

VERITAS Storage Foundation™ 4.1

Release Notes

Solaris

March 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

Storage Foundation Product Suites	2
Product and Feature Descriptions	6
VERITAS Volume Manager	6
VERITAS File System	8
VERITAS Storage Foundation for Databases	9
VERITAS Volume Replicator	10
VERITAS Cluster Server	10
New and Enhanced Features	11
VERITAS Storage Foundation	11
VERITAS Volume Manager	11
VERITAS File System	14
VERITAS Storage Foundation for Databases	16
System Requirements	22
Solaris Operating System Requirements	22
Solaris Patch Requirements	22
VERITAS Patches	27
VxVM and VEA Patches	27
Software Limitations	28
VERITAS Volume Manager Software Limitations	28
VERITAS File System Software Limitations	29
VERITAS Storage Foundation for Databases Software Limitations	31



Fixed Issues	36
VERITAS Volume Manager Fixed Issues	36
VERITAS File System Fixed Issues	38
VERITAS Storage Foundation for Databases Fixed Issues	39
Software Issues	47
VERITAS Storage Foundation	47
VERITAS Volume Manager Software Issues	47
VERITAS File System Software Issues	76
VERITAS Storage Foundation for Databases Software Issues	81
No Longer Supported	88
Available Documentation	88
VERITAS Documentation Disc	89
Release Notes and Installation Guides	89
VERITAS Storage Foundation Guides	89
Manual Pages	92
Getting Help	93



VERITAS Storage Foundation Release Notes

This document provides release information about the products in the VERITAS Storage Foundation 4.1 Solaris product line:

- ◆ VERITAS Storage Foundation (QuickStart, Standard, Standard HA, Enterprise, and Enterprise HA)
- ◆ 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)
- ◆ VERITAS Volume Manager (VxVM)
- ◆ VERITAS File System (VxFS)

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.

This document does not contain release information for VERITAS Volume Replicator (VVR) or VERITAS Cluster Server (VCS). For release information on these products, see the *VERITAS Volume Replicator Release Notes* and the *VERITAS Cluster Server Release Notes*.

Note Throughout 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.



Storage Foundation Product Suites

The following table lists the VERITAS products and optionally licensed features available with each Storage Foundation product suite.

Storage Foundation Version	Products and Features
Storage Foundation QuickStart	VERITAS Volume Manager (Base feature set) VERITAS File System (Base feature set)
Storage Foundation Standard	VERITAS Volume Manager VERITAS File System Optionally licensed features: VERITAS Volume Replicator Option
Storage Foundation Standard HA	VERITAS Volume Manager VERITAS File System VERITAS Cluster Server Optionally licensed features: VERITAS Volume Replicator Option
Storage Foundation Enterprise	VERITAS Volume Manager VERITAS File System VERITAS FlashSnap Option Optionally licensed features: VERITAS Volume Replicator Option
Storage Foundation Enterprise HA	VERITAS Volume Manager VERITAS File System VERITAS Cluster Server VERITAS FlashSnap Option Optionally licensed features: VERITAS Volume Replicator Option



Storage Foundation Version	Products and Features
Storage Foundation <i>for Oracle</i> Standard	VERITAS Storage Foundation <i>for Oracle</i> VERITAS Volume Manager VERITAS File System VERITAS Quick I/O Option VERITAS Extension for Oracle Disk Manager Option
Storage Foundation <i>for Oracle</i> Enterprise	VERITAS Storage Foundation <i>for Oracle</i> VERITAS Volume Manager VERITAS File System VERITAS FlashSnap Option VERITAS Quick I/O Option VERITAS Extension for Oracle Disk Manager Option VERITAS Storage Checkpoint Option VERITAS Storage Mapping Option Optionally licensed features: VERITAS Volume Replicator Option
Storage Foundation <i>for Oracle</i> Enterprise HA	VERITAS Storage Foundation <i>for Oracle</i> VERITAS Volume Manager VERITAS File System VERITAS Cluster Server VERITAS FlashSnap Option VERITAS Quick I/O Option VERITAS Extension for Oracle Disk Manager Option VERITAS Storage Checkpoint Option VERITAS Storage Mapping Option Optionally licensed features: VERITAS Volume Replicator Option
Storage Foundation <i>for DB2</i> Standard	VERITAS Volume Manager VERITAS Storage Foundation <i>for DB2</i> VERITAS File System VERITAS Quick I/O Option



