowed</pre> <p>This is due to a problem with the swapadd scripts, which use a faulty swap -l output when device names are longer than a particular number of characters. The swap devices, however, are correctly added and with no ill effects on the system. To avoid seeing this message, rename the swap volumes (other than swapvol) to swap1, swap2, and so on, instead of using the current swapvol[0-9] names.</p>

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, cd 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 as 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>, 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>



VERITAS Incident Number	Description
55735	While VxVM works with AP, occasionally multiple entries for the same AP metanode are visible to VxVM. This leads to duplicate error records for the same device in <code>vxdisk list</code> output. Each time <code>vxctl enable</code> is invoked, a new error record gets added to the <code>vxdisk list</code> output. A reboot is required each time this problem is encountered.



VMSA

VMSA consists of a server and a client. The VMSA client can be run on any machine that supports the Java 1.1 Runtime Environment (including Solaris, HP-UX, or Windows).

Note This version of VxVM supports VMSA version 3.2 only. If you have an older version of VMSA, you must install VMSA version 3.2.

This release of VxVM and VMSA supports and has been tested on Solaris 2.6, 7, and 8.

The VMSA server and client have been tested on Solaris 2.6, 7, and 8. The VMSA client has also been tested on Windows NT, Windows 2000, Windows 98 and WindowsME.

Note The VMSA 3.2 GUI client is not compatible with earlier versions of the VMSA server. When installing VMSA 3.2, both the VMSA server and the VMSA GUI client must be updated.

Software Fixes in VMSA 3.2

The following issues have been fixed in this release of VMSA.

VERITAS Incident Number	Description
62941	The problem with the display not showing the result of an operation, requiring user to scroll up to view it, has been fixed.
65639	VMSA server no longer hangs when object view is opened on a disk group and <code>vxvg destroy diskgroup</code> command is used.
66341, 66342	When creating two disk groups with identical volume names and attempting to mount a vxfs file system to a different mount point, the VMSA GUI encountered confusion and the operation failed. This is now fixed.



Software Limitations and Problems in VMSA 3.2

The following problems and issues exist in this release of VMSA.

VERITAS Incident Number	Description
17772	VMSA does not support destroying deported disk groups. Workaround: Use <code>vxdg</code> to import the disk group, then deport the diskgroup.
18338	VMSA does not allow the user to specify volume usage types.
20495	Launching online help under a modal dialog prevents access to the help window. You can use the help window's Close button to close the help window or you can close the corresponding dialog.
22456	The maxsize operation always returns sizes in sectors.
22482	The VMSA server can hang on Solaris while waiting for <code>ncsd(1M)</code> to reply to a name service lookup. This can occur when the system is set up for DNS and a DNS server is not set up. Workaround: Remove the <code>dns</code> entry from the <code>hosts: files dns</code> line in the <code>/etc/nsswitch.conf</code> file.
22730	Online help is not supported for non-English locales. If the user's <code>\$LANG</code> environment variable is set to a non-English value, the online help files can have problems printing and following their hyperlinks. Workaround: Create a symbolic link using the following command: <pre># cd /opt/VRTSvmsa/vxvm/java ln -s help help_locale</pre> where <i>locale</i> is the appropriate locale abbreviation.
23730	The splitter cursor does not always go away. This can prevent the wait cursor from being displayed. This problem is caused by a Java bug. Workaround: Move the cursor outside the main window and then back into the main window.
24701	When a dialog box such as a disk/space allocation dialog box is displayed, a message similar to the following is displayed: <pre>Warning: Name: scrollbar Class: XmScrollbar The scrollbar page increment is less than 1.</pre> This message can be ignored.



