

VERITAS Storage Foundation™ 4.1

Release Notes

Solaris

Maintenance Pack 1

November 2005

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.

VERITAS Legal Notice

Copyright © 2005 VERITAS Software Corporation. All rights reserved. VERITAS, VERITAS Storage Foundation and the VERITAS Logo are trademarks or registered trademarks of VERITAS Software Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

VERITAS Software Corporation
350 Ellis St.
Mountain View, CA 94043
USA
Phone 650-527-8000
Fax 650-527-2908
www.veritas.com

Third-Party Legal Notices

Data Encryption Standard (DES)

Copyright© 1990 Dennis Ferguson. All rights reserved. Commercial use is permitted only if products which are derived from or include this software are made available for purchase and/or use in Canada. Otherwise, redistribution and use in source and binary forms are permitted. Copyright© 1985, 1986, 1987, 1988, 1990 by the Massachusetts Institute of Technology. All Rights Reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting. WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

GNUPlot 3.7

Copyright © 1986-1993, 1998 Thomas Williams, Colin Kelley License: Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided the above copyright notice appear in all copies and both that copyright notice and this permission notice appear in supporting documentation. Permission to modify the software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you 1. Distribute the corresponding source modifications from released version in the form of a patch file along with the binaries, 2. Add special version identification to distinguish your version in addition to the base release version number, 3. Provide your name and address as the primary contact for the support of your modified version, and 4. retain our contact information in regard to use of the base software. Permission to distribute the releases version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions. This software is provided "as is" without express or implied warranty to the extent permitted by applicable law. Authors: Original Software: Thomas Williams, Colin Kelley. Gnuplot 2.0 additions: Russell Lang, Dave Kotz, John Campbell. Gnuplot 3.0 additions: Gershon Elber and many others.



Apache Jakarta Lucene

Apache Jakarta Lucene. Copyright © 2000 The Apache Software Foundation. All rights reserved. This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

Sun Microsystems Trademarks

Sun, Solaris, SunOS, Java, Sun Java System Cluster, Sun StorEdge, Solstice DiskSuite, Sun Fire, Sun Enterprise, Online: Backup, and Netra are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.





Contents

Introduction	1
System Requirements	2
Solaris Operating System Requirements	2
Database Support	2
No Longer Supported	2
Component Product Release Notes	3
Patches	3
VERITAS Patches	3
Solaris Patch Requirements	7
Language Patches	11
Installing the Patches	13
Installing for the First Time	14
Removing Patches	14
Software Limitations	15
VERITAS Volume Manager Software Limitations	15
VERITAS Volume Replicator Software Limitations	16
VERITAS File System Software Limitations	17
VERITAS Storage Foundation for Databases Software Limitations	19
Fixed Issues	24
VERITAS Volume Manager Fixed Issues	24
VERITAS FlashSnap Agent for Symmetrix Fixed Issues	28
VERITAS Volume Replicator Fixed Issues	29
VERITAS File System Fixed Issues	30



VERITAS Storage Foundation for Databases Fixed Issues	31
Open Issues	35
VERITAS Storage Foundation Open Issues	35
VERITAS Volume Manager Software Issues	36
VERITAS FlashSnap Agent for Symmetrix Software Issues	41
VERITAS Volume Replicator Software Issues	42
VERITAS File System Software Issues	43
VERITAS Storage Foundation for Databases Software Issues	48
Documentation Issues	49
Localization Issues	50
Getting Help	50



VERITAS Storage Foundation

Release Notes

Introduction

This document provides release information for VERITAS Storage Foundation 4.1 Maintenance Pack 1 (MP1) for Solaris and includes information for these products:

- ◆ VERITAS Storage Foundation (QuickStart, Standard, Standard HA, Enterprise, and Enterprise HA)
- ◆ VERITAS Volume Manager (VxVM)
- ◆ VERITAS FlashSnap Agent for Symmetrix (VxFAS)
- ◆ VERITAS Volume Replicator (VVR)
- ◆ VERITAS File System (VxFS)
- ◆ VERITAS Storage Foundation *for Oracle* (Standard, Enterprise, and HA Editions)
- ◆ VERITAS Storage Foundation *for DB2* (Standard, Enterprise, and HA Editions)
- ◆ VERITAS Storage Foundation *for Sybase* (Standard, Enterprise, and HA Editions)

Each of these products is activated by a single license key. You must obtain a license key before installing the product. For information on obtaining a license key, see the *VERITAS Storage Foundation Installation Guide*.

Note For the latest information on updates, patches, and software issues regarding this release, see the following TechNote on the VERITAS Technical Support website:
<http://support.veritas.com/docs/272714>.

Review this entire document before installing your VERITAS Storage Foundation product.

Note For product descriptions and other information see the *VERITAS Storage Foundation Release Notes 4.1 for Solaris* document.

Note In this document, reference numbers for issues are included in square brackets, for example [xxxxxx], where xxxxxx is a number. If you contact VERITAS Technical Support about an issue, please provide the reference number where available.



System Requirements

Solaris Operating System Requirements

The VERITAS Storage Foundation 4.1 product line operates on the following Solaris operating systems:

Solaris 8 (32-bit and 64-bit)

Solaris 9 (32-bit and 64-bit)

Solaris 10 (64-bit)

Storage Foundation verifies that the target machine is running a required version of the Solaris operating system. Storage Foundation installation will fail if the product discovers an incorrect Solaris version.

Database Support

VERITAS Storage Foundation 4.1 *for Sybase* now supports Sybase ASE 15.

VERITAS Storage Foundation 4.1 for Oracle now supports Oracle 10gR2.

No Longer Supported

- ◆ Sun Microsystems has announced the End of Support Life for Solaris 2.6. Contact Sun customer support for more information.
- ◆ With VERITAS Storage Foundation 4.1, Solaris 2.7 is no longer supported.
- ◆ With VERITAS Storage Foundation 4.1, Oracle 8i is no longer supported. Any references to Oracle 8i in the *VERITAS Storage Foundation for Oracle Administrator's Guide* should be disregarded.
- ◆ VERITAS Storage Foundation 4.1 is the last release to support Oracle Disk Manager for raw devices.
- ◆ VERITAS Storage Foundation 4.1 is the last release to support the text-based VxDBA menu interface.
- ◆ VERITAS Storage Foundation 4.0 was the last release to support the VERITAS Space Capacity Planning utility for Storage Checkpoints in the GUI. The 4.1 release supports the VERITAS Space Capacity Planning utility with the CLI only.
- ◆ VERITAS Storage Foundation 4.0 was the last release to support scheduling Storage Checkpoint creation.

- ◆ VxFS disk layout versions prior to Version 4 cannot be mounted. Use the `vxfsconvert` command to convert them to a disk layout version that can be mounted. See the *Upgrading VxFS Disk Layout Versions* section in the *VERITAS Storage Foundation 4.1 Installation Guide* for more information.
- ◆ The VxFS QuickLog feature is not supported on disk layout Version 6 nor on Solaris 10. To use the features of disk layout Version 6 and the benefits of the QuickLog feature, see the *Converting From QuickLog to MVS* section of *Chapter 9, Multi-Volume File Systems* in the *VERITAS File System 4.1 Administrator's Guide*. QuickLog will not be supported in the next release.
- ◆ In VERITAS Storage Foundation 4.1 *for Oracle*, snapshot plexes created by the `vxassist` command are not supported. A combination of snapshot plexes created by `vxassist` and `vxsnap` is also not supported.

Component Product Release Notes

Release notes for component products in all versions of the VERITAS Storage Foundation are located under the `storage_foundation/release_notes` directory of the VERITAS Storage Foundation disc. It is important that you read the relevant component product release notes before installing any version of VERITAS Storage Foundation:

VERITAS Storage Foundation Release Notes (`sf_notes.pdf`)

VERITAS Cluster Server Release Notes (`vcs_notes.pdf`)

Because product release notes are not installed by any packages, VERITAS recommends that you copy them to the `/opt/VRTSproduct_name/doc` directory after the product installation so that they are available for future reference.

Patches

VERITAS Patches

The following VERITAS Storage Foundation patches are included with this MP1 release:

Product	Patch ID	Package
VERITAS Enterprise Administrator Service	115209-xx	VRTSob
VERITAS Enterprise Administrator	115210-xx	VRTSobgui



Product	Patch ID	Package
VERITAS Volume Manager	117080-xx	VRTSvxvm
VERITAS File System (Solaris 8)	119300-xx	VRTSvxfs
VERITAS File System Manual Pages	119300-xx	VRTSfsman
VERITAS File System (Solaris 9)	119301-xx	VRTSvxfs
VERITAS File System Manual Pages	119301-xx	VRTSfsman
VERITAS File System (Solaris 10)	119302-xx	VRTSvxfs
VERITAS File System Manual Pages	119302-xx	VRTSfsman
VERITAS Oracle Disk Manager	119303-xx	VRTSodm
VERITAS Oracle Disk Manager	119304-xx	VRTSodm
VERITAS Oracle Disk Manager	119305-xx	VRTSodm
VERITAS File System Management Services Provider	119306-xx	VRTSfspro
VERITAS Device Discovery Layer Services Provider	119306-xx	VRTSddlpr
VERITAS Volume Manager Intelligent Storage Provisioning	119746-xx	VRTSalloc
VERITAS Storage Foundation for Oracle	112014-xx	VRTSdbed
VERITAS Storage Foundation GUI for Oracle	112015-xx	VRTSorgui
VERITAS Storage Foundation for DB2	120116-xx	VRTSdb2ed
VERITAS Storage Foundation GUI for DB2	120117-xx	VRTSd2gui
VERITAS Storage Foundation for Sybase	120118-xx	VRTSsybed
VERITAS FlashSnap Agent for Symmetrix	120143-xx	VRTSfas
VERITAS Volume Manager Management Services Provider	120151-xx	VRTSvmpro
VERITAS Array Providers	120156-xx	VRTSvail

Note xx is the two-numeral extension following the patch number.



