

VERITAS Volume Manager™ 3.1

Release Notes

Solaris

August 2000
32-000000-399


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 © 1999-2000 VERITAS Software Corporation. All rights reserved. VERITAS is a registered trademark of VERITAS Software Corporation in the US and other countries. The VERITAS logo and VERITAS VxVM are trademarks of VERITAS Software Corporation. All other trademarks or registered trademarks are the property of their respective owners.

Printed in the USA, August 2000.

VERITAS Software Corporation
1600 Plymouth St.
Mountain View, CA 94043
Phone 650-335-8000
Fax 650-335-8050
www.veritas.com



Contents

Installation Notes	5
Organization	5
Getting Help	6
Conventions	7
Installing the Volume Manager	8
Overview	8
Installing Volume Manager Packages	8
Setting Up the Volume Manager	8
Upgrading the Volume Manager	9
Disk Group Versioning	9
New Features	10
Volume Manager	10
Cluster Functionality (Optional)	12
Termination of Support Statements	12
Software Problems Fixed in Volume Manager 3.1	13
Software Limitations and Problems in Volume Manager 3.1	15
Firmware Upgrades	15
FastResync	16
Installation Issues	16
Upgrade Issues	17
Utility Issues	18
Device Issues	20
Hot-Relocation Issues	21



Relayout Issues	22
DMP Issues	22
Cluster Functionality Issues	22
Miscellaneous Issues	26
Solaris Issues	28
Software Limitations and Problems in Storage Administrator Release 3.1	33
Encapsulating and Mirroring the Root Disk	37
Booting From DMP Devices	38
VxVM and Multi-Host Failover Configurations	38
Available Documentation	40
Displaying Documentation Online	40
English Versions of the Documentation	41
Printing Documentation	42



Important Release Information

Installation Notes

This document provides release information for the VERITAS Volume Manager™ (VxVM®) Release 3.1. This Release includes the Volume Manager Storage Administrator (VMSA) Release 3.1 graphical user interface.

Note Before you install the packages, read this entire document.

This version of Volume Manager supports VMSA version 3.1 only. If you have an older version of VMSA, you must install VMSA version 3.1.

The Volume Manager Storage Administrator consists of a server and a client. The Storage Administrator server must be run on a UNIX machine running Solaris Release 2.5.1, or higher. The Storage Administrator client can be run on any machine that supports the Java 1.1 Runtime Environment (including Solaris, HP-UX, or Windows).

This release of the Volume Manager supports and has been tested on:

- ◆ Solaris 2.6
- ◆ Solaris 7
- ◆ Solaris 8

The Storage Administrator server and client have been tested on Solaris Release 2.6 and higher. The Storage Administrator client has also been tested on Windows NT, Windows 2000, Windows 98, and Windows 95.

Organization

This guide is organized with the following sections:

- ◆ [Installation Notes](#)
- ◆ [Organization](#)
- ◆ [Getting Help](#)
- ◆ [Installing the Volume Manager](#)



-
- ◆ [Upgrading the Volume Manager](#)
 - ◆ [New Features](#)
 - ◆ [Termination of Support Statements](#)
 - ◆ [Software Problems Fixed in Volume Manager 3.1](#)
 - ◆ [Software Limitations and Problems in Volume Manager 3.1](#)
 - ◆ [Software Limitations and Problems in Storage Administrator Release 3.1](#)
 - ◆ [Encapsulating and Mirroring the Root Disk](#)
 - ◆ [Encapsulating and Mirroring the Root Disk](#)
 - ◆ [Booting From DMP Devices](#)
 - ◆ [VxVM and Multi-Host Failover Configurations](#)
 - ◆ [Available Documentation](#)

Getting Help

For information about VERITAS® service packages, contact VERITAS Customer Support:

US Customers: 1-800-342-0652
International Customers: +1-650-335-8555
Fax: 1-650-335-8428
Email: support@veritas.com

For license information:

Phone: 1-650-318-4265
Email: license@veritas.com
Fax: 1-650-335-8428

For software updates:

Phone: 1-650-526-2549
Email: swupdate@veritas.com

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

www.veritas.com

For additional information about the Knowledge Base and Technical Notes & Alerts, visit the Technical Support Web Site:

www.support.veritas.com



Conventions

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

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	<code># mount -F vxfs /h/filesys</code>
<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; you do not type this character	<code># mount -F vxfs \ /h/filesys</code>
[]	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>



Installing the Volume Manager

Overview

Volume Manager installation consists of three parts:

1. Installing the packages onto the system.
2. Configuring and setting up the Volume Manager.
3. Setting up the Storage Administrator.

Installing Volume Manager Packages

See the *VERITAS Volume Manager Installation Guide* for complete instructions on how to install VxVM using the `pkgadd` command.