VERITAS Incident Number	Description
25089	<p>If you remove a VERITAS Volume Manager Storage Administrator 1.x package and then install a 3.x package, the new VMSA server may not start properly.</p> <p>Workaround: Before you remove the 1.x package, run the <code>server.sh -k</code> utility to stop the 1.x server.</p>
25361	<p>When the window manager's interactive placement setting is turned on, VMSA windows may not be sized properly. The windows can be very small or very large.</p> <p>Workaround: Turn off the interactive placement setting. For example:</p> <pre>Mwm*interactivePlacement: False</pre>
27348	<p>When the VMSA is run with the <code>fvwm</code> window manager, window decorations appear above the top of the screen.</p>
27349	<p>When the Volume to Disk Mapping window is updated, the contents of the window may become inaccurate.</p> <p>Workaround: Close the Volume to Disk Mapping window and then reopen it.</p>
29621	<p>When disks are added, the order in which the VxVM disk names are assigned may not match the order of the selected or specified devices. The disks are processed in the order in which they appear in the grid.</p>
31029	<p>The <code>VRTSvmsa</code> client-only installation output states that <code>VRTSvxvm</code> is a prerequisite. The <code>VRTSvxvm</code> package must be installed on the machine where you install the server portion of the <code>VRTSvmsa</code> package, but <code>VRTSvxvm</code> is <i>not</i> a prerequisite for the VMSA client.</p>
32052	<p>VMSA does not support communication between the client and server across a firewall.</p>



VERITAS Incident Number	Description
32599	<p>VMSA may have trouble connecting to a host machine (server) if multiple host names are associated with a single IP address. VMSA displays the following message:</p> <pre>Summary: There is no such server (<i>host1</i>) Detail: java.net.UnknownHostException: Unknown host : [<i>host2</i>:32839]; nested exception is: java.net.UnknownHostException: <i>host2</i></pre> <p>In this example, the administrator specifies <i>host1</i> for the VMSA connection, but the server host machine is identified as <i>host2</i>.</p> <p>In some cases, this is caused by a problem with the way DNS is set up. A DNS reverse lookup (by IP address) can return a host name that differs from the host name provided to VMSA at startup (<i>host1</i>). In this case, make appropriate changes to DNS to ensure that the names are consistent. In other cases, the host name is different because <i>host1</i> is not the first host in the list of hosts for the associated IP address in the <code>/etc/hosts</code> file.</p> <p>Workaround: Make sure <i>host1</i> shows up as the first host in the list of names for the address of <i>host1</i> in <code>/etc/hosts</code>.</p>
33367	<p>VMSA hangs on Solaris 7 when the Wnn input method server or the CS00 input server is active under the Japanese locale environment. Wnn is the default. Sun also provides CS00, ATOK as X input methods and HTT as the Japanese Kana characters input server. Wnn and CS00 use the HTT as the input server.</p> <p>Workarounds:</p> <ol style="list-style-type: none"> 1. When using the Wnn Japanese input method server, the following Wnn patches are required to run VMSA on Japanese Solaris 7 or Solaris 8. 107636-04 for Solaris 7 108773-03 for Solaris 8 Patches may be downloaded from http://sunsolve.sun.com. 2. No patches are needed if using the ATOK input method server. Use the Solaris 7 ATOK input method server. ATOK is executed by selecting ATOK from the input method desktop menu item. Terminate the WnnWnn input method, log out, and log back in. ATOK is now the default input method server for the session. 3. Export <code>LANG=C</code> before executing VMSA. <p>Note VMSA strings are displayed in English.</p>
39544	<p>With Windows 98, the Selected menu may overwrite other menus.</p>

VERITAS Incident Number	Description
46077	<p>The following server-side exception can occur when terminating the client:</p> <pre>java.rmi.UnmarshalException: error unmarshalling return header.</pre> <p>You can ignore this exception.</p>
none	<p>To prevent core dumps on NCD terminals, the <code>jre/lib/font.properties</code> file has been renamed to <code>font.properties-</code>. On some machines, this has a negative effect on the appearance of default fonts for online help and other text areas. If you <i>do not</i> use NCD terminals, you can improve the appearance of these fonts by renaming <code>font.properties-</code> to <code>font.properties</code>.</p>
none	<p>The following X Window System error may occur when VMSA is started:</p> <pre>Xlib: connection to "hostname:0.0" refused by server Xlib: Client is not authorized to connect to Server java.lang.InternalError: Can't connect to X11 window server using hostname:0.0'as the value of the DISPLAY variable. at sun.awt.motif.MToolkit.<init>(MToolkit.java:48) at java.awt.Toolkit.getDefaultToolkit(Toolkit.java:244)</pre> <p>Workaround: Type <code>xhost + [hostname]</code> to allow X server access.</p>
none	<p>If a volume with an unmounted file system is resized, the file system may not be resized.</p>
none	<p>Under some circumstances, the contents of the Object View window do not display correctly.</p> <p>Workaround: Either resize the Object View window or close the Object View window and then reopen it.</p>
none	<p>VMSA cannot be started from a Web browser. Early releases of VMSA were run from a Web browser. However, due to the nature of Web browsers, running VMSA from a Web browser was not recommended because it resulted in security constraints, performance problems, and other issues.</p>
57571	<p>The watch cursor may remain in the Browse Volumes window.</p>
69157	<p>VMSA may not work correctly under the Gnome window manager and may cause system crashes. You must use a different window manager.</p>
72742	<p>Unable to mirror rootdisk via <code>vxdiskadm</code>. See the section on “Full Disk Mirroring Considerations” on page 43 for more information. (Sun Bug I.D. 4466429)</p>