Storage Foundation Patches for Databases

You must install the following patches before installing a VERITAS Storage Foundation 4.1 for database product on the indicated operating system versions:

Operating System	Patch ID	Package
Solaris 8	115209-xx	VRTSob
	115210-xx	VRTSobgui
	117080-xx	VRTSvxvm
	119300-xx	VRTSvxfs
	119303-xx	VRTSodm
	119306-xx	VRTSfspro
	119745-xx	VRTSddlpr
	119746-xx	VRTSalloc
	120114-xx	VRTSdbed
	120115-xx	VRTSorgui
	120116-xx	VRTSdb2ed
	120117-xx	VRTSd2gui
	120118-xx	VRTSsybed
	120143-xx	VRTSfas
	120144-xx	VRTSvcs
	120151-xx	VRTSvmpro
	120156-xx	VRTSvail



Operating System	Patch ID	Package
Solaris 9	115209-xx	VRTSob
	115210-xx	VRTSobgui
	117080-xx	VRTSvxvm
	119301-xx	VRTSvxfs
	119304-xx	VRTSodm
	119306-xx	VRTSfspro
	119745-xx	VRTSddlpr
	119746-xx	VRTSalloc
	120114-xx	VRTSdbed
	120115-xx	VRTSorgui
	120116-xx	VRTSdb2ed
	120117-xx	VRTSd2gui
	120118-xx	VRTSsybed
	120143-xx	VRTSfas
	120145-xx	VRTSvcs
	120151-xx	VRTSvmpro
	120156-xx	VRTSvail
Solaris 10	115209-xx	VRTSob
	115210-xx	VRTSobgui
	117080-xx	VRTSvxvm
	119302-xx	VRTSvxfs
	119305-xx	VRTSodm
	119306-xx	VRTSfspro
	119745-xx	VRTSddlpr
	119746-xx	VRTSalloc
	120114-xx	VRTSdbed
	120115-xx	VRTSorgui
	120116-xx	VRTSdb2ed
	120117-xx	VRTSd2gui
	120118-xx	VRTSsybed
	120143-xx	VRTSfas
	120146-xx	VRTSvcs
	120151-xx	VRTSvmpro
	120156-xx	VRTSvail

Note *xx* is the two-numeral extension following the patch number.

In addition, the following VERITAS Cluster Server patches must be installed for the indicated database agents:

- ◆ 120147-*xx* (VCS Oracle Agent)
- ◆ 120161-*xx* (VCS Sybase Agent)

Add the following patches for use in an HA environment:

- ◆ 120871-*xx*
- ◆ 121372-*xx* (Solaris 10 only)

Solaris Patch Requirements

Solaris Patches for VxVM

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 on your system. For more information, see the `showrev(1M)` manual page.

If the patches shown in the required list are not already installed, go to <http://sunsolve.sun.com> to download them. You must 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.

Operating System	Sun Patch Number	Notes
Solaris 10	119254-09 (or later)	<p>Install Sun patch 119254-09 before installing Storage Foundation 4.1 MP1. This same patch is required to remove VERITAS Storage Foundation products installed on Solaris 10 operating systems.</p> <p>An additional Solaris OS patch may be required to completely fix this problem. Solaris patches are available from http://sunsolve.sun.com/</p> <p>To verify the latest information on support for Oracle database versions including 10gR2, see the VERITAS Technical Support TechNote: http://support.veritas.com/docs/272714.</p>



Operating System	Sun Patch Number	Notes
Solaris 9	SUNWscpu package	The SUNWscpu package should be shipped as part of the OS; however, if you do not have this package, it is available from http://www.sun.com/storage/san/index.html SUNWscpu is a package and should be installed using pkgadd.
Solaris 8	SUNWsan package (see your vendor to obtain) SUNWscpu package 111413-06 (or later)	Regardless of whether you have an A5x00 drive connected, you must install these A5x00 patches for VxVM to function. You must install the SUNWsan package and Sun Patch 109529 before installing Sun Patch 111413. You can then install VxVM. This order is important. Note To get the SUNWsan package, go to http://www.sun.com/software/download/allproducts.html#S and then scroll down to "Sun StorEdge SAN 4.4." Click the link to download. You must be a registered user. The download is free. SUNWsan and SUNWscpu are packages and should be installed using pkgadd.
	108528-29 (or later)	Patch 108528-18 requires patches 112396-02 (or later), 108987-09 (or later), 111293-01 (or later), 111310-01 (or later), and 111111-03 (or later). These patches must be installed in the order listed.
	108993-18 (or later)	
	110722-01 (or later)	If you have Sun's Alternate Pathing (AP), you must install Solaris Patch 110722-01 for VxVM 4.1 to co-exist with the AP driver version 2.3.1. However, you must also upgrade AP. See " DMP Coexistence with Alternate Pathing " on page 9 before installing VxVM.

DMP Coexistence with Alternate Pathing

This note applies only if you are using Solaris Operating System 8. VERITAS Volume Manager 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 or later. VxVM version 4.1 does not support earlier versions of AP.

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

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

A5x00 Devices

Whether or not you have A5x00 devices connected to your system, you must add the Solaris patches, shown in the table in “[Solaris Patches for VxVM](#)” on page 7, to this release of VxVM. If you have A5x00 devices attached to your system, you need this patch to use them; if you don't have A5x00 devices attached to your system, VxVM uses the patches to discover whether such devices exist.

Claiming Devices in the SENA Category

The A5x00 disk arrays are claimed by DMP under the SENA category only if the required libraries are present on the system at the time of VxVM installation.

The libraries are present by default on Solaris 2.8 and Solaris 2.9. The patches must be installed before installing VxVM.

The libraries are present by default on Solaris 2.10. No patch installation is necessary.

Solaris Patches for VEA

It is recommended that you install Sun Solaris Package SUNW1of if you plan to use the VERITAS Enterprise Administrator client. SUNW1of is a package required by JRE 1.4. The package contains fonts that are needed to support various locales. For more information, see <http://java.sun.com/j2se/1.4/font-requirements.html>.

Note The name of the Sun package should contain the number one (“1”), not the letter l, after the lower case “i.”



Solaris Patches for VVR

VVR requires the following Solaris patches in addition to the ones listed for VERITAS Volume Manager. The following list gives the minimum revision levels for the required patches:

Solaris 8: 108981-03, 108806-02

VVR will *not* work properly if the following patches are installed on the system:

Solaris 8: 108434-11

Solaris 9: 111711-05

The above patches are withdrawn by Sun Microsystems. Contact Sun Microsystems for more information about these patches and recommended actions.

Do Not Use Solaris 8 Patch 110934-10 or Solaris 9 Patch 113713-01

Solaris 8 patch 110934-10 and Solaris 9 patch 113713-01 prevent the installation of VCS, VxVM, and GLM patches.

By using the `showrev -p` command, you can display the currently installed patches and their levels. For example, to check for patch 110934-10, enter:

```
# showrev -p | grep 110934
```

If you have patch 110934-10 (Solaris 8) or patch 113713-01 (Solaris 9) installed, you must either upgrade or remove it using the `patchrm` command:

```
# patchrm 110934-10
```

If the patch cannot be removed from the system, you will receive the following error message:

```
Patch 110934-10 was installed without backing up the original
files. It cannot be backed out.
```

If you have this problem with Solaris 8 patch 110934-10, refer to TechNote 252441 at <http://support.veritas.com/docs/252441.htm>.

If you have this problem with Solaris 9 patch 113713-01, rename the space file as follows:

```
# mv /var/sadm/pkg/VRTSvxvm/install/space \
/var/sadm/pkg/VRTSvxvm/install/space.org
```

After you remove the patch, you can replace it with the latest patch, using the `patchadd` command, that has been verified by VERITAS.

The following patch levels have been verified:

- ◆ 110934-14 (Solaris 8)
- ◆ 113713-11 (Solaris 9)

To install the latest revision of a patch, use the `patchadd` command as follows:

```
# patchadd 110934-14
```

For Solaris 8, you can use patch 110934-08 or lower if you choose not to upgrade to patch 110934-14. You can successfully install the VERITAS packages without either patch.

If you are installing VERITAS Storage Foundation 4.1 MP1 in a language other than English, you must install the required language patches after installing the English patches. If you are planning to use the GUI, you must install patches for the VEA client.

Language Patches

The following language patches are available for this release.

Product	Patch ID	Package
VERITAS Volume Manager	120584-xx	VRTSjavmc VRTSfrvmc VRTSzhvmc
VERITAS Volume Manager Management Services Provider	120152-xx	VRTSmuvmp
VERITAS Intelligent Storage Provisioning (VRTSalloc)	119747-xx	VRTSmualc
VERITAS File System Management Services Provider	119307-xx	VRTSmufsp
VERITAS File System Maintenance Language Patch	119752-xx	VRTSjafsc
Storage Foundation for DB2 Message Catalog	120121-xx 120122-xx	VRTSjad2g VRTSjadb2
Storage Foundation for Oracle Message Catalog	120123-xx 120124-xx	VRTSjadbe VRTSjaorg
VERITAS Enterprise Administrator	115212-xx 115213-xx	VRTSmuob VRTSmuobg
VERITAS Cluster Server	120153-xx	VRTSjacs, VRTSjacsd, VRTSjacsj, VRTSjacsm, VRTSjacsu, VRTSjacsw



VERITAS Cluster Server DB2 Enterprise Extension	120569-xx	VRTSjacsb
VERITAS Cluster Server Sybase Enterprise Extension	120571-xx	VRTSjacss
VERITAS Cluster Server Oracle Enterprise Extension	120155-xx	VRTSjacso

Note xx is the two-numeral extension following the patch number.

You must have superuser (root) privileges to install the VERITAS software.

▼ **To install the product language packages on the server**

1. Make sure the VEA service is not running.

```
# /etc/init.d/isisd status
Current state of server : RUNNING
```

2. If the VEA service is running, stop it using the `isisd stop` command.

```
# /etc/init.d/isisd stop
```

3. Insert the “Language” disc into the disc drive. If you are using Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.