Storage Foundation Version	Products and Features
Storage Foundation <i>for DB2</i> Enterprise	VERITAS Storage Foundation <i>for DB2</i> VERITAS Volume Manager VERITAS File System VERITAS FlashSnap Option VERITAS Quick I/O Option VERITAS Storage Checkpoint Option VERITAS Storage Mapping Option Optionally licensed features: VERITAS Volume Replicator Option
Storage Foundation <i>for DB2</i> Enterprise HA	VERITAS Storage Foundation <i>for DB2</i> VERITAS Volume Manager VERITAS File System VERITAS Cluster Server VERITAS FlashSnap Option VERITAS Quick I/O Option VERITAS Storage Checkpoint Option VERITAS Storage Mapping Option Optionally licensed features: VERITAS Volume Replicator Option
Storage Foundation <i>for Sybase</i> Standard	VERITAS Storage Foundation <i>for Sybase</i> VERITAS Volume Manager VERITAS File System VERITAS Quick I/O Option
Storage Foundation <i>for Sybase</i> Enterprise	VERITAS Storage Foundation <i>for Sybase</i> VERITAS Volume Manager VERITAS File System VERITAS Quick I/O Option VERITAS FastResync Option Optionally licensed features: VERITAS Volume Replicator Option



Storage Foundation Version	Products and Features
Storage Foundation <i>for Sybase</i> Enterprise HA	VERITAS Storage Foundation <i>for Sybase</i> VERITAS Volume Manager VERITAS File System VERITAS Cluster Server VERITAS Quick I/O Option VERITAS FastResync Option Optionally licensed features: VERITAS Volume Replicator Option



Product and Feature Descriptions

VERITAS Volume Manager

VERITAS Volume Manager is a storage management tool that removes the physical limitations of disk storage so that you can configure, share, manage, and optimize storage I/O performance online without interrupting data availability. VxVM also provides easy-to-use, online storage management tools to reduce planned and unplanned downtime.

VxVM 4.1 Licenses

The following table shows the levels of licensing in VERITAS Volume Manager 4.1 and the features supported at each level:

VxVM License	Description of Supported Features
Full	Concatenation, spanning, rootability, volume resizing, multiple disk groups, co-existence with native volume manager, striping, mirroring, DRL logging for mirrors, striping plus mirroring, mirroring plus striping, RAID-5, RAID-5 logging, Smartsync, hot sparing, hot-relocation, online data migration, online relayout, volume snapshots, Intelligent Storage Provisioning, FastResync with Instant Snapshots, Storage Expert, Device discovery Layer, multipath DMP, and VEA.
Add-on Licenses	Features that augment the Full VxVM license such as FlashSnap™ (FastResync and Dynamic Disk Group Split and Join).

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

To see the license features that are enabled in VxVM 4.1, enter the command:

```
# vxdctl license
```



CDS Licensing

The ability to import a CDS disk group on a platform that is different from the platform on which the disk group was last imported is controlled by a CDS license. CDS licenses are included as part of the VERITAS Storage Foundation license.