The VERITAS CD-ROM contains the following packages:

- ◆ VRTSvxvm—Volume Manager Software (driver and utilities)
- ◆ VRTSvmdev—Developer Kit
- ◆ VRTSvmmman—Manual Pages
- ◆ VRTSvmdoc—Volume Manager Documentation
- ◆ VRTSvmsa—Storage Administrator Software
- ◆ VRTSvrdoc—Volume Replicator Documentation

Note VxVM is a licensed product; you must obtain a license key before you install VxVM. To obtain a license key, complete a License Key Request Form and fax it to VERITAS Customer Support (see “[Getting Help](#)” on page 6 for contact information). See the section on obtaining a license key in the *VERITAS Volume Manager Installation Guide* for additional information.

Setting Up the Volume Manager

See the *VERITAS Volume Manager Installation Guide* for information on how to initialize the Volume Manager (using `vxinstall`).

After the initialization is complete, you may be asked to reboot your machine to allow the kernel configuration to be updated. If necessary, you must reboot before any VxVM features are usable on your system. **Setting Up the Storage Administrator**



See the *VERITAS Volume Manager Installation Guide* for information on how to set up and start the Storage Administrator server and client.

Note The Storage Administrator Release 3.1 server is not backward compatible with Storage Administrator Release 1.x clients, so you must upgrade any existing client(s) to Release 3.1.

Upgrading the Volume Manager

See the *VERITAS Volume Manager Installation Guide* for instructions on how to upgrade to Volume Manager Release 3.1 and/or compatible releases of Solaris.

Disk Group Versioning

All disk groups have a version number associated with them. Each Volume Manager 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 Volume Manager release:

Volume Manager Release	Disk Group Version	Supported Disk Group Versions
1.2	10	10
1.3	15	15
2.0	20	20
2.2	30	30
2.3	40	40
3.0	60	20-60
3.1	70	20-70

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

```
# vxvg list dgname
```

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

VxVM upgrades the disk group to the highest version supported by the release of VxVM that is currently running. To upgrade a disk group, use the command:



New Features

Note Volume Manager Release 3.1 and Storage Administrator Release 3.1 are Year 2000 compliant. For additional information, see the Year 2000 certification statement on the VERITAS Web site (www.veritas.com).

Volume Manager

New features with VxVM 3.1 and higher include:

- ◆ Unrelocate
 - ◆ The new unrelocate feature can be used to return subdisks that have been relocated by the hot-relocation feature back to their original disk locations after the original disk is repaired or replaced
- ◆ Fast Mirror Resynchronization (FMR)
 - ◆ This feature requires a license.
 - ◆ The Fast Mirror Resynchronization (FMR) feature speeds up the resynchronization of mirrors in a volume.

The Fast Mirror Resynchronization (FMR) feature, as described in the VxVM documentation, is now also known as FastResync.
 - ◆ Significant changes have been made to both the `vxassist` and `vxplex` commands as they both pertain to the current snapshot functionality. Also, the `vxprint` and `vxkprint` commands have been changed to display useful information. The `vxvol` command has also been changed.
- ◆ DMP Enhancements
 - ◆ By default, the restore daemon checks the condition of only the disabled paths. A new option allows you to change the policy of the restore daemon to check the condition of all the paths connected to the host. To set the policy, use the following command:


```
# vxdmpadm start restore policy=check_all
```


To set the restore daemon's policy to check all the paths, comment the following line in `/etc/init.d/vxvm-sysboot`:


```
# restore_daemon_opts="interval=300 \  
policy=check_disabled"
```


and uncomment the following line:
 - ◆ DMP for JBOD and JBOD-type arrays require Solaris 2.6. or higher.
 - ◆ Support added for Sun StorEdge T3 disk array



-
- ◆ Storage Replicator for Volume Manager (optional)
 - ◆ This feature requires a license.
 - ◆ This release of Volume Manager includes Storage Replicator for Volume Manager (SRVM), which is also known as Volume Replicator (VR). SRVM is a data replication tool designed to contribute to an effective disaster recovery plan. SRVM is a separately licensed feature of Volume Manager. For release information, see the *Storage Replicator for Volume Manager Release Notes* (`vr_notes.pdf` and `vr_notes.ps`). The SRVM documentation set is available in the `pkgs/VRTSvrdoc` directory.

Note Volume Manager Storage Administrator 3.1 does not support SRVM. However, a future release of Storage Administrator will support SRVM.