4. Install the language patches using the `install_vp` command.

```
# cd /cdrom/cdrom0
# ./install_vp
```

5. When prompted, reboot the system. When you reboot the system, the VEA service is automatically restarted.

Installing the Patches

If you install VERITAS Storage Foundation using the installation script, patches are added for you. However, if you choose to install the product using `pkgadd`, you need to manually add the patches after installing.

You must have superuser (`root`) privileges to load the VERITAS software.

▼ To install the patches using `install_vp`

1. Stop the VERITAS Enterprise Administrator (VEA) server before installing patches for `VRTSob`, `VRTSobgui`, `VRTSvmp`, `VRTSfspro` and `VRTSorgui`. See [“VERITAS Patches”](#) on page 3 for details.

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

2. Insert the patch disc into the disc drive. If you are using Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.

3. Install the patches using the `install_vp` command.

```
# /cdrom/cdrom0/install_vp
```

4. Copy all the patches from the `storage_foundation/patches` directory to the desired directory on your system.

```
# cd /working_directory
# cp /cdrom/cdrom0/storage_foundation/patches/* .
```

It is recommended that you install all the patches.

5. Install the patches in the following order:
119302 119300 119301 117080 120144 120145 120146
120114 120115 120116 120117 120118 119306 119737
119738 119739 119303 119304 119305 120148 120149
120150 120147 120143 120156 120120 115209 115210
119745 119746 120151 120161 119735

6. Use the `showrev` command to verify patch installation.

```
# showrev -p | grep patch_number
```

7. Shut down the system.

```
# shutdown -i6 -g0 -y
```

Note If you are installing using `patchadd`, use `patchadd -G` for patches 119735, 119746, 119745, and 120871.



Installing for the First Time

If you are installing VERITAS Storage Foundation for the first time you must first install version 4.1 and then upgrade to MP1.

Removing Patches

▼ To remove an installed patch

1. Log into the machine as root.
2. Remove the patch using the `patchrm` command.

```
# patchrm 119735 -G
```

If you are removing all the patches for the MP1 release, remove them in the following order:

```
119735 120161 120151 119746 119745 115210 115209 120120
120156 120143 120147 120150 120149 120148 119305 119304
119303 119739 119738 119737 119306 120118 120117 120116
120115 120114 120146 120145 120144 117080 119301 119300 119302.
```

Note Use the `patchrm -G` command when removing patches 119735, 119746, 119745, and 120871.

3. Shut down the system.

```
# shutdown -i6 -g0 -y
```

▼ To remove an installed patch from a cluster

1. Log into one of the nodes as superuser.
2. Separate the node from the cluster.

```
# hastop -local
```

3. Stop the `vxconfigd` daemon on that node.

```
# vxdctl stop
```

4. Remove the patch(es).

```
# patchrm -G 119735
```

5. Shut down the node.

```
# shutdown -i6 -g0 -y
```

Note Repeat the steps for each node in the cluster.

Software Limitations

The following sections describe VERITAS Storage Foundation software limitations that exist in this release.

VERITAS Volume Manager Software Limitations

No Support for Local Zones

VERITAS Volume Manager does not support local zones.

vxddladm addsupport Command Limitations

The `vxddladm addsupport` command could cause your system to hang when using a Sun SCSI Enclosure Service (SES) Driver. This situation can be caused by stale entries in the `/dev/es` file. A stale entry is a device link that is present in the `/dev/es` file, but the device itself is not connected to the machine. [i115323, 140441].

In some circumstances, the `pkgadd vxvm` command might cause a system hang because it also executes the `vxddladm addsupport` command.

▼ If the system hangs, perform the following workaround

1. Remove entries from the `/dev/es` file.

You must remove all of the entries because it is not obvious which entries are stale and which are valid.

2. Run the `devfsadm` command as follows:

```
# devfsadm -C
```

This command recreates `/dev/es` with valid entries.



VERITAS Volume Replicator Software Limitations

Support for Local Zones

VVR is supported in the global zone only. Note that although VVR cannot be executed in a non-global zone, VVR can support applications that run in non-global zones, provided the data resides on VxVM volumes in a global zone. For example, if an application is mounted on VERITAS File System in a non-global zone and the file system uses VxVM volumes in a global zone, VVR can be used to replicate such an application's data.

VVR and VxVM commands must be executed in the global zone. If you are replicating an application running in a non-global zone, the global zone administrator requires application context to perform tasks such as snapshots and IBCs.

Unsupported Internet Protocol Versions

VVR does not support IPv6. All nodes in a VVR environment must use IPv4. [135504]

RAID-5 Volume

VVR does not support Volume Manager RAID-5 volumes as part of RVGs. Hardware RAID-5 is supported.

Disk Group Split and Join

VVR does not support Disk Group Split and Join. This means that you cannot use Disk Group Split and Join on data volumes in an RVG. However, you can take snapshots of data volumes and use DGSJ on the snapshots.

Importing a Disk Group with VVR Objects

If a private disk group containing VVR objects is imported on two nodes, the SRL is corrupted and a full resynchronization of the Secondary is required.

Volumes in Boot Disk Group

It is not recommended to have replicated volumes in the boot disk group (`bootdg`).

Length of RLINK Names

VVR does not support RLINK names that exceed 31 characters. If the host name and domain name together exceed 31 characters, the `vxmake rlink` command fails. [151336]

VERITAS File System Software Limitations

Local Zone Support

The following features are supported with local zones:

- ◆ Access to a VxFS file system in the global zone from the local zone of a lofs file system
- ◆ Access to ODM files from local zones
- ◆ Concurrent I/O with files from local zones
- ◆ `ktrace` command in the global zone only
- ◆ `odmstat` command in both global and local zones

The following features are unsupported with local zones:

- ◆ Admin `ioctl`s
- ◆ Administration commands
- ◆ VSM
- ◆ VFS/VxMS
- ◆ Quick I/O and CQIO
- ◆ Cluster File System (CFS)

Enabling ODM Access

The following procedure must be performed to enable ODM (Oracle Disk Manager) access from a local zone.

▼ To enable ODM access from a local zone

1. Install a license in a global zone by exporting the `/etc/vx/licenses/lic` directory to the local zone as an lofs. This is done by adding the following resource to a zone's configuration in addition to other resources in the zone:

```
# zonecfg -z zone1
add fs
set dir=/etc/vx/licenses/lic
set special=/etc/vx/licenses/lic
set type=lofs
end
```



2. Create an `odm` subdirectory under the `zonepath/local_zone/dev` directory. The `zonepath` for a given zone can be obtained using the `zonecfg` command.
3. Mount the `/dev/odm` directory in the local zone.

```
# mount -F odm /dev/odm /dev/odm
```

Note The ODM is not automatically mounted after the zone is booted. Use the above command to mount the ODM after the zone is booted.

Mounting a VxFS File System in the Local Zone

To mount a VxFS file system in the local zone, add the `fs` resource type to the local zone:

```
# zonecfg -z zone1
add fs
set dir=/mnt1
set special=/mnt1
set type=lofs
end
```

The value of *dir* is a directory in the local zone, while the value of *special* is a directory in the global zone to be mounted in the local zone.

Accessing a File for Concurrent I/O

A process can read from or write to a file concurrently with other processes. This implies that a process in a local zone can access the file concurrently with other processes in the local or global zone. An application must perform the following functions:

```
fd=open(filename, oflag)
ioctl(fd, VX_SETCACHE, VX_CONCURRENT)
write(fd, buff, numofbytes)
```


VERITAS Storage Foundation for Databases Software Limitations

Oracle, DB2, and Sybase Software Limitations

No Support for Local Zones

The Standard, Standard HA, Enterprise, and Enterprise HA versions of VERITAS Storage Foundation *for Oracle*, VERITAS Storage Foundation *for DB2*, and VERITAS Storage Foundation *for Sybase* do not support local zones. [268530]

No Support for Intelligent Storage Provisioning

The Standard, Standard HA, Enterprise, and Enterprise HA versions of VERITAS Storage Foundation *for Oracle*, VERITAS Storage Foundation *for DB2*, and VERITAS Storage Foundation *for Sybase* do not support Intelligent Storage Provisioning (ISP).

Disk Layouts Version 5 and Earlier Do Not Display Storage Checkpoint Quotas in the GUI

In VERITAS Storage Foundation *for Oracle* and VERITAS Storage Foundation *for DB2*, VxFS disk layouts Version 5 and earlier do not display Storage Checkpoint quotas in the GUI. [34432 and 34433]

If you attempt to retrieve quota information for a Storage Checkpoint in Version 5 or earlier using the GUI, a message similar to the following displays:

```
DBED4646:ERROR: Unknown error code -30391108 (getquota) for DBED
46464....
```

VERITAS Storage Foundation for Oracle Software Limitations

Storage Checkpoint Limitations

- ◆ You cannot create a clone database using a mounted Storage Checkpoint. [32726]
- ◆ You must run the `dbed_update` command after upgrading to VERITAS Storage Foundation 4.1 *for Oracle* from a previous release. This will allow you to roll back to a Storage Checkpoint that was created prior to this release. [86431]
- ◆ If you create an Oracle instance using the `spfile` option, you must run the `dbed_update` command before you can successfully perform any Storage Checkpoint or Database FlashSnap functions.



Database Cloning Limitation

- ◆ When cloning an Oracle instance using the `dbed_clonedb` or `dbed_vmclonedb` command, the clone database's ORACLE_SID can be only eight characters or less. You will receive an error (ERROR V-81-5713) if the ORACLE_SID is more than eight characters. [345573]

VEA and VxDBA Menu Utility Limitations

- ◆ VERITAS Enterprise Administrator (VEA) and the VxDBA menu utility do not display tablespace information when the `v$table` column names are changed using the SQL*Plus profile facility. [34446]
- ◆ VEA may display system fonts incorrectly. On a Japanese desktop, VEA may incorrectly display system fonts. Japanese characters may not be properly displayed when you select the non-default font for the VEA GUI.