An example of a CDS enabling license (output by the `vxlicrep` command) is:

```

License Key                = XXXX-XXXX-XXXX-XXXX-XXXX-XXXX
Product Name               = VERITAS Storage Foundation
Standard
License Type              = DEMO_EXTENSION
OEM ID                    = 4095
Demo End Date             = Fri Mar 18 00:00:00 2005
                          (21.5 days from now).
Editions Product          = YES

Features :=
VxVM#VERITAS Volume Manager    = Enabled
CPU Count#VERITAS Volume Manager = Not Restricted
VXFS#VERITAS File System       = Enabled
CPU Count#VERITAS File System   = Not Restricted
QLOG#VERITAS File System       = Enabled
PGR#VERITAS Volume Manager     = Enabled
VERITAS Foundation Suite       = Enabled

SPC Lite#VERITAS SANPoint Control = Enabled
Storage Expert#VERITAS Volume Manager = Enabled
Platform                     = un-used
Version                      = un-used

Product Name                = VERITAS Volume Manager
License Type                = DEMO
Demo End Date               =

Features :=
Dynamic Lun Expansion        = Enabled
Hardware assisted copy       = Enabled
Cross-platform Data Sharing  = Enabled
PGR                          = Enabled
PGR_TRAINING                 = Enabled
VVS_CONFIG                   = Enabled
[...]
```



Photon Licensing

VxVM automatically generates a Photon license when it detects an A5k array. In addition, VxVM 4.1 allows you to remove your Photon device for servicing for up to two weeks. To do this, when the current license is due to expire, VxVM automatically removes the old license, and re-issues a new Photon license with another two-week expiration date.

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

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

```
License Key =
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XX
Product Name = VERITAS Volume Manager
License Type = DEMO_NODE_LOCK
OEM ID = 4090
Demo End Date = Fri Aug 30 01:00:00 2002 <=====
EXPIRATION DATE
(13.6 days from now).
Node Lock Type = (Hostid and Architecture ID)
Features :=
PHOTON                      = Enabled
CPU Count                   = Not Restricted
```

Note This message does not mean that your Photon array feature license will soon expire permanently. You will *not* lose the VxVM features associated with the license in 13.6 days time.

VERITAS File System

VERITAS File System provides high performance and online management capabilities to facilitate the creation and maintenance of file systems. File systems are a collection of directories organized into a structure that enable you to locate and store files.

The primary purposes of a file system are to:

- ◆ Provide shared access to data storage.
- ◆ Provide structured access to data.
- ◆ Control access to data.
- ◆ Provide a common, portable application interface.
- ◆ Enable the manageability of data storage.

VERITAS Storage Foundation for Databases

VERITAS Storage Foundation *for Oracle*, VERITAS Storage Foundation *for DB2*, and VERITAS Storage Foundation *for Sybase* combine the strengths of the core VERITAS products with database-specific enhancements to offer unrivaled performance, availability, and manageability for databases. The database feature options are as follows.

VERITAS FlashSnap Option

The VERITAS FlashSnap option, when applied to databases, lets you create, resynchronize, and reverse resynchronize volume snapshots for databases. The snapshots can be used on a second host. Also, database administrators can perform these tasks without root privileges.

This option is available with the Enterprise versions of VERITAS Storage Foundation *for Oracle* and VERITAS Storage Foundation *for DB2*.

VERITAS Quick I/O Option

VERITAS Quick I/O is a VERITAS File System feature that improves the throughput for databases built on VERITAS File Systems. Quick I/O delivers raw device performance to databases run on VxFS, providing the administrative advantages of using file systems without performance penalties. VERITAS Cached Quick I/O further enhances database performance by leveraging large system memory to selectively buffer the frequently accessed data.

This option is available with both the Standard and Enterprise versions of VERITAS Storage Foundation *for Oracle*, VERITAS Storage Foundation *for DB2*, and VERITAS Storage Foundation *for Sybase*.

VERITAS Extension for Oracle Disk Manager Option

VERITAS Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle9i and 10g. Oracle Disk Manager improves the performance and manageability of system bandwidth through an improved API that contains advanced kernel support for file I/O.

This option is available with both the Standard and Enterprise versions of VERITAS Storage Foundation *for Oracle*.



VERITAS Storage Checkpoint Option

VERITAS Storage Checkpoint technology lets you create a point-in-time image of a file system. Storage Checkpoints are treated like any other VxFS file system and can be created, mounted, unmounted, and removed.

This option is available with the Enterprise versions of VERITAS Storage Foundation *for Oracle* and VERITAS Storage Foundation *for DB2*.

VERITAS Storage Mapping Option

Storage Mapping lets you map datafiles to physical devices and display storage object I/O statistics. Both storage object I/O statistics and the storage structure can be displayed for a specific file.