The VERITAS Volume Manager Storage Administrator is the graphical user interface for the Volume Manager. New features with VMSA 3.1 include:

- ◆ Multiple Host Support
 - ◆ The Storage Administrator client can provide simultaneous access to multiple host machines. The administrator can use a single Storage Administrator client session to connect to multiple hosts, view the objects on each host, and perform administrative tasks on each host. Each host machine must be running the Storage Administrator server.
- ◆ Object View Window
 - ◆ The new Object View window displays a graphical view of volumes, disks, and other objects in a given disk group.
- ◆ Command Log View Window
 - ◆ The new Command Log View window displays a history of Volume Manager Storage Administrator tasks performed in the current session and previous sessions.
- ◆ Enhanced Read-only Mode
 - ◆ Administrators can run the Storage Administrator in read-only mode, which is useful for monitoring, training, or browsing purposes. Read-only mode allows administrators to view objects on the system, but prevents administrative actions from taking effect.
- ◆ Performance Monitoring
 - ◆ The Volume to Disk Mapping window has a performance monitoring feature that ranks volume response time.



Cluster Functionality (Optional)

This Volume Manager release includes an *optional* cluster feature that enables VxVM to be used in a cluster environment. Cluster functionality is discussed in Chapter 6, of the *VERITAS Volume Manager Administrator's Guide*.

- ◆ This feature requires a license.
- ◆ With cluster support enabled, this release of VxVM supports up to four nodes per cluster. However, support for more than two nodes is currently available only if VxVM is used with a Sun StorEdge A3000/A5000.

Note The new features introduced in Volume Manager 3.1 are available in private disk groups, but are not yet supported for shared disk groups.

The `logtype=seq` feature introduced in CVM Release 2.2.1, used for Dirty Region Logs, is not supported in this release because it requires changes to the on-disk layout.

The following VxVM features are now supported in a cluster environment:

- ◆ 32- and 64-bit Support
VxVM now provides 32- and 64-bit support and works with the Solaris 7 and Solaris 8 32- and 64-bit kernels in a cluster environment.
- ◆ Rolling Upgrade
The rolling upgrade feature can be used to upgrade a cluster to a new software version while the cluster remains online. Each node in the cluster is upgraded separately while the rest of the cluster remains online. When all nodes in the cluster are upgraded, the cluster is brought up to the current software version.

Termination of Support Statements

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



-
- ◆ Solaris 2.3 operating system
 - ◆ Solaris 2.4 operating system
 - ◆ Solaris 2.5 & 2.5.1 operating system
 - ◆ VERITAS Volume Manager no longer supports the Sun-4c product line:
 - ◆ SPARCstation 1
 - ◆ SPARCstation 1+
 - ◆ SPARCstation 2
 - ◆ SPARCstation IPC
 - ◆ SPARCstation IPX
 - ◆ SPARCstation SLC

The following software is no longer available with the VERITAS Volume Manager:

- ◆ Volume Manager Visual Administrator (VxVA)

Software Problems Fixed in Volume Manager 3.1

The following problems have been fixed in this release:

VERITAS Incident Number	Description
25581	config and log copy states are not distributed.
25722	vxrelocd does not select hot spares according to the "closeness" value. (Sun Bug I.D. 4154269)
26644	vxmake of root usage type volume gives odd error.
26802	File descriptor leak in vxconfigd when DMP is disabled. (Sun Bug I.D. 4187714)
27543	Disks are powered off, but the DA records are still online.
27635	vxunroot should be able to recreate the /export/home slice. (Sun Bug I.D. 4170778)
27787	Failed root disk cannot be replaced if not named rootdisk.
27794	Growing degraded RAID-5 volume will corrupt data.
30837	vxconfigd hangs during boot when AP devices are active. (Sun Bug I.D. 4238826)
30886	vxbootsetup does not properly make partitions on boot disks. (Sun Bug I.D. 4305586)



31195	A potential security problem that assigned incorrect permissions to files created by <code>vmsa_server</code> has been corrected.
33128	There is discrepancy between <code>vxdisk list</code> and <code>vxprint</code> output in CVM.
33905	Security problem in <code>/usr/sbin/vxprint</code> . (Sun Bug I.D. 4277942)
36416	Performance problems with VxVM 3.0.2. (Sun Bug I.D. 4255085)
36454	<code>slib junk</code> must be updated when OS patches are applied.
39097	<code>vxdiskadm</code> allows initialization of bootdisk.
39899	Negative percentage figures when SRL fills.
40201	<code>vxresize</code> fails if non-English language Solaris is installed.
40365	Sequential I/O performance for A5X00 disk arrays with VxVM 3.0.x is worse by 25% than with VxVM 2.5.x. (Sun Bug I.D. 4309793)
40971	(27095,45678) <code>vxvm</code> 3.0.4 giving away free licenses with no SSA/SENA present. (Sun Bug I.D. 4318311 & 4243845)
41858	Can read stale data from an attaching plex which is marked as "preferred:". (Sun Bug I.D. 4278679)
42211/45774	After <code>vxinstall</code> and reboot all controllers might not be seen by DMP. (Sun Bug I.D. 42533763)
44318	<code>vxinstall</code> sees multiple paths to A5000 device. (Sun Bug I.D. 4336678)
42878	<code>vxdiskadm</code> replaces the disk but errors with unknown error while writing the <code>vtoc</code> .
43178	Oracle Error ORA-27061 complains about missing AIOs
43320	DMP should not use stale <code>devinfo</code> pointers when used with DR. (Sun Bug I.D. 4338705)
44662	VMSA does not grow a volume of <code>use_type</code> "gen" (Sun Bug I.D. 4328528)
44712	world has <code>rw</code> on some files under VXVM. (Sun Bug I.D. 4244390)
44745	In <code>vxvrfy</code> the " <code>dg/volume</code> " argument is only 31 total characters. (Sun Bug I.D. 4341777)
45263	Documentation should state that rootability volumes cannot be grown (Sun Bug I.D. 4214541)
28202/46566	AP handling code does not completely work for photons. (Sun Bug I.D. 4332453)