VxDBA Monitoring Agent Limitation

- ◆ The Monitoring Agent fails to start if the setting for the Warning/Grow Threshold for all file systems is less than the actual file system size. (For example, if the file system is 75% full and you set the threshold to 60%.) To work around this limitation, only specify warning or growth thresholds that are greater than your current file system usage amount. [423312]

Database FlashSnap Limitations

- ◆ The Database FlashSnap feature does not support RAID-5 volumes. [34570]
- ◆ When cloning a database using Database FlashSnap, the Oracle database must have at least one mandatory archive destination, otherwise `dbed_vmchecksnap` results in this error message:

```
SFORA dbed_vmchecksnap ERROR V-81-5677 Could not find a
mandatory, primary and valid archive destination for database
PROD.
```

```
Please review the LOG_ARCHIVE_DEST_n parameters and check
v$archive_dest.
```

This example shows how to establish a mandatory archive destination using SQL*Plus:

```
alter system set log_archive_dest_1 =
'LOCATION=/ora_mnt/oracle/oradata/PROD/archivelogs MANDATORY
[REOPEN] [scope=both];'
```

For more information about Oracle parameters for archiving redo logs, see your Oracle documentation. [270905]

Note that the `[scope=both]` syntax shown above does not work on Oracle 10gR2. [423718]

- ◆ After running `dbed_vmsnap -o reverse_resync_commit`, your primary database is started using a pfile. If your original primary database used an spfile, you need to shut down the database and restart it using the spfile. Then, run `dbed_update` to update the repository.

Oracle Disk Manager Limitations

- ◆ If you want to run Oracle 10g on a local zone and use Oracle Disk Manager, the Oracle version should be 10.1.0.3 or higher. To enable Oracle Disk Manager file access from local zones with VERITAS File System, follow the procedure in “[Enabling ODM Access](#)” on page 17.
- ◆ Because Oracle Disk Manager uses the Quick I/O driver to perform asynchronous I/O, do not turn off the Quick I/O mount option, which is the default.
- ◆ Using Oracle Disk Manager with Cached Quick I/O enabled is not supported and could cause your system to panic [34281]. To avoid a system panic, ensure the following:
 - ◆ If you are using Oracle Disk Manager, do not enable Cached Quick I/O on your file system.
 - ◆ If you are converting from Quick I/O to Oracle Disk Manager, make sure you disable Cached Quick I/O.

Oracle Managed Files Limitation

- ◆ Storage Foundation *for Oracle* tools do not support Oracle Managed Files (OMF). [299774]

VERITAS Storage Foundation for DB2 Software Limitations

VEA May Display System Fonts Incorrectly

- ◆ On a Japanese desktop, VEA may incorrectly display system fonts. Japanese characters may not be properly displayed when you select the non-default font for the VEA GUI.

Command Line Interface Limitations

- ◆ The following commands do not support multipartition databases (SMP):
 - ◆ `db2ed_clonedb`



- ◆ db2ed_mon
- ◆ db2ed_vmchecksnap
- ◆ db2ed_vmclonedb
- ◆ db2ed_vxdba
- ◆ The following scripts can be run at the partition level, but not in db2_all. These scripts use mkqio as the input/output filename, so the file will be overwritten by multiple partitions:
 - ◆ qio_getdbfiles
For example: export DB2NODE=1; export DB2DATABASE=mydb3;qio_getdbfiles;
 - ◆ qio_convertdbfiles
For example: export DB2NODE=1; export DB2DATABASE=mydb3;qio_convertdbfiles;

Disabled Monitoring Agent

- ◆ In this release of VERITAS Storage Foundation *for DB2*, the monitoring agent is disabled for multiple partition databases in the GUI.

VEA Server Must Be Restarted After Configuring a Multiple Partition

- ◆ After configuring a multiple partition, you must stop, then restart the VEA server.

Database Cloning Limitations

- ◆ If you clone a database and mount it, ensure that the directory where the mount point resides is owned by the instance owner of the cloned database.

If the directory where the mount point resides is not owned by the instance owner, an error message is displayed when you attempt to remove and unmount the cloned database.

For example:

```
$ db2ed_clonedb -I inst01 -S prod -T clone -c \  
Checkpoint_1105997700 -m /mnt
```

where /mnt is created by root and the owner has been changed to inst01.

When you attempt to remove and unmount the clone database, you will get the following error message:

```
$ db2ed_clonedb -T clone -o umount -d  
rm: Unable to remove directory /mnt Permission denied
```



This error message does not affect the functionality of `db2ed_clonedb`. The clone database has been removed and unmounted even when you receive this error message.

To avoid this error, create a directory under `/` as root and change the owner of the directory to the instance owner. Then, specify a mount point under the newly created directory. For example, instead of using the mount point `/mnt` as in the above example, specify a mount point under `/mnt`, such as `/mnt/clone`:

```
$ db2ed_clonedb -I inst01 -S prod -T clone -c \  
Checkpoint_1105997700 -m /mnt/clone
```

Note If your mount point is under the directory `/tmp`, you will not encounter this problem. [285139]

- ◆ When cloning a DB2 database using the `db2ed_clonedb` or `db2ed_vmclonedb` command, the name of the clone database can be only eight characters or less. You will receive an error (ERROR V-81-5713) if the clone database name is more than eight characters. [345584]



Fixed Issues

The following issues have been fixed since the release of VERITAS Storage Foundation 4.1 for Solaris. To view fixed issues for that release, see the *VERITAS Storage Foundation 4.1 Release Notes for Solaris*.

VERITAS Volume Manager Fixed Issues

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

Incident	Description
152942	<code>libvxvm</code> needs interface to disable/enable linking with global device fs. [Sun Bug ID 4899376]
153010	Unable to boot system in absence of <code>bootdg</code> link to <code>rootdg</code> . [Sun Bug ID 6230224]
157549	<code>cvmstep</code> timeouts on 8 node RAC [Sun Bug ID 4919180]
211810	<code>vx dg split</code> fails after pulling and reconnecting primary path cable.
217093	<code>lun</code> size shrinks after writing <code>cdsdisk</code> label in VM 4.0. [Sun Bug ID 6278993]
225179	<code>vx sat</code> only displays <code>rootdg</code> if <code>-g</code> option was omitted. [Sun Bug ID 6198468]
229538	Multiple enclosures of the same array types generate unexpected names in the form of "xxx0, xxx01, xxx012".
230827	VxVM 4.0 reports two separate A1000 luns as one multipathed device. [Sun Bug ID 5041496]
263569	<code>vx dmp</code> creates the <code>dmp_restore_daemon()</code> with a <code>stksize</code> that is too small causing panic at trap level 2. [Sun Bug ID 6212692]
267062	<code>vxconfigd</code> goes into infinite loop if <code>/etc/vx/array.info</code> file is corrupted.
272263	Changing a disk's I/O policy while the disc is active can cause <code>vxconfigd</code> to hang.
272626	Snapshots for a RAID-5 volume were required to be in multiples of 128 blocks.
272891	Sometimes <code>vxnotify</code> did not receive events relating to the change of controller states in DMP.

Incident	Description
273314	The volume recovery operation will incorrectly use the DRL to recover the volume when the system is restarted.
273850	If you have installed Solaris 10, VxVM cannot boot from an encapsulated DMX 1000 array.
293862	DMP PGR not registering multiple primary keys.
298317	In the case of a failed path, the restore daemon should not be stopped in a cluster and the <code>vxctl enable</code> command is not needed.
299680	<code>vxconfigd</code> core dumps in <code>devintf_dm_reassoc_da</code> on slave.
300098	System panics in <code>voliodone_fun</code> . [Sun Bug ID 6229878]
300979	S-VOL Devices did not show in the <code>vxdisk list</code> output on HDS with TrueCopy enabled.
301931, 308975	Failures were observed on the Hitachi 9990 (Genesis) arrays where the disk geometry data was handled incorrectly by <code>vxconfigd</code> , resulting in error message during <code>vxctl enable</code> or <code>vxconfigd</code> startup. This message does not affect VxVM use of the array. [Sun Bug ID 6221005]
302750	CVM: Point-in-time of the snapshot is not as expected.
303078	VxVM: Support to enable caching of WRITES in kernel caused due to UFS logging. [Sun Bug ID 5105913]
303216	Solaris disk label failed for <code>/dev/vx/rdmp/GENESIS_0</code> (err 22) [Sun Bug ID 6221005]
304344	VxVM/CVM volumes within shared disk groups become unresponsive while resyncing. [Sun Bug ID 6226804]
304618	No swapvol causes <code>vxbootsetup</code> to not partition some volumes. [Sun Bug ID 5063043]
304618	<code>vxbootsetup</code> assumes that first plex of swapvol matches rootvol. [Sun Bug ID 5063799]
304662	Both <code>vxtranslog</code> and <code>vxcmdlog</code> commands write to <code>stderr</code> , not <code>stdout</code> . [Sun Bug ID 6182640]



Incident	Description
310784	Encapsulating boot file system that spans two disks does not complete. [Sun Bug ID 6230678]
311530	<code>vxconfigd</code> dumps core if a user changes an enclosure name to a new name which is longer than the current one. [Sun Bug ID 6230031]
311531	<code>vxconfigd</code> dumped core during an enclosure name change. [Sun Bug ID 6230031]
311664	<code>vxconfigd/dmp</code> hang due to a problem in <code>dmp_reconfig_update_cur_pri</code> function's logic.
312026	CVM: <code>kmsg</code> sender and receiver threads should be interlocked for access of outbound queue for first sends.
312561	4.1- "disk not found" message shows during encapsulation via <code>vxdiskadm</code> . [Sun Bug ID 6230678]
313521	When rootdisk spans multiple disks only one is encapsulated successfully. Only / is added to <code>bootdg</code> . [Sun Bug ID 6234544]
314860	Changes to VxVM upgrade script to convert PP metadevices from foreign/define records to auto-configured records.
316096	Reconnection to <code>vxconfigd</code> failed: Configuration daemon is not accessible. [Sun Bug ID 6212772]
316168	CVM reconfig timeout caused cluster nodes to panic. [Sun Bug ID 5097042]
318405	CVM reconfig timeout caused cluster nodes to panic.
318531	<code>vxedvtoc</code> unknown error happens and system cannot boot from replaced boot disk. [Sun Bug ID 6246560]
325018	VxVM with encapsulated root volume and UFS logging enable in S9u7 fails to boot. [Sun Bug ID 5105913]
325018	Persistent warning messages after encapsulating root logging disk with VxVM. [Sun Bug ID 5045578]
330394	If a volume is mounted as <code>bootdg</code> , the <code>vxresize</code> command fails.[Sun Bug ID 6281539]
336192	Panic in <code>voldr1_unlog</code> when I/Os were happening against <code>cachevol</code> and <code>vxcache stat</code> command is issued in parallel.