VERITAS Incident Number	Description
73958	After renaming a mounted volume, it will fail to remount under the new volume name. The <code>/etc/vfstab</code> file may need to be manually updated with the new volume name.
74297	VMSA fails to enable FastResync on RAID-5 volumes. Use the command line interface to enable FastResync.

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.

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 VMSA.

Note If the root disk is full, you may not be able to use the `vxdiskadm` menu. See [“Full Disk Mirroring Considerations”](#) on page 43 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 section on rootability in the *VERITAS Volume Manager Administrator's Guide*.

Full Disk Mirroring Considerations

VxVM 3.2's larger private region size may create mirroring scenarios which require manual instructions in order to work. Three such scenarios are:

- ◆ 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.
- ◆ 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)`.
- ◆ Mirroring a full VxVM disk (not a root disk) that was encapsulated in 3.2 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#t#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 one of the following choices.
 - ◆ If the source disk is a root disk, get the source disk's private region length using:

```
# vxprint <rootdisk>Priv
```

where `<rootdisk>` is the disk media name of the root disk. This is typically `rootdisk`. In the displayed output, find the length value, for example 3071.
 - ◆ 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, find the length value for the `private` field. Use the value shown in `len=xxxx`. In this case, the value of this field is 3071.

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

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

where XXXX is one more than the size of the source disk's private region. For example, if the size of the disk's private region is 3071, then the value of XXXX is 3072.

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

```
# vxdg adddisk -g <diskgroup> <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 becomes accessible to the host. If the administrator or system software wants to privately use the same disk group from another host, the host that already has the disk group imported (*importing host*) must “deport” (give up access to) the disk group. Once deported, the disk group can be imported by another host.

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

Import Lock

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

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

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

Note Since VxVM uses the host name as the host ID (by default), it is advisable to change the host name of one machine if another machine shares its host name. To change the host name, use the `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 `vxdg(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 `vxldg 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 `vxldg` prevents automatic re-imports on reboot and is necessary when used with a host monitor (such as VCS) that controls imports itself, rather than relying on automatic imports by VxVM.

Enclosure-based Naming

Starting with version 3.2, 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:

- ◆ 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 `vxdmpadm` command to administer enclosure names.)
- ◆ Disks in the DISKS category (formerly known as JBOD disks) are named using the *Disk_#* format.
- ◆ Disks in the OTHER_DISKS category are named as follows:
 - Non-fabric disks are named using the *c#t#d#s#* format.
 - Fabric disks are named using the *fabric_#* format.
- ◆ Enclosure names are limited to 25 characters. Disk access names include these 25 characters plus an underbar and numbering (*_nnnnn*) so the total number is 31 characters.
- ◆ You cannot mix naming schemes. Disks names must be in either *c#t#d#s#* format or enclosure-based name format.

See “[Issues Regarding Simple/Nopriv Disks with Enclosure-Based Naming](#)” on page 53 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
```

Using Enclosure-based Naming on Simple or Nopriv Disks

For users with simple or nopriv disks, enclosure-based naming can cause errors in two situations: after upgrading to VxVM 3.2, and after changing your disk naming scheme.

If you have upgraded your system to VxVM 3.2 and you have simple or nopriv disks, they may be in an error state after upgrading. To recover these disks, you can run `vxdarestore` after you run the `upgrade_finish` script. For more information, refer to “[Issues Regarding Simple/Nopriv Disks with Enclosure-Based Naming](#)” in the *VERITAS Volume Manager Administrator’s Guide*. See also `vxdarestore (1M)`.