Lists of incidents fixed in previous releases of VxVM can be found in the applicable Release Notes.

Software Limitations and Problems in Volume Manager 3.1

The following problems and issues exist in this release of the Volume Manager.

Firmware Upgrades

Workaround for SUN Bug I.D. 4164338

In order to do Firmware upgrade for disk controllers, there is no need to reboot the system and unload the VxVM in-kernel drivers `vxdmp`, `vxio` and `vxspec` drivers to guarantee data availability during the firmware upgrade procedure to a disk participating in the RAID mirror configuration. The procedure is outlined below.

With DMP

With DMP installed on the system and the configuration is setup as follows:

System is having mirrored volume across the controller.

Dmp is enabled

1. Disable the plex associated with the disk device.

```
#/usr/sbin/vxplex -g dg_name det plex_name
```

2. Stop the IOs to all disks through the specified HBA by executing

```
#/usr/sbin/vxdmpadm disable ctlr=1st_cntlr_name
```

For last controller for this device

```
#/usr/sbin/vxdmpadm -f disable ctlr=2nd_cntlr_name
```

3. Perform the upload of firmware on those disks for which the *ctlr* has been disabled

```
#./download
```

4. After doing upload, enable all the controllers by executing

```
#/usr/sbin/vxdmpadm enable ctlr=2nd_cntlr_name
```

5. Enable the plex associated with the device.

```
#/usr/sbin/vxplex -g dg_name att Volume_name plex_name
```

The above command will take some time depending upon the size of the mirror set.



Without DMP

With NO DMP installed on the system and the configuration is setup to be as follows:

System is having mirrored volume (Not necessary across the controller)

Dmp is not enabled

1. Chose `vxdiskadm` Option 4, Remove a disk for replacement and chose the disk on which firmware download has to be perform.
2. Chose `vxdiskadm` Option 11, offline the device.
3. Download the firmware on to the disk device.
4. Chose `vxdiskadm` Option 10, online the device.
5. Chose `vxdiskadm` Option 5, Replace a failed or removed disk.

FastResync

VERITAS Incident Number	Description
47485	For <code>fsgen</code> type volumes: during I/O, if a user tries to detach a plex (via <code>vxplex det plex-name</code> operation, Fast Resync does not get turned on.

Installation Issues

VERITAS Incident Number	Description
none	<p>Remove a VM disk from a disk group after evacuating any data on the disk. You can permanently remove this disk from Volume Manager control by removing the VM metadata partition from that disk Use the VM low-level command, <code>vxdiskunsetup</code>, as follows:</p> <pre># /usr/lib/vxvm/bin/vxdiskunsetup c#t#d#s2</pre> <p>Note that this command permanently removes a disk from Volume Manager control and should be used with caution only by a system administrator who is trained and knowledgeable in the use Volume Manager.</p>



VERITAS Incident Number	Description
none	Note any of the disks you are planning to initialize that were previously under Volume Manager control. If the disks were previously under Volume Manager control and were used on the same host system, and the proper deinstallation procedures were <i>not</i> followed, the disk groups they represent are imported automatically during the installation process. An attempt during the installation to initialize or encapsulate disks that were previously under Volume Manager control fails. After the installation, if you no longer desire to use those disk groups, use the destroy option of the <code>vxdg (1M)</code> command to remove those disk groups. Alternately, you can use <code>vxdiskunsetup (1M)</code> to remove the disks from Volume Manager control. Be aware that these options can result in data loss if used incorrectly.