Incident	Description
336852	Panic in voldco_needupdate_instant.
338867	vxplex att causes core-dump with ECOPY enabled.
340532	Reconnection to vxconfigd failed: Configuration daemon is not accessible.
340590	PGR strange behavior with VERITAS Storage Foundation standard license. [Sun Bug ID 6220218]
343878	Qual, HDS 95xx with v3 ASL: write i/o errors on subdisks in A/P array transition (NODEVICE on good devices).
344626	DMP PGR not registering keys on newly activated secondary paths during failover.
344873	DMP failover - use temporary state (NODE_SUSPECT) for all LUNs on failed path.
345491	vxprint with -F & e option shows the STATE=SNAPDONE/DCOSNP after the disk is removed.
348104	System crash creating a space-opt snapshot.
354268	Need to backout UFS workaround fix now that permanent VxVM support enabling of cache writes due to UFS logging.
361347	vxdisk resize hangs.
361950	Conflict between global and local failover.
361952	Very slow I/O with space-optimized volume snapshot using 16K region size.
365919	Encapsulation fails in VM 4.1 with A/P luns having multiple primary/secondary paths.
370277	vxdiskadm option 5 errors: disk.repl: =__disk: not found [Sun Bug ID 6290314]
374171	DMP: gen_select_aa_path() should choose not NODE_QUIESCED path if available.
374270	vxprint -Vh coredumps on DCO object.
374427	Checkin snappoint support for 4.1 MP1 release.
375680	VxMS: Differential mapping support for VM plugin.



Incident	Description
380403	DDL: Lun serial no. extraction for EMC arrays with page 83 support broken in 4.1.
389127	I/O fails during NDU process on Cx600.
398147	VxVM patch install does not copy the APM files.
398247	DMP: Memory leak observed in <code>dmpgen_da.c</code> .
414352	VM: Not able to boot from the replaced boot device.
419371	CVM I/O quiesce time during <code>vxclust</code> step 4 too long.
425012	<code>vxsparecheck</code> command doesn't work in Japanese locale OS.
425281	cvm/RAC- panic after long loop failing doing <code>ucmm_reconf</code> step 3 in campus cluster. [Sun Bug ID 5049370]

VERITAS FlashSnap Agent for Symmetrix Fixed Issues

The following issues have been fixed in this release of VxFAS:

Incident	Description
375590	VxFAS: Support RAID-5 STD devices in VxFAS timefinder operations
375592	Qualify VxFAS on Solaris 10 for 4.1MP1
405228	Support RAID-5 BCVs in Timefinder operations



VERITAS Volume Replicator Fixed Issues

The following issues have been fixed in this release of VVR:

Incident	Description
278838	When replicating in TCP mode, the default Nagle algorithm results in very slow replication because it attempts to coalesce small messages from the Secondary, which delays their transmission.
279096	After a disk failure on the Secondary, the <code>secondary_log_err</code> flag on the Primary was not cleared even after the disk on the Secondary was fixed.
296372	Under certain conditions, I/O could hang when issued during a checkpoint attach.
299421	An attempt to attach an Rlink too soon after configuring the replication IP address fails, because VVR does not attempt a sufficient number of retries.
304305	When replicating using TCP, the algorithm for choosing the number of connections has been enhanced.
308169	There is potential for a panic when an SRL I/O error occurs during a pending transaction that is waiting for I/Os to finish.
310204	You cannot resize the SRL while automatic synchronization or DCM replay is in progress. Also, the SRL resize operation does not preserve checkpoint information.
312769	For disk group version 90, if the SRL is disassociated and associated again before the Secondary assumes the role of Primary, the <code>srl_config</code> table is corrupted or not initialized.
316409	Under certain conditions, VVR could hang when connecting to the remote host during a transaction.
319992	In a shared disk group environment, VVR on the master can hang if another host joins the cluster during automatic synchronization.
370768	In certain cases, with a high volume of concurrent writes, replication can stop because VVR on the Primary sends too many messages that are not in the order expected on the Secondary.
383627	The <code>vxrlink stats</code> command incorrectly reports certain errors as timeout errors but actually they should be reported as memory errors.



VERITAS File System Fixed Issues

The following issues have been fixed in this release of VxFS:

Incident	Description
i147918	There was a potential race condition with NFS clients caused by nfs3_fhtovp and a forced unmount, which could result in a panic.
i147939	Unnecessary limit checks were being performed, degrading performance.
i151890	VxFS incorrectly updated the vmstat pi counter on Solaris 10.
i152189	If multiple VxFS threads tried to modify the value of ml_flag at the same time, the value of ml_flag might not have been correctly modified.
i152278	The DNLC cookie was changed from 96-bit to 128-bit to allow 16 bits instead of 8 bits for the CPUID.
i152379	Intermittently, file permissions were seen with all zeroes from a secondary node.
i152455	Performance degraded in the write codepath of VxFS 4.0 compared to VxFS 3.5.
i152770	There was a panic in voldrl_unlog due to opening the device with OTYP_CHR.
i152809	When a large number of attribute inodes were allocated, the system could have panicked due to over-consumed kernel memory.
i152830	Improved vxfsckd sleeping time.
i153001	A data region with zeroes was seen in a CFS shared file in a mix of mmap I/Os with MAP_PRIVATE and POSIX I/Os from different cluster nodes.
i153375	bmap routines for IORG_NONE inodes now return an error.
i153701	A bug in VxFS bmap code paths could cause file system operations to hang.
i153702	When a file system was full, truncation of file operations could fail due to ENOSPC conditions, leading to the file being marked bad and the file system being set up for a full fsck.
i153719	There was an intent log version mismatch issue with BCV backup.
i154159	There was a panic in vx_dqswapfs() when upgrading from disk layout Version 5 to Version 6.

Incident	Description
i154416	Storage Checkpoints did not report all files.

VERITAS Storage Foundation for Databases Fixed Issues

VERITAS Storage Foundation for Oracle Fixed Issues

The following issues have been fixed in this release of VERITAS Storage Foundation *for Oracle*:

Incident	Description
295399	The VxDBA monitoring agent will fail to obtain Oracle tablespace information if the home directory of the Oracle VxDBA user is different than the actual value of ORACLE_HOME.
322948	When running <code>dbed_update</code> after mounting an VxFS file system to <code>/etc/vx/vxdba/\$ORACLE_SID</code> in an HA environment, the following error message would be shown: <code>/etc/vx/vxdba/SID1/checkpoint.dir: No such file or directory</code> Yet, the <code>checkpoint_dir</code> directory was created, and subsequent <code>dbed_update</code> commands did not produce the error message. This has been fixed.
323251	<code>dbed_vmclonedb</code> no longer requires an archivelog mount point to exist when snapshot archivelog is not selected.
324339	When creating or validating a snapplan with <code>dbed_vmchecksnap</code> , the command previously failed to detect snapshot plexes created by the <code>vxassist</code> command. This has been fixed. (Note that <code>vxassist</code> snapshot plexes are no longer supported, even though they are now detected by <code>dbed_vmchecksnap</code> .)
331509	Previously, man pages displayed copyright information on top of each page. This information is no longer shown.
344855	Previously, Storage Foundation <i>for Oracle</i> 4.1 did not support Solaris Liveupgrade. To use Liveupgrade, contact VERITAS support to obtain a customized package for Solaris Liveupgrade support for your OS version.
354652	Previously, <code>vxsvc</code> core dumps when the Oracle SID is greater than 8 bytes. This has been fixed.



Incident	Description
354746	Changed column aliases in <code>login.sql</code> confuse DBED scripts, leading to wrong results.
361996	Database FlashSnap previously failed during resync and/or reverse_resync operations when one of the mirrors of multi-snapshot is disabled. This is now fixed.
376728	Database FlashSnap previously failed with error "5679: Could not obtain database information" when the database had <code>nls_numeric_character = ' , . '</code> This has been fixed.

VERITAS Storage Foundation for DB2 Fixed Issues

The following are fixed issues in this release of VERITAS Storage Foundation *for DB2*:

Incident	Description
315325	The <code>qio_getdbfiles</code> command, which extracts information on Quick I/O files used by the database and stores the names of these files in <code>mkqio.dat</code> , now works with EEE databases. The <code>qio_getdbfiles_all</code> command has been added to work with every database partition for a partitioned DB2 database. See “DB2 EEE Support for Quick I/O Commands” on page 33 for more information.
344858	Previously, Storage Foundation <i>for DB2</i> 4.1 did not support Solaris Liveupgrade. To use Liveupgrade, contact VERITAS support to obtain a customized package for “Solaris live upgrade support for same OS version.”
354636	<p>Previously, <code>db2ed_clonedb</code> and <code>db2ed_ckptrollback</code> worked for DB2 7.2 and for DB2 8.1 up to Fixpack 6, but did not work for DB2 8.2 or above, due to new DB2 configuration files that were introduced.</p> <p>When <code>db2ed_clonedb</code> runs against a version 8.2 DB2 database, it aborts with the following error message:</p> <pre>ERROR V-81-7211 Relocate failed</pre> <p>When <code>db2ed_ckptrollback</code> runs against an 8.2 DB2 database, <code>db2ed_ckptrollback</code> executes successfully, but the configuration data is not correctly restored. This may cause problems rolling the database forward. This problem has been fixed.</p>

VERITAS Storage Foundation for Sybase Fixed Issue

The following is a fixed issue in this release of VERITAS Storage Foundation *for Sybase*:

Incident	Description
320940	There should be a helpful error message when attempting to get a list of data files that reside on raw devices with the <code>qio_getdbfiles</code> command, even though this command should not work on raw device files. Such an error message is now produced.