This option is available with the Enterprise versions of VERITAS Storage Foundation *for Oracle* and VERITAS Storage Foundation *for DB2*.

In addition, with VERITAS Storage Foundation *for Oracle*, mapping information showing which tablespaces reside on which physical disks can be obtained for a specified database.

VERITAS Volume Replicator

VERITAS Volume Replicator is data-replication software designed to contribute to an effective disaster recovery plan by maintaining an exact or consistent copy of application data at one or more remote locations.

VERITAS Cluster Server

VERITAS Cluster Server provides an open systems clustering solution that eliminates both planned and unplanned downtime, facilitates server consolidation and failover, and effectively manages a wide range of applications in heterogeneous environments.

New and Enhanced Features

The following new features and enhancements have been incorporated into VERITAS Storage Foundation, VERITAS Volume Manager, VERITAS File System, VERITAS Storage Foundation *for Oracle*, VERITAS Storage Foundation *for DB2*, and VERITAS Storage Foundation *for Sybase*.

Note For information about new features in VVR, see the *VERITAS Volume Replicator Release Notes*. For information about new features in VCS, see the *VERITAS Cluster Server Release Notes*.

VERITAS Storage Foundation

The following new features and enhancements are included with all five Storage Foundation product suites:

- ◆ Support for Solaris 10
- ◆ JumpStart compliance

VERITAS Volume Manager

This release of VERITAS Volume Manager includes the following new features and enhancements.

Bootting from Fabric Devices

Disks that are connected to a Storage area Network (SAN) are now bootable. In this release, you have the ability to encapsulate a disk that is in a SAN environment. VxVM should theoretically be able to encapsulate any device in a SAN environment; however, only those arrays that have been formally tested by VERITAS will be supported. Refer to the Hardware Compatibility list for the supported arrays.

Enhancements to the VxVM Upgrade Procedures

The upgrade procedures for VxVM have been simplified by:

- ◆ Reducing the number of reboots by using the Solaris OS Live upgrade feature.
- ◆ Reducing the amount of downtime when upgrading a system with an encapsulated root disk.
- ◆ Reducing the number of upgrade paths.



Enhancements to Intelligent Storage Provisioning (ISP)

ISP has been enhanced with the following features:

- ◆ Support for VVR to replicate intent-preserved volumes
VVR requires that DCM logs be allocated and associated with a volume. This release of VxVM provides support for DCM logs so that VVR can replicate intent-preserved volumes.
- ◆ Support for traditional Dirty Region Logging
ISP was designed to support Dirty Region logging (DRL) that was based on FMR-3 DRL. ISP has now been enhanced to support traditional DRL that will provide better scaling.
- ◆ Migrating `vxassist` volumes to ISP
Volumes created with `vxassist` can now be migrated for use with ISP. After migration, volumes can no longer be used with `vxassist`.
- ◆ Enhancements to the ISP language
The ISP language has been enhanced so that you can state a preferred order of rules. To take advantage of this feature, you should upgrade your Disk Group version to 120.
- ◆ Simpler syntax rules
`vxvoladm` supports a simpler syntax for specifying some of the rules.

Extensible Firmware Interface (EFI) Support

The Solaris 9 and 10 64-bit kernel Operating Systems provide support for disks larger than 1 terabyte. Disks of this size are formatted with the EFI disk label rather than the VTOC disk label. A number of commands and libraries have been modified to provide this support. EFI disks are supported on Solaris 10 only. [Sun Bug ID 6226760, 303294]

Coexistence of VERITAS DMP with Third Party Drivers

With this release, the generic framework within the Device Discovery Layer (DDL) enables VERITAS DMP to coexist with other third party multi-pathing drivers such as EMC PowerPath.

Enhancements to DDL-DMP

There are a number of enhancements in the area of DDL-DMP:

- ◆ CVM DMP changes

Failback in CVM is now a cluster-wide operation which provides significant performance improvements.

A/PF (active/passive failover) type arrays are now supported by DMP in a CVM environment.