Upgrade Issues

VERITAS Incident Number	Description
none	<p>If you have third-party multipathing software configured (such as Sun Enterprise Server Alternate Pathing), insert the following step in the upgrade sections of the <i>VERITAS Volume Manager Installation Guide</i>:</p> <p>After completion of the upgrade <code>start_script</code> and before rebooting the system (for example, by using <code>/etc/shutdown</code>), deconfigure the third-party software multipathing upgrades before continuing with the VERITAS upgrade.</p> <p>Once you have completed deconfiguring the Solaris and third-party multipathing software upgrades, return to the next step in the VERITAS upgrade procedure.</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: tuti10 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>eeprom</code> option <code>use-nvramrc?</code> is set to <code>true</code>, boot the system by specifying <code>vx-diskname</code>.</p>
48041	<p>The <code>/etc/system</code> file contains extra lines with the list of <code>forceload</code> entries before upgrade being different from the list of entries after upgrade. The difference is that there are duplicate entries for <code>set vxio:vol_rootdev_is_volume=1</code> and <code>rootdev:/pseudo/vxio@0:0</code>. Functionality is not affected and the duplicates can safely be removed.</p>

Utility Issues

VERITAS Incident Number	Description
6154	<p>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.</p>
11286	<p>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.</p>



VERITAS Incident Number	Description
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. This problem will be addressed in a future release.</p>
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 data-bbox="657 787 1218 814"># 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 data-bbox="657 919 1055 947"># vxassist mirror volume ...</pre> <pre data-bbox="657 957 1055 984"># vxassist addlog volume ...</pre> <p>This problem will be addressed in a future release.</p>
45668	<p>Due to the current implementation to handle the resize of layered volumes, it is recommended not to <code>grow</code> or <code>shrink</code> layered volumes (stripe-mirror, concat-mirror, and so on) while resynchronization is ongoing.</p> <p>Internally, the Volume Manager 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>relayout</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> <p>Although this release supports layered volumes, creating volumes with mixed layout types is not recommended. For example, adding a mirror using the <code>vxassist mirror</code> command specifying <code>layout=mirror-stripe nmirror=1</code> to an existing volume with stripe-mirror layout is not recommended. (Sun Bug I.D. 4339626)</p>



VERITAS Incident Number	Description
none	<p>While doing <code>relayout</code> 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 1024 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>
none	Resize of a mixed plex volume is currently not supported.

Device Issues

VERITAS Incident Number	Description
none	<p>Disks with insufficient space (less than 1024 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 1024 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 (1024 blocks) for proper disk identification.</p>



VERITAS Incident Number	Description
5316	<p>The Volume Manager tracks disks using long unique identifiers that VxVM 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>

Hot-Relocation Issues

VERITAS Incident Number	Description
14894	<p>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>



Relayout Issues

VERITAS Incident Number	Description
47375	Assuming that all columns fit on a single spindle, then the relayout spindle requirement is: MAXIMUM OF (the greater of the source volume number of columns or the destination volume number of columns) + 2 Note that the two extra spindles must be different from the source volume and destination volume spindles. The two extra spindles are required because more space is needed on two spindles to allocate the <code>tmp plex</code> .

DMP Issues

VERITAS Incident Number	Description
29959	<code>vxdmpadm 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 Reference Guide</i> .

Cluster Functionality Issues

VERITAS Incident Number	Description
none	The new features in Volume Manager Release 3.0.x are not supported in shared disk groups.



VERITAS Incident Number	Description
40055	<p>In VxVM 3.x, the volume layout policy defaults to a layered volume when the requested size for a mirrored volume is equal to or greater than one gigabyte. For example, to create a layered volume, use the following command:</p> <pre data-bbox="699 569 1390 594"># vxassist make volx 1g layout=striped nmirror=2</pre> <p>However, for shared disk groups, layered volumes are not yet supported. For example, if the volume were part of a shared disk group, the command can fail with the following message:</p> <pre data-bbox="699 701 1289 753">vxvm:vxassist: ERROR:Association count is incorrect</pre> <p>For a volume within a shared disk group, to explicitly specify the layout as layered, use the following command:</p> <pre data-bbox="699 835 1357 888"># vxassist make volx 1g layout=stripe-mirror \ nmirror=2</pre> <p>However, the command can fail with the following message:</p> <pre data-bbox="699 940 1403 966">vxvm:vxassist: ERROR: Cannot assign minor number</pre> <p>Therefore, for shared disk groups, when creating a volume of size one gigabyte or greater, specify the layout= mirror-stripe option in vxassist, as shown in the following command:</p> <pre data-bbox="699 1073 1357 1125"># vxassist make volx 1g layout=mirror-stripe \ nmirror=2</pre> <p>Note The default value that vxassist uses to create a layered volume is one gigabyte. However, this is a user configurable parameter. For more info on this, see the vxassist(1M) man page.</p> <p>VMSA will <i>always</i> attempt to create a layered volume when a mirrored volume's size is specified as equal or greater than one gigabyte. When using VMSA to create large mirrored volumes in a shared disk group, you must select the No Layered Volumes option in the Create Volume dialog box. This ensures that the resulting volume has a non-layered layout.</p> <p>If you accidentally create a layered volume, you should convert the volume layout to a non-layered layout (using the Change Volume Layout task) before importing the disk group as shared.</p>
none	The Volume Manager 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 cluster functionality in this release of VxVM has not been qualified on Sun SPARC PCI machines