DB2 EEE Support for Quick I/O Commands

Previously, the `qio_getdbfiles` command did not support partitioned databases. To correct this, a new command (`qio_getdbfiles_db2_all`) has been added. The `qio_getdbfiles_db2_all` command creates a file `mkqio.dat` for each node under the specified directory.

Usage Notes

- ◆ Note that before invoking `qio_getdbfiles_db2_all`, `DB2INSTANCE` and `DB2DATABASE` environment variables must be set.
- ◆ For a non-partitioned database, you can use the `qio_getdbfiles` or `qio_getdbfiles_db2` commands the same as before.
- ◆ If `qio_getdbfiles_db2_all` is invoked with no options, a directory will be created for each node under the current directory and a `mkqio.dat` file will be created underneath each `./NODEnum` directory.

For example, if your database has 3 partitions (0, 1, and 2), the `qio_getdbfiles_db2_all` command will create the files `./NODE0000/mkqio.dat`, `./NODE0001/mkqio.dat`, and `./NODE0002/mkqio.dat`.

- ◆ If `qio_getdbfiles_db2_all` is invoked with the `-d` option, the `mkqio.dat` files will be created in the directory that you specify.

For example, if your database has 3 partitions (0, 1, and 2), the `qio_getdbfiles_db2_all -d /tmp/proddb` command will create the files `/tmp/proddb/NODE0000/mkqio.dat`, `/tmp/proddb/NODE0001/mkqio.dat`, and `/tmp/proddb/NODE0002/mkqio.dat`.



Note The `-d <directory>` option has also been added for the `qio_getdbfiles` and `qio_getdbfiles_db2` commands so that you can specify the directory where the `mkqio.dat` file resides.

Note There is no change in the `qio_convertdbfiles` or `qio_recreate` command syntax. You can use the `qio_convertdbfiles` or `qio_recreate` command as usual.

Converting Database Files to Quick I/O Files and Recreating Quick I/O Files for Partitioned Databases

Once you have created the `mkqio.dat` files for each partition using the `qio_getdbfiles_db2_all` command, there are two ways to convert and recreate the files:

1. Consolidate all `mkqio.dat` files into one `mkqio.dat` file, put the consolidated `mkqio.dat` file in your current directory, then execute the `qio_convertdbfiles` or `qio_recreate` command.

This method only works if all your database files for all nodes are on the same machine.

2. You can change directory (`cd`) to each `directory/NODenum` directory and execute the `qio_convertdbfiles` or `qio_recreate` command for each node separately.

Open Issues

VERITAS Storage Foundation Open Issues

License Package Not Completely Removed From Local Zones

Some files from the VERITAS licensing package (VRTSvlic) may not be removed from a local zone that was created after VRTSvlic was originally installed. An error message is displayed if all files are not removed. If the error was encountered, after the package removal process ends, run the following command from the global zone to remove any remaining VRTSvlic files:

```
# rm -rf zonepath/root/var/sadm/pkg/VRTSvlic
```

If you are upgrading a product, and local zones are configured, instead of selecting the installer upgrade procedure (or running the script from the command line), perform the upgrade in steps: Uninstall the product, uninstall the infrastructure packages (uninstallinfr script), then reinstall the product.

Removing Patch From a Clustered Node

Patch 117080-03 can cause a core dump when the patch is removed from a node that is part of a cluster. [422513]

▼ To remove patch 117080-03 from a clustered node:

1. Log into the node as superuser.
2. Separate the node from the cluster.

```
# hastop -local
```
3. Stop the vxconfigd daemon on that node.

```
# vxdctl stop
```
4. Remove the patch.

```
# patchrm 117080-03.
```
5. Shut down the node.

```
# shutdown -i6 -g0 -y
```

Note Repeat the steps for each node in the cluster.



VERITAS Volume Manager Software Issues

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

Increase UDP Buffer Size to Reduce Long I/O Quiesce Time

In a Sun Cluster environment when a node joins a cluster originally comprised only of a master node, a long I/O quiesce time may be observed. In this case to reduce the long I/O quiesce time, tune the UDP receive buffer size from 8K to 64K by doing the following:

```
# ndd -set /dev/udp udp_rcv_hiwat 65536
```

[429461] [Sun Bug ID 6331482]

Enclosure Type FAKE_ARRAY May be Displayed

Initially, the boot disk is marked as FAKE array type. If the boot disk was not excluded from VxVM then the boot disk will migrate to the correct enclosure type depending on discovery and the FAKE array type is deleted from it. In the case where boot disk is excluded from VxVM, the disk does not migrate. The FAKE array still has a disk, it can not be deleted and so the FAKE_ARRAY enclosure type status is still visible to the user. This is by design.

In this circumstance, when a user runs the command `vxddmpadm listenclosure all` the root disk shows as "FAKE_ARRAY" enclosure type.

```
# vxddmpadm listenclosure all
```

ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS	ARRAY_TYPE
=====				
FAKE_ARRAY0	FAKE_ARRAY	FAKE_ENCLR_SNO	DISCONNECTED	-
Disk	Disk	DISKS	CONNECTED	Disk
HDS99100	HDS9910	EC05	CONNECTED	A/A
T30	T3	60020f200000c2f80000	CONNECTED	A/PF-T3PLUS
GENESIS0	GENESIS	02740	CONNECTED	A/A

[492268, 492269] [Sun Bug ID 6179424]

Repaired Device Path Failures Before Membership Change

After repairing a failed device path, a user must run the `vxddctl enable` command prior to a planned or unplanned membership change in the cluster, or the previous bad state of the storage paths will persist. If the node that has had the path failures becomes the new CVM master, the same paths may fail on all nodes, even though the path has been repaired.

[Sun Bug ID 6346291] [499260]



Inconsistencies in Storage Object State Reporting

When path failures occur, commands that show the state of VERITAS storage objects (e.g., `vxprint`, `vxdisk list`, etc.) will show the state from the local node perspective rather than the global clusterwide state. This perspective can result in the state being reported differently, depending upon which node the command is executed.

CVM Tunables for Sun Cluster

The `vol_kmsg_send_period` and `vol_kmsg_resend_period` tunables are measured in seconds. To accommodate better reconfiguration times for `vxclust step4`, finer resolution of the send and resend periods was required. To accomplish this, two new tuning variables have been introduced that have micro-second resolution, with new default values.

These new variables and defaults are:

```
vol_kmsg_send_period_usec 10
vol_kmsg_resend_period_usec 100
```

With this scheme there will be two sets of variables possible to tune each value in the `vxio.conf` file. If none of them have been specified, the new finer grained variables and defaults will be used. If only the 1 second granularity tunables are specified in the `vxio.conf` file they will be used. If the new fine-grained values only are specified, then they will be used. If both are specified a warning will be reported, and the coarse variables will be used.

▼ To modify a tunable

1. Log in as superuser.
2. Go to the `/kernel/drv/vxio.conf` file.


```
# cd /kernel/drv/vxio.conf
```
3. Modify the file with the following line:


```
vol_kmsg_resend_period=1;
```
4. Reboot the system.

Note It is important to include the semicolon (;) in Step 3.

[419371, 419372]



Replacing an Encapsulated Root Disk

The following is a procedure to replace any disk in the `rootdg` which is mirrored.

▼ To replace any mirrored disk in the `rootdg`

1. Enable mirroring.
2. Replace the mirror disk.
3. Remove the root mirror for replacement.

```
# vxdbg -g myrootdg -k rmdisk myrootdg02
```
4. Verify the volumes in the `myrootdg`.

```
# vxprint -g myrootdg
```
5. Make the removed root mirror offline and remove from the VM config.

```
# vxdisk offline <c0t1d0>
# vxdisk rm <c0t1d0>
```
6. Set the default disk layout on the removed disk

```
# prtvtoc /dev/rdisk/<c0t1d0s2> | grep -v ^* | tee /tmp/vtoc
```
7. Edit the `/tmp/vtoc` and remove all the slices except the slice 2 definition.

```
# fmthard -s /tmp/vtoc /dev/rdisk/<c0t1d0s2>
# prtvtoc /dev/rdisk/<c0t1d0s2> | grep -v ^*
```
8. Get back the removed disk in to VM config.

```
# vxdisk list
# vxdctl enable
# vxdisk list
# vxdisksetup <c0t0d0> format=sliced
# prtvtoc /dev/rdisk/<c0t1d0s2> | grep -v ^*
```
9. Add the disk to `myrootdg` replacing the earlier removed disk and recover the mirrored disk.

```
# vxdbg -g myrootdg -k adddisk myrootdg02=<c0t0d0>
# vxprint -g myrootdg
# vxrecover -v -g myrootdg myrootdg02
# prtvtoc /dev/rdisk/<c0t1d0s2> | grep -v ^*
# vxbootsetup -g myrootdg myrootdg02
# prtvtoc /dev/rdisk/<c0t1d0s2> | grep -v ^*
```

```
# vxprint -g myrootdg
```

The mirrored disk is now replaced.