- ◆ Persistent DMP Policies

It has been possible to change dynamically DMP I/O policies, but these changes were not persistent after a reboot. This release of VxVM solves this problem.

- ◆ Enhancements to the naming scheme

Previous releases of VxVM supported two naming schemes; OS-Native Scheme and Enclosure-based Naming. Although VxVM provided the means of changing between the two schemes dynamically, it was necessary to restart `vxconfigd` which is a time-consuming operation. With this release, the `vxconfigd` restart operation is no longer necessary.

VxVM now supports persistent device names. The disk names, once assigned, will remain constant across reconfiguration and rebooting.

- ◆ Performance Enhancements

Device discovery is now multi-threaded. Performance improvements have also been made to DMP.

Enhancement of Local Detach Policy

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



VERITAS File System

This release of VERITAS File System includes the following new features and enhancements.

Concurrent I/O Interface

Support is now available for `VX_CONCURRENT` caching advisory for a file. When the advisory is set, a write system call on the file takes the inode read/write lock in shared mode instead of exclusive. Also the file system avoids caching the file data for these writes in the page cache. The required concurrency control is performed by the application.

The behavior of the write system call is changed only on the file descriptor on which the `VX_CONCURRENT` advisory is set. Write system calls to the same file through other file descriptors take an exclusive inode read/write lock and cache the pages. For information on setting the `VX_CONCURRENT` advisory, see the *VERITAS File System Administrator's Guide*.

This interface is licensed and is enabled through the QIO license.

Support for Solaris 10 Including Fine-Grain Permissions and Global Zones

VxFS 4.1 supports Solaris 10 and the accompanying fine-grain permissions. Local zones are also supported.

Support for Local Zones

VxFS is supported in local zones. A VxFS file system is mounted as a loopback file system (LOFS) in local zones. Concurrent I/O and ODM are supported in local zones. See “[Local Zone Support](#)” on page 29.

SANVM Compatibility

Previously, VxFS provided two interfaces to SANVM, depending on whether or not the mounted file system was a cluster file system. There is now a single `ioctl` interface that freezes and thaws file systems regardless of how they are mounted. See the *VERITAS File System Administrator's Guide* for details.

Application Templates (Provider Only)

Application templates provide a parameterized approach to storage allocation and management. The VxFS graphical user interface integrates with the application template infrastructure to facilitate VxFS administration. See the *VERITAS File System Administrator's Guide* for details.

Software Developer's Kit (SDK)

The VERITAS File System Software Developer's Kit provides developers with information on how to use the application programming interfaces (APIs) provided with the VxFS Software Developer's Kit.

SDK features include:

- ◆ File Change Log
- ◆ Multi-volume support
- ◆ Named Data Streams
- ◆ VxFS I/O

See the *VERITAS File System 4.1 Programmer's Reference Guide* for details.

Other VxFS Enhancements

- ◆ File Change Log (FCL) is now fully supported.
- ◆ VxFS support for UIDS now exceeds 67 million.
- ◆ The `fsadm` command now allows fragmentation reporting and defragmentation operations on an individual file or directory, or on Storage Checkpoints.
- ◆ The `vxupgrade` command now upgrades only Version 4 and Version 5 disk layouts. Disk layout versions prior to Version 4 cannot be mounted.
- ◆ The `largefiles` option is now the default file size option for the `mount` command and `mkfs` command.
- ◆ The `delaylog` option is now the default intent logging mount option. The change of the default mount option from `log` to `delaylog` does not increase the risk of data loss, but allows VxFS to cache data to improve performance. See the *VERITAS File System Administrator's Guide* for more information.
- ◆ The default Storage Checkpoint creation mode is now removable.
- ◆ VxFS is now JumpStart compliant.
- ◆ Operation of the intent log replay was improved to increase the speed of recovery after a file system failure.



- ◆ The `histlog` function was implemented in the `fsdb_vxfs` command. The history log records structural changes to the file system to aid in product support.
- ◆ Two new tunable parameters, `inode_aging_count` and `inode_aging_size`, for use with the Storage Checkpoint API, were added to the `vxtunefs` command. See the `vxtunefs(1M)` manual page for more information.
- ◆ The `vxfsu_get_iooffsets` library call was renamed `vxfs_get_iooffsets`.
- ◆ More VxFS functions can be performed from the VERITAS Enterprise Administrator GUI (see the *VERITAS Enterprise Administrator (VEA 500 Series) Getting Started* guide for more information).
- ◆ The Lite version of VxFS now supports the `setext` command.

VERITAS Storage Foundation for Databases

VERITAS Storage Foundation for Oracle

This release of VERITAS Storage Foundation *for Oracle* includes the following new features and enhancements.

Storage Topology Analyzer

This release introduces functionality that maps tablespaces to physical disks and retrieves this mapping information, including the percentage of disk space used by a tablespace. This functionality is available using the CLI only.

Unified Message Identifiers

Logging, error, and warning messages are now saved to a log file with a standard-format Unique Message Identifier. An example of the standardized format is:

```
SFORA rollback INFO V-81-3000 VxDBA Rollback for Database to a
Storage Checkpoint.
```

Support for Multivolume File Systems in Storage Checkpoint Policy Administration

Storage Checkpoint policy administration functionality now allows database administrators to specify the volumes on which Storage Checkpoints can be created for each file system within the database. This allows you to control the space used by Storage Checkpoints and prevent Storage Checkpoints from fragmenting the space in the primary fileset. This functionality is available using the CLI.

Storage Checkpoint Quota Administration

With VERITAS Storage Foundation for Oracle, you can now administer file system quotas for Storage Checkpoints for a database or for multiple file systems using the CLI.

GUI Enhancements for Cloning a Database

The GUI has been enhanced to offer a single wizard for cloning a database using either a Storage Checkpoint or Database FlashSnap, as well as a single wizard for removing a clone database using either a Storage Checkpoint or Database FlashSnap.

New CLI Commands

The following VERITAS Storage Foundation *for Oracle* CLI commands and corresponding manual pages have been added:

Command Name	Description
<code>dbed_analyzer</code>	Maps tablespaces to physical disks and retrieves the information, including the percentage of disk space used by a tablespace. This option is not available through the VxDBA utility menu or through the GUI.
<code>dbed_ckptpolicy</code>	Creates and administers Storage Checkpoint allocation policies for a Multi-Volume File System (MVS). You can display, create, update, assign, and remove Storage Checkpoint allocation policies using this command. This option is not available through the VxDBA utility menu or the GUI.
<code>dbed_ckptquota</code>	Administers quotas for Storage Checkpoints. This option is not available through the VxDBA utility menu or the GUI.
<code>edgetmsg2</code>	Manages message log files. This option is not available through the VxDBA utility menu or the GUI.



Enhanced CLI Commands

The following VERITAS Storage Foundation *for Oracle* CLI commands have been enhanced:

Command Name	Description
<code>dbed_clonedb</code>	New options have been added that allow you to restart and remove a clone database created with a Storage Checkpoint. These options are <code>-o restartdb</code> and <code>-o umount</code> .
<code>dbed_ckptcreate</code>	A new <code>-p ckpt_policy[,ckpt_metadata]</code> option has been added that allows you to assign an allocation policy for Storage Checkpoint data. You can also optionally assign an allocation policy for Storage Checkpoint metadata.

VERITAS Storage Foundation for DB2

This release of VERITAS Storage Foundation *for DB2* includes the following new features and enhancements.

Support for DB2 8.2

This release of VERITAS Storage Foundation *for DB2* offers support for DB2 8.2.

Unified Message Identifiers

Logging, error, and warning messages are now saved to a log file with a standard-format Unique Message Identifier.

Support for Multivolume File Systems in Storage Checkpoint Policy Administration

Storage Checkpoint policy administration functionality now allows database administrators to specify the volumes on which Storage Checkpoints for each file system within the database. This allows you to control the space used by Storage Checkpoints and prevent Storage Checkpoints from fragmenting the space in the primary fileset. This functionality is available using the CLI.

Storage Checkpoint Quota Administration

With VERITAS Storage Foundation *for DB2*, you can now administer file system quotas for Storage Checkpoints for a database or for multiple file systems using the CLI.

Database FlashSnap

With this release of VERITAS Storage Foundation *for DB2*, Database FlashSnap operations are available via the GUI and CLI.

GUI Enhancements

The GUI has been enhanced to offer a single wizard for cloning a database using either a Storage Checkpoint or Database FlashSnap, as well as a single wizard for removing a clone database using either a Storage Checkpoint or Database FlashSnap.

New CLI Commands

The following VERITAS Storage Foundation *for DB2* commands and corresponding manual pages have been added:

Command Name	Description
db2ed_ckptpolicy	<p>Creates and administers Storage Checkpoint allocation policies for a Multi-Volume File System (MVS). You can display, create, update, assign, and remove Storage Checkpoint allocation policies using this command.</p> <p>This option is not available through the VxDBA utility menu or the GUI.</p>
db2ed_ckptquota	<p>Administers quotas for Storage Checkpoints.</p> <p>This option is not available through the VxDBA utility menu or the GUI.</p>
db2ed_checkconfig_all	<p>Checks the configuration of a DB2 Universal Database (UDB) Enterprise Server Edition (ESE) database with the Data Partitioning Feature (DPF) (that is, a partitioned DB2 database) in a VERITAS Storage Foundation <i>for DB2</i> environment.</p> <p>Performs the same operation from the command line, as Display Database/VxDBA Information > Examine Volume/File System/Database Configuration in the VxDBA utility menus.</p> <p>Performs the same operation from the command line, as right clicking on a database and selecting Check System Configuration in the GUI.</p>



Command Name	Description
db2ed_saveconfig_all	<p>Saves the configuration of a DB2 UDB DPF database in a VERITAS Storage Foundation <i>for DB2</i> environment.</p> <p>This option is not available through the VxDBA utility menu or the GUI.</p>
db2ed_vmchecksnap	<p>Creates and validates a snapplan that the db2ed_vmsnap command uses to create a volume snapshot of a DB2 database. The snapplan specifies snapshot scenarios (such as <code>online_snapshot</code>, <code>online_mirror</code>, or <code>offline</code>). The command can also be used to validate, list, copy, and remove a snapplan.</p> <p>This option is not available through the VxDBA utility menu.</p> <p>Performs the same operation from the command line, as the options within Snapplans > Create Snapplan... and Snapplan > Modify/Validate Snapplan in the GUI.</p>
db2ed_vmclonedb	<p>Mounts the file systems on the snapshot volumes and starts a clone database from snapshot volumes. You can also use this command to shut down or restart the clone database, unmount the file systems, or deport the clone database's volumes.</p> <p>This option is not available through the VxDBA utility menu.</p> <p>Performs the same operation from the command line, as the options within DB2 instances > Create Snapshot Database... in the GUI.</p>
db2ed_vmsnap	<p>Creates a snapshot image of a DB2 database by splitting the mirror volumes used by the database. You can also use this command to resynchronize the snapshot image back to the current database. The command also allows you to reverse resynchronize a snapshot image of a DB2 database.</p> <p>This option is not available through the VxDBA utility menu.</p> <p>Performs the same operation from the command line, as the options within Snapplan > Create Snapshot in the GUI.</p>
edgetmsg2	<p>Manages message log files.</p> <p>This option is not available through the VxDBA utility menu or the GUI.</p>

Enhanced CLI Commands

The following VERITAS Storage Foundation *for DB2* CLI commands have been enhanced:

Command Name	Description
db2ed_clonedb	New options have been added that allow you to restart and remove a clone database created with a Storage Checkpoint. These options are <code>-o restartdb</code> and <code>-o umount</code> .
db2ed_ckptcreate	A new <code>-p ckpt_policy[,ckpt_metadata]</code> option has been added that allows you to assign an allocation policy for Storage Checkpoint data. You can also optionally assign an allocation policy for Storage Checkpoint metadata.

VERITAS Storage Foundation for Sybase