VERITAS Incident Number	Description
none	<p>If Sun Cluster 2.0 and NETDisk are installed, up to four nodes are supported per cluster. Otherwise, only two nodes are currently supported per cluster.</p> <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, Volume Manager output and GUI 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 Volume Manager cluster functionality with a SPARCStorage Array, you must use firmware level 3.4 or higher.</p>



VERITAS Incident Number	Description
none	<p>Dynamic Multipathing (DMP) enables the Volume Manager to use multiple host-to-disk paths in some multiported disk arrays. DMP enhances reliability by doing path failover in the event of the loss of one or more paths, and increases performance by doing I/O load balancing across multiple I/O paths.</p> <p>VxVM does not currently support the Dynamic Multipathing feature in a shared write access disk environment. It is not possible to configure the disk arrays in a VxVM cluster to have multiple paths from a single host. It is therefore recommended that you disable DMP when VxVM is used in a cluster environment.</p> <p>To disable DMP, use the procedure in the section called “To Disable DMP” on page 23 of the Volume Manager Hardware Notes.</p>
none	<p>In the Sun Cluster, when a disk error occurs on a node, the disk is detached. This is the case even if the other node can access the disks successfully. As a result, an error in the path from a given node to a controller results in the loss of all access to the disks on that controller.</p>
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, re synchronize the volume manually with the following command</p> <pre># vxvol -f resync volume</pre> <p>[Sun #4087612]</p>

Miscellaneous Issues

VERITAS Incident Number	Description
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>
9936	<p>RAID-5 volumes cannot currently be mirrored.</p>



VERITAS Incident Number	Description
none	On machines with low memory (32 megabytes or less), under heavy I/O stress conditions against high memory usage volumes (i.e., 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.
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.
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 volname</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	The way that Volume Manager handles minor numbers for volume devices can cause problems when upgrading from 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 effect disk group versioning.



VERITAS Incident Number	Description
32576	Using an A5x00 array on a PCI bus as an encapsulated root (boot) disk or an alternate root disk is not supported at this time. 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. [Sun #4259045]

Solaris Issues

VERITAS Incident Number	Description
6211	Driver close calls should take priority on SVR4 derivative operating systems. If this is not done, it is possible to receive a simultaneous last-close and first-open operation, swap 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).
none	Since the disk label is stored in block 0 of the disk, block 0 must not be used (i.e., no application should write any information in block 0). Special protection has been built into VxVM to protect block 0 from being overwritten.
6914	Boot disks are to be replaced with disks of similar geometry. When replacing a boot disk using the "Remove a disk for replacement" or "Replace a failed or removed disk" menus from the <code>vxdiskadm</code> utility, ensure that the replacement device has the same disk geometry as the failed boot disk. For disks other than boot disks, this restriction does not apply. A better solution for boot disks is to move all volumes from the boot disk to an alternate disk and to then remove the old disk.
none	The UNIX <code>dd</code> command 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.



VERITAS Incident Number	Description
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>
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>If multiple swap partitions are encapsulated on your disks, the Volume Manager 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:</p> <pre>/dev/vx/dsk/swapvol2 : Overlapping swap files are not allowed</pre> <p>This is due to a problem with the <code>swapadd</code> scripts, which use a faulty <code>swap -l</code> 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 <code>swapvol</code>) to <code>swap1</code>, <code>swap2</code>, and so on, instead of using the current <code>swapvol[0-9]</code> 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 data-bbox="652 657 1356 709">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 Reference 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 data-bbox="652 877 1052 903"># ls -l vxio* vxspec* vxdmp*</pre> <p>The release-specific versions of the kernel modules are stored as <code>module.OS_release</code>, where <code>OS</code> and <code>release</code> are the result of running the <code>uname -s</code> and <code>uname -r</code> commands on the system, respectively.</p> <p>(continued)</p>