10. Boot with the replaced disk.

```
# init 0
ok setenv boot-device vx-myrootdg02 vx-myrootdg01
ok boot
```

11. Login after boot and continue the process of replacing the main root disk and boot with the newly replaced disk.

```
# vxdisk list
```

12. Remove the main root disk for replacement. Remove the disk from VM config and set default layout on the disk.

```
# vxdg -g myrootdg -k rmdisk myrootdg01
# vxdisk list
# vxdisk offline <c0t0d0>
# vxdisk rm <c0t0d0>
# vxdisk list
# prtvtoc /dev/rdisk/<c0t0d0s2> | grep -v ^* | tee /tmp/vtoc
```

13. Edit the file to remove all the slice definitions except the slice 2.

14. Set the default layout on the disk

```
# fmthard -s /tmp/vtoc /dev/rdisk/<c0t0d0s2>
# prtvtoc /dev/rdisk/<c0t1d0s2> | grep -v ^*
# vxdisk list
```

15. Bring back the removed disk to VCM config

```
# vxdctl enable
# vxdisk list
# vxdisksetup -i <c0t0d0> format=sliced
# prtvtoc /dev/rdisk/<c0t0d0s2> | grep -v ^*
# vxdisk list
```

Verify if the root disk was created using the entire disk space. In this case the VM while encapsulating creates a subdisk for Block 0 (myrootdg01-B0) to protect zeroth block. That block has to be removed when a new disk is replacing the myrootdg01 failing which the vxbootsetup utility would fail. An option "-p" in the vxdg command would remove the subdisk.

```
# vxprint -g myrootdg
# vxdg -g myrootdg -pk adddisk myrootdg01=<c0t0d0>
# vxprint -g myrootdg
```



```
# vxedit -g myrootdg -rf rm myrootdg01Priv
# vxrecover -v -g myrootdg myrootdg01
# vxprint -g myrootdg
# prtvtoc /dev/rdisk/<c0t0d0s2> | grep -v ^*
# vxbootsetup -g myrootdg myrootdg01
# init 0
ok setenv boot-device vx-myrootdg01 vx-myrootdg02
ok boot
```

[Sun Bug ID 6277202] [336387]

Using MPxIO and DMP

In this mode, for single node/host (non-cluster) the Sun StorEdge Traffic Manager (SSTM) driver or DMP can be used to handle multipathing.

In multi-host configurations, the Sun StorEdge T3 must be configured to the MPxIO mode (explicit failover or nonauto_trespass). In this mode, the Sun StorEdge Traffic Manager Software (STMS) driver on the host handles multipathing. VERITAS Volume Manager (VxVM) will not see multiple paths to any device on the array.

Users must upgrade the Sun StorEdge T3 firmware to 1.17B. The patch for this upgrade is 109115-09. The latest firmware for StorEdge T3 and T3+ Arrays is available in patches 109115-xx and 112276-xx.

For Solaris versions 8 or 9, the `/kernel/drv/scsi_vhci.conf` file must have the MPxIO mode enabled (**`mpxio-disable="no"`**). For Solaris 10, the `/kernel/drv/fp.conf` file must have **`mpxio-disable="no"`**. Use the command `sys mp_support=mpxio` on the Sun StorEdge T3 array to enable the MPxIO mode.

Note All hosts in the clustered environment must use the MPxIO mode.

[Sun Bug ID 6266369] [389407] [492435]

Incorrect Stale Records Displayed

In some cases when enclosure based naming is used, the result of a `vxdbg free` command shows incorrect stale records of deleted disks after the devices have been disabled with `luxadm offline` operation followed by `vxddctl enable`.

In this case use `vxconfigd -k` instead of `vxddctl enable`.

[Sun Bug ID 5107437] [496130]

Master Node Fails to Import Disk Group

In a clustered environment if a master node fails over and the second node fails, a third node takes over as master. If the new master does not automatically import the disk group, the user must manually use the `vxdctl enable` command on the new master node to import the disk group. If `vxdctl enable` doesn't import the disk group then the user must import the disk group using the `vxdg` command. [Sun Bug ID 5049370] [496350]

Special Note for UFS Logging

This release supports UFS logging on a VxVM encapsulated root disk. If you have previously disabled UFS logging (by setting the "nologging" option in `/etc/vfstab`), you may now restore UFS logging after installing this release.

To change the nologging setting, edit the `/etc/vfstab` file (and any other boot volumes that are created by VxVM, for example, `"/", "/usr", "/var", "/opt"`) and remove "nologging" to the mount option of those volumes, as in the following example:

```
/dev/vx/dsk/bootdg/rootvol /dev/vx/rdisk/bootdg/rootvol /ufs 1 no -
/dev/vx/dsk/bootdg/usr /dev/vx/rdisk/bootdg/usr /usr ufs 1 no -
/dev/vx/dsk/bootdg/var /dev/vx/rdisk/bootdg/var /var ufs 1 no -
/dev/vx/dsk/bootdg/opt /dev/vx/rdisk/bootdg/opt /opt ufs 2 yes -
```

After you have changed the nologging setting, reboot the system. [325018, 303078, 259150]

VERITAS FlashSnap Agent for Symmetrix Software Issues

The following are known issues in this release of VERITAS FlashSnap Agent for Symmetrix.

When devices are configured in fabric naming schemes (WWN names), VxFAS operations are only supported if the VxVM naming scheme is enclosure based.

When devices are configured in fabric naming schemes (WWN names), VxVM changes the default device names to enclosure based names without actually converting the naming scheme into enclosure based. This type of configuration is not supported by VxFAS. If you are using the devices configured in WWN names you should manually switch the naming scheme to enclosure based. [412155]

VxFAS CLI May Fail With TPD Naming

When an EMC Powerpath (EMCpower) package is installed, VxVM uses the TPD (third party driver) naming convention by default. Sometimes the VxFAS CLI (command line interface) may fail with these names after running `vxdisk init` or after `vxsysm split` because of a correlation issue. The user must run the `vxdctl enable` command to make the correlation complete.



We recommend that users configure the EMC Symmetrix devices with enclosure based naming instead of TPD. The VxFAS CLI does not have a conflict with enclosure based naming. Users can change from TPD to enclosure based naming using the VxVM command `vxdiskadm option 20`. The user may also configure EMC Symmetrix devices as foreign to avoid a VxFAS CLI conflict. [412155]

VERITAS Volume Replicator Software Issues

The following are known issues in this release of VERITAS Volume Replicator:

RLINKs Fail to Connect After Logowner Failover

In a shared environment, replicating over a firewall, RLINKs may fail to connect after the logowner fails over to another host, because the VVR heartbeats continue from the original logowner host.

On both the original logowner host and the host to which it is heartbeating, stop and restart `vxnetd`, using the following steps.

1. Stop `vxnetd`:

```
# vxstart_vvr stop
```

2. Restart `vxnetd`:

```
# vxstart_vvr start
```

[346443]

VVR Hangs During SRL Resize

Under certain conditions, VVR may hang while executing the `vradm resize` command. Workaround: Reboot the host. [497131]

VVR Hangs After Logowner Move

In a shared environment, if a logowner has been moved repeatedly between the slave and master, VVR may hang. Workaround: Reboot both hosts. [496922]

VxVM ISP Stripe-mirror Volumes Cannot be Associated to an RVG

You cannot associate VxVM ISP volumes of the type stripe-mirror to an RVG. [497136]

VERITAS File System Software Issues

The following are known issues in this release of VERITAS File System:

API for Manipulating Disk Quotas

VxFS now implements the quota Application Program Interface (API) documented in the Solaris `quotactl(7I)` manual page. Users who have written their own quota tools based on the `Q_QUOTACTL ioctl` can now use those tools on VxFS file systems. However, you cannot administer VxFS file system quotas using the `Q_QUOTACTL ioctl` from a client which mounts VxFS over NFS. This capability will not be available until a modification to the RPC quota daemon (enabling quotas on file systems other than UFS) is implemented on the Solaris operating system.

Stack Size Change

When installed on Solaris 8, Solaris 9, and Solaris 10, VxFS changes the default stack size to 24K for 64-bit systems. In 32-bit mode, VxFS can operate with a stack size of 16K. The stack size is designated in the Solaris configuration file `/etc/system`.

Storage Checkpoints Do Not Operate With HSM Products

Storage Checkpoints cannot be created on a file system where the VERITAS Storage Migrator™ is active, or with other hierarchical storage management (HSM) products that use the DMAPI interface. [451132]

VxFS Incompatible With Some HSM Applications

VxFS does not operate with VERITAS Storage Migrator versions 4.5 and earlier. A patch for VERITAS Storage Migrator 4.5 is available from VERITAS support on the VERITAS Customer Support website:

<http://support.veritas.com/docs/258566.htm>

Other HSM applications may also require a patch. Contact your HSM vendor for product-specific information.

The `ustat` Command Returns an Error for VxFS File Systems Larger than One Terabyte

The `ustat` command returns an `EOVERFLOW` error for VxFS file systems larger than one terabyte because the variable used to store file system size overflows. See the `ustat(2)` manual page.



Commands Must be Large-File Aware to Operate Correctly on File Systems Larger than One Terabyte

For utilities to operate correctly on large-file systems, they must be large file aware. This applies even if commands are invoked on small files in a large file system. See the information regarding disk layout in the *VERITAS File System Administrator's Guide*.

Inode Limitation on File Systems Without Large File Support

For a file system to have more than 8 million inodes, you must create it using the `largefiles` option of `mkfs` (the `fsadm` utility can also be used to set the `largefiles` flag on the file system). See the `mkfs_vxfs(1M)` and `fsadm_vxfs(1M)` manual pages for details. The `largefiles` option is enabled by default on VxFS 4.1. In previous VxFS releases, `nolargefiles` was the default mount option.

Large Files Should Be Mounted Only on Systems With Sufficient Memory

When a file system is mounted, VxFS keeps certain data structures in the kernel. As the size of the file system increases, the amount of data structures stored by VxFS also increases. The file system typically keeps approximately 128 bytes per allocation unit (32,768 file system blocks). This translates to a usage of 512K per 1 TB for an 8K block size file system (4 MB per 1 TB for a 1K block size file system). Therefore, large file systems must be mounted only on systems that have sufficient memory. The memory requirements for mounting large file systems are shown in the tables below.