If you change your VxVM 3.2 system to enclosure-based naming using `vxdiskadm`, you need to check any simple or nopriv disks, as they may be in an error state after you change disk naming. If they are, use `vxdarestore` to restore these disks.



Note The next version of VxVM will use enclosure-based disk access names by default.

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, uninitialize, 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` command 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` command 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:

- ◆ *VERITAS Volume Manager Release Notes* (this document)
- ◆ *VERITAS Volume Manager Hardware Notes*
- ◆ *VERITAS Volume Manager Installation Guide*
- ◆ *VERITAS Volume Manager Administrator's Guide*
- ◆ *VERITAS Volume Manager Storage Administrator Administrator's Guide*
- ◆ *VERITAS Volume Manager Troubleshooting Guide*
- ◆ *VERITAS Volume Replicator Release Notes*
- ◆ *VERITAS Volume Replicator Configuration Notes*
- ◆ *VERITAS Volume Replicator Administrator's Guide*
- ◆ *VERITAS Volume Replicator Administrative Services Administrator's Guide*
- ◆ *VERITAS Volume Manager Storage Administrator for VVR Administrator's Guide*
- ◆ Online manual pages

VMSA 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) and PostScript formats. You can view the documents online in either of these formats.



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>.

To view PostScript documents, you can use the Solaris Image Tool (`imagetool`) or any PostScript previewer.

English Versions of the Documentation

The manuals are provided on the CD-ROM under the `/CD_path/pkgs/VRTSvmdoc/reloc/VRTSvxvm/docs` directory. If you have installed the VRTSvmdoc package, the documents are available in the following locations:

- ◆ *VERITAS Volume Manager Installation Guide*

`/opt/VRTSvxvm/docs/igbook.ps`
`/opt/VRTSvxvm/docs/igbook.pdf`

- ◆ *VERITAS Volume Manager Administrator's Guide*

`/opt/VRTSvxvm/docs/admin.ps`
`/opt/VRTSvxvm/docs/admin.pdf`

- ◆ *VERITAS Volume Manager Troubleshooting Guide*

`/opt/VRTSvxvm/docs/tshoot.ps`
`/opt/VRTSvxvm/docs/tshoot.pdf`

- ◆ *VERITAS Volume Manager Hardware Notes*

`/opt/VRTSvxvm/docs/hwnotes.ps`
`/opt/VRTSvxvm/docs/hwnotes.pdf`

- ◆ *VERITAS Volume Manager Storage Administrator Administrator's Guide*

`/opt/VRTSvxvm/docs/vmsaguide.ps`
`/opt/VRTSvxvm/docs/vmsaguide.pdf`

Unformatted manual pages related to the VxVM are located in the `VRTSvmman` directory on the CD-ROM. If you have installed the `VRTSvmman` package, the manual pages can be found in the `/opt/VRTS/man` directory. The VMSA manual pages are in the `/opt/VRTS/man` directory. If you add these directories 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:

- ◆ *VERITAS Volume Replicator Administrator's Guide*

`/opt/VRTSvxvm/docs/vr_ag.ps`
`/opt/VRTSvxvm/docs/vr_ag.pdf`

◆ ***VERITAS Volume Replicator Configuration Notes***

/opt/VRTSvxvm/docs/vr_config.ps
/opt/VRTSvxvm/docs/vr_config.pdf

◆ ***VERITAS Volume Manager Storage Administrator for VVR Administrator's Guide***

/opt/VRTSvxvm/docs/vr_vmsa_ag.ps
/opt/VRTSvxvm/docs/vr_vmsa_ag.pdf

◆ ***VERITAS Volume Replicator Administrative Services Administrator's Guide***

/opt/VRTSvras/docs/vr_as_ag.ps
/opt/VRTSvras/docs/vr_as_ag.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 in the following ways:

- ◆ Use the print options in your PostScript previewer to print one or more pages.
- ◆ Use the print options in your Acrobat Reader viewer to print one or more pages.
- ◆ Print entire chapters using the `lp` command and your PostScript printer.