VERITAS Incident Number	Description
13312 (continued)	<p data-bbox="651 434 1422 499">For example, on a misconfigured system running Solaris 2.6, the listing is similar to the following:</p> <pre data-bbox="695 506 1422 1493"> -rw-r--r-- 1 root other 339224 Mar 27 07:17 vxdmp -rw-r--r-- 1 root sys 324568 Mar 24 15:46 vxdmp.SunOS_5.6 -rw-r--r-- 1 root sys 327420 Mar 24 16:23 vxdmp.SunOS_5.7 -rw-r--r-- 1 root sys 339224 Mar 24 16:26 vxdmp.SunOS_5.8 -rw-r--r-- 1 root sys 1036 Mar 24 15:46 vxdmp.conf -rw-r--r-- 1 root other 1682424 Mar 27 07:17 vxio -rw-r--r-- 1 root sys 1647664 Mar 24 15:45 vxio.SunOS_5.6 -rw-r--r-- 1 root sys 1661340 Mar 24 16:23 vxio.SunOS_5.7 -rw-r--r-- 1 root sys 1682424 Mar 24 16:26 vxio.SunOS_5.8 -rw-r--r-- 1 root sys 1001 Mar 24 15:45 vxio.conf -rw-r--r-- 1 root other 14928 Mar 24 07:17 vxspec -rw-r--r-- 1 root sys 14252 Mar 24 15:45 vxspec.SunOS_5.6 -rw-r--r-- 1 root sys 14540 Mar 24 16:23 vxspec.SunOS_5.7 -rw-r--r-- 1 root sys 14928 Mar 24 16:26 vxspec.SunOS_5.8 -rw-r--r-- 1 root sys 1325 Mar 24 15:45 vxspec.conf </pre> <p data-bbox="651 1499 1422 1612">Note The size of the kernel modules being used (those without suffixes) match the <i>driver</i>:SunOS_5.8 versions. To correct the problem, copy the SunOS_5.6 versions to the “in-use” module names, as follows:</p> <pre data-bbox="695 1619 1104 1692"> # cp vxio.SunOS_5.6 vxio # cp vxspec.SunOS_5.6 vxspec </pre> <p data-bbox="651 1698 1422 1724">The root file system is then unmounted and the system can be rebooted.</p>



VERITAS Incident Number	Description
13388	<p>During encapsulation, the Volume Manager 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 is <i>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>
24619	<p>To install and test the Solaris-based Volume Manager 3.1 package, you <i>must first</i> install the Sun patches according to the patch matrix shown in Table 1.</p>

Table 1. Solaris Operating System Patch Requirements

Sun Operating System	Sun Array	Sun Patch
Solaris 7	N/A	106541
Solaris 2.6	SSA A5X00	105223 105357

Sun patches are available through the Sun web site sunsolve.sun.com.

Volume Manager 3.x `pkgadd` scripts for the `VRTSvxvm` have been modified to produce an informational message if the required Sun patches are not present on your system.



Software Limitations and Problems in Storage Administrator Release 3.1

The following problems and issues exist in this release of the Volume Manager Storage Administrator:

VERITAS Incident Number	Description
17772	The Volume Manager Storage Administrator does not support destroying deported disk groups. Workaround: Import the disk group, then destroy it.
18338	The Volume Manager Storage Administrator does not allow the user to specify volume usage types.
22456	The maxsize operation always returns sizes in sectors.
22482	The Volume Manager Storage Administrator 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.
22752	On Windows, attempts to print online help are silently ignored if there is no default printer on the system.
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: Warning: Name: scrollbar Class: XmScrollBar The scrollbar page increment is less than 1. This message can be ignored.



VERITAS Incident Number	Description
25089	If you remove a Volume Manager Storage Administrator 1.x package and then install a Storage Administrator 3.x package, the new Storage Administrator server may not start properly. Workaround: Before you remove the Storage Administrator 1.x package, run the <code>server.sh -k</code> utility to stop the 1.x server.
25361	When the window manager's interactive placement setting is turned on, Storage Administrator windows may not be sized properly. The windows can be very small or very large. Workaround: Turn off the interactive placement setting. For example: <code>Mwm*interactivePlacement: False</code>
26269	Menu shortcuts only work when input focus is in the tree or grid. Workaround: To enable shortcuts, select an item in the tree or grid.
27291	Under some circumstances, moving the mouse over a toolbar button results in an exception. If this happens, the status area may not display the correct toolbar button description.
27348	When the Storage Administrator is run with the <code>f_vwm</code> window manager, window decorations appear above the top of the screen.
27349	When the Volume to Disk Mapping window is updated, the contents of the window may become inaccurate. Workaround: Close the Volume to Disk Mapping window and then reopen it.
29621	When disks are added, the order in which the Volume Manager 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.
31029	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 Storage Administrator client.
32052	The Storage Administrator does not support communication between the client and server across a firewall.



VERITAS Incident Number	Description
32599	<p>The Storage Administrator may have trouble connecting to a host machine (server) if multiple host names are associated with a single IP address. The Storage Administrator displays the following message:</p> <pre>Summary:There is no such server (host1) Detail: java.net.UnknownHostException: Unknown host: [host2:32839]; nested exception is: java.net.UnknownHostException: host2</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 LANG=C before executing VMSA. <p>Note VMSA strings are displayed in English.</p>
34293	Displayed time is always GMT when in a Japanese locale.
39544	With Windows 98, the Selected menu may overwrite other menus.



VERITAS Incident Number	Description
40641	With Windows 2000, the Customize window tabs do not always display properly. Workaround: After selecting a tab, resize the window so that it displays correctly.
45734	Sparse plexes of striped volumes, when viewed in Object View, may display incorrectly if the missing subdisk is the last of several subdisks in a column.
46098	VMSA allows the creation of a ufs filesystem with a 4096 blocksize. The Sun4u systems do not support this option. However, VMSA will present this option even for Sun4u systems. Avoid selecting a file system blocksize of 4096 when using Sun4u systems. (Sun Bug ID 4250332)
46077	The following server-side exception can occur when terminating the client: <pre>java.rmi.UnmarshalException: error unmarshalling return header.</pre> Users can ignore this exception.
47481	VMSA uses the default number of columns for the volume change-layout dialog. Please make sure that this number is correctly entered before executing the relayout.
47736	Cancelling the VMSA Connect to Host dialog does not result in cancelling VMSA. It only cancels the current connection. Therefore, VMSA may be run while connected to no hosts.
none	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> .
none	The Volume Manager Storage Administrator does not support statistics and analysis.



VERITAS Incident Number	Description
none	<p>The following X Window System error may occur when the Storage Administrator 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>The Storage Administrator cannot be started from a Web browser. Early releases of Storage Administrator were run from a Web browser. However, due to the nature of Web browsers, running the Storage Administrator from a Web browser was not recommended because it resulted in security constraints, performance problems, and other issues.</p>

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 Volume Manager control through encapsulation. The root disk can be encapsulated either during the `vxinstall` process (while installing VxVM), from the `vxdiskadm` menus (after VxVM is installed), or from the Storage Administrator. Once encapsulated, the root disk can be mirrored using `vxdiskadm`.

Note 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 bootdisk (`rootvol`, `usr`, `var`, `opt`, `swapvol`, etc.) because these map to a physical underlying partition on the disk and must be contiguous. Refer to *Volume Manager Rootability* in Chapter 3 of the *VERITAS Volume Manager Administrator's Guide*.

Booting From DMP Devices

When the root disk is placed under Volume Manager 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.

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).

Note The import lock contains a host ID (with 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. 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.

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 Reference Guide*.

Caution If `vxdg 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 a large number of 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 wide 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 Chapter 2 of the *VERITAS Volume Manager Reference 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 `vxdg` 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.

Available Documentation

The following documents accompany this Volume Manager 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 Reference Guide*
- ◆ *VERITAS Volume Manager Storage Administrator Administrator's Guide*
- ◆ *VERITAS Storage Replicator for Volume Manager Administrator's Guide (3.1)*
- ◆ *VERITAS Storage Replicator for Volume Manager Release Notes (3.1)*
- ◆ *VERITAS Storage Replicator for Volume Manager Configuration Guide (3.1)*
- ◆ Online manual pages

The Storage Administrator provides online help files. 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 VERITAS Volume Manager guides are provided on the CD-ROM under the `pkgs/VRTSvmdoc` directory. If you have installed the `VRTSvmdoc` package, the documents are available in the following locations:

- ◆ *VERITAS Volume Manager Installation Guide*
 - `/opt/VRTSvxvm/docs/installguide.ps`
 - `/opt/VRTSvxvm/docs/installguide.pdf`
- ◆ *VERITAS Volume Manager Reference Guide*
 - `/opt/VRTSvxvm/docs/ref.ps`
 - `/opt/VRTSvxvm/docs/ref.pdf`
- ◆ *VERITAS Volume Manager Administrator's Guide*
 - `/opt/VRTSvxvm/docs/admin.ps`
 - `/opt/VRTSvxvm/docs/admin.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`
- ◆ *VERITAS Volume Replicator Administrator's Guide*
 - `/opt/VRTSvrdoc/docs/vr_ag.pdf`
 - `/opt/VRTSvrdoc/docs/vr_ag.ps`
- ◆ *VERITAS Volume Replicator Configuration Guide*
 - `/opt/VRTSvrdoc/docs/vr_config.pdf`
 - `/opt/VRTSvrdoc/docs/vr_config.ps`



Unformatted manual pages related to the VERITAS Volume Manager are located in the `VRTSvmmman` directory on the CD-ROM. If you have installed the `VRTSvmmman` package, the manual pages can be found in the `/opt/VRTSvxvm/man` directory. The Storage Administrator manual pages are in the `/opt/VRTSvmsa/man` directory. If you add these directories to your `MANPATH` environment variable, you can view these man pages with the `man(1)` command.

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.