Memory Usage for a File System With a 1K Block Size

File System Size	128 GB	1 TB	8 TB	64 TB	256 TB
Memory Usage	1 MB	4 MB	32 MB	N/A	N/A

Memory Usage for a File System With a 2K Block Size

File System Size	128 GB	1 TB	8 TB	64 TB	256 TB
Memory Usage	512K	2 MB	16 MB	128 MB	N/A

Memory Usage for a File System With a 4K Block Size

File System Size	128 GB	1 TB	8 TB	64 TB	256 TB
Memory Usage	256K	1 MB	8 MB	64 MB	N/A

Memory Usage for a File System With an 8K Block Size

File System Size	128 GB	1 TB	8 TB	64 TB	256 TB
Memory Usage	128K	512K	4 MB	32 MB	128 MB

While performing a full `fsck`, the system keeps certain data structures in the core for validating the space usage and inode usage. The space needed depends on the number of inodes and the number of blocks in the file system. The `fsck` command needs approximately 16 MB per 1 TB for an 8K block size file system (128 MB per 1 TB for a 1K block size file system) and 32 MB per million inodes. Sufficient memory and swap space should be configured on the system before running a full `fsck` on a large file-enabled system. If the system is booted through a 32-bit kernel, a full `fsck` of file systems that have a large number of blocks or large number of inodes may fail, as the total address space available for a 32-bit process is limited.

A replay `fsck` does not need a significant amount of memory and does not have these issues.

Quick I/O Files Cannot Be Sparse Files

If you try to convert a sparse file to a Quick I/O file, the Oracle instance can fail if Oracle tries to write into an unallocated block. Specifically, datafiles used by the Oracle8i and Oracle9i temporary tablespace may be sparse files, so do not convert these to Quick I/O files. See the *VERITAS Storage Foundation 4.1 for Oracle Database Administrator's Guide* for more information.

Some Disk Quota Operations Do Not Function on NFS

When VxFS file systems are exported via NFS, quotas on the file system apply to users when accessing the file system from NFS clients. However, neither the Solaris nor the VxFS quota commands on the NFS client can be used to query or edit quotas. The VxFS quota commands can be used on the server to query or edit quotas.

fscdstask validate Error With ja_JP.UTF-8-Encoded File Names

The `fscdstask validate` command returns an error when files on the specified mount point have names with the `jp_JP.UTF-8` encoding, but the locale has been changed to `ja_JP.eucJP` or `ja_JP.PCK`. The error is as follows:

```
xargs: Input file is corrupt. : Incorrect byte order
```

Files should be created with the same locale encoding as the file system on which they reside.



Non-Standard Command Behavior When Using Access Control Lists

The output of the `ls -l` command on VxFS file systems shows mask/CLASS_OBJ in place of group permissions if ACLs are in use on a file or a directory. You can determine the effective group permissions by using the command.

The `chmod` command changes mask/CLASS_OBJ instead of the group permissions if ACLs are in use on a file or a directory. GROUP_OBJ is not changed by `chmod`, and because effective group permissions are determined by GROUP_OBJ and CLASS_OBJ, the default group may not receive the permissions specified by `chmod`. Because `ls -l` shows mask only (which is changed by `chmod`), it only appears that the group permissions are changed as specified in `chmod`. On files with ACLs, use the command to manipulate permissions. See the following manual pages for ACL-related information: `chmod(1)`, `ls(1)`, and `umask(1)`.

Files and Directories

To maximize VxFS performance, do not exceed 100,000 files in the same directory. Use multiple directories instead.

100% Full File System Cannot Be Resized

In some circumstances, the `fsadm` and `fsvoladm` commands cannot resize a 100% full file system due to lack of space for updating structural information. Check VxFS file systems on a regular basis and increase their size if they approach 100% capacity. This problem can also occur if the file system is very busy. Free up space or reduce activity on the file system and try the resize again.

Data Integrity Issues With Disks and Disk Arrays With Write-Back Caches

Disk drives configured to use a write-back cache, or disk arrays configured with a volatile write-back cache, can 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.

Disable QuickLog Device Logging Before Upgrading to Disk Layout Version 5

Because of the VxFS Version 5 disk layout change, you must disable QuickLog logging on any file systems mounted with the `mount qllog=` option before upgrading from disk layout Version 4. See the *VERITAS Storage Foundation Installation Guide* for information on upgrading from older disk layout versions.

Note QuickLog does not operate on the Version 6 disk layout used by VxFS 4.0 and 4.1.

JumpStart Enterprise Toolkit Not Supported

The JumpStart Enterprise Toolkit is not supported in this release, but will be supported in a future release.

DTrace Warnings May Display on First Boot After Installation

On the Solaris 10 operating system, DTrace warnings may display when the system is booted for the first time after VxFS is installed. The warnings are similar to the following:

```
Configuring devices.  
Hostname: MyHost.MyCompany.com  
WARNING: couldn't allocate SDT table for module vxfs  
.  
.  
.  
WARNING: couldn't allocate SDT table for module vxfs  
WARNING: couldn't allocate FBT table for module vxfs  
Loading smf(5) service descriptions: 2/2
```

These warnings indicate that the SDT and FBT DTrace probes may not be available for the `vxfs` module until the next reboot. The `vxfs` module will still load and work correctly.

These warnings do not display on subsequent reboots.



VERITAS Storage Foundation for Databases Software Issues

VERITAS Storage Foundation for Oracle Software Issues

The following are known issues in this release of VERITAS Storage Foundation *for Oracle*:

Incident	Description
none	ORAMAP-FMON might not start in Oracle 10gR2. To work around this issue, add the following entry to <code>init.ora</code> and try to start the ORAMAP: <code>_filemap_dir=\$ORACLE_HOME/rdbms/filemap</code>
301174	The column heading TBSNAME in the output for the <code>dbed_analyzer</code> command is not localized. If you are running in an environment other than English, note that TBSNAME means “Tablespace Name.”
303238	When installing the VRTSdbed package using JumpStart, you may see the following warning: <code>installing <a/opt/VRTS/man/man1m/qio_convertdbfiles.1m></code> <code>with default mode of 644</code> It is safe to ignore this warning. The permission for this file is correctly set as 644.
421816	The following table headings produced by <code>dbed_analyzer</code> are not aligned with the data shown below them: DATAFILE, DEVICE, SIZE(sectors), and TBSNAME. This can lead to data being aligned under the wrong heading. Also, if there is a Japanese tablespace and the tablespace name exceeds nine characters, then the data line for the DEVICE column will be shifted about five characters.
424436	In the Japanese version of VERITAS Storage Foundation <i>for Oracle</i> , if an invalid Japanese language tablespace name is given as input to the <code>dbed_analyzer</code> command (with option <code>-t</code>), the resulting error message (ERROR V-81-5908) includes only the first character of the invalid Japanese tablespace name instead of its entire name.
430072	In the Japanese version of VERITAS Storage Foundation <i>for Oracle</i> , if the NLS setting is incorrect, when trying to start a database via the GUI, an error window displays the message “Can't find the error message” and the GUI does not show the database information.
497035	In the Japanese version of VERITAS Storage Foundation <i>for Oracle</i> , database information under the Storage Rollback Administration menu is misaligned.

VERITAS Storage Foundation for DB2 Software Issue

The following is a known issue in this release of VERITAS Storage Foundation *for DB2*:

Incident	Description
417505	Executing an offhost <code>db2ed_vmclonedb</code> command with the <code>-o umount</code> option occasionally produces a Segmentation Fault message for both online_snapshot and offline modes. This message can be ignored. The cloning operation has successfully completed.

Documentation Issues

- ◆ In the *Storage Foundation for Oracle 4.1 Database Administrators Guide*, on page 229 under the section “Establishing a Mandatory Archive,” the SQL*Plus example misplaces the second single-quote character. This single-quote should be located at the end of the command sequence, as shown below:


```
alter system set log_archive_dest_1 =
'LOCATION=/ora_mnt/oracle/oradata/PROD/archivelogs MANDATORY
[REOPEN] [scope=both];'
```
- ◆ In the *Storage Foundation for Oracle 4.1 Database Administrators Guide*, on page 204 under the section “To Shut Down the Clone Database and Unmount All Snapshot File Systems,” the `db2ed_vmclonedb` command used in the example is missing the `-D DB2DATABASE` option in its syntax. The correct command example should be as follows:


```
$ /opt/VRTS/bin/db2ed_vmclonedb -D DB2DATABASE -o \
umount,new_db=db_name -f SNAPPLAN [-r relocate_path]
```
- ◆ In the *Storage Foundation for DB2 4.1 Database Administrators Guide*, on page 181, in the second example for two-host configuration, the `db2ed_vmchecksnap` command is shown without a `-p` option preceding the `PRODTAG`. The correct command syntax is as follows:


```
$ /opt/VRTS/bin/db2ed_vmchecksnap -D PROD -f snap2 -o setdefaults \
-p PRODTAG -t host2
```
- ◆ In the *Storage Foundation 4.1 Installation Guide*, the language package `VRTSjafad` is not listed.
- ◆ In the *Storage Foundation 4.1 Installation Guide*, the package `VRTSmuobg` is incorrectly identified as being for “Windows client only.”



- ◆ The *VERITAS Volume Manager Hardware Notes* contains errors about MPxIO and DMP. See “[Using MPxIO and DMP](#)” on page 40. [Sun Bug ID 6266369] [389407] [492435]
- ◆ In the *VERITAS Volume Manager 4.1 Administrator’s Guide* a reference in the “Effect of Disk Connectivity” section on page 375 refers the reader to the same section. The reference should be to the “Connectivity Policy of Shared Disk Groups” section on page 371. [Sun Bug ID 6331461] [429475, 429898, 429899, 429934, 492920]

Localization Issues

In Japanese locales, the VAIL configuration fields for configuring Out of Band arrays in the VEA GUI are too small to enter data. To work around this, use the `vaildiag` utility to configure these arrays instead of the GUI. This issue applies to all database products. [277495]

Getting Help

For technical assistance, visit <http://support.veritas.com> and select phone or email support. This site also provides access to resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and the VERITAS customer email notification service. Use the Knowledge Base Search feature to access additional product information, including current and past releases of product documentation.

Diagnostic tools are also available to assist in troubleshooting problems associated with the product. These tools are available on disc or can be downloaded from the VERITAS FTP site. See the `README.VRTSspt` file in the `/support` directory for details.

For license information, software updates and sales contacts, visit <https://my.veritas.com/productcenter/ContactVeritas.jsp>. For information on purchasing product documentation, visit <http://webstore.veritas.com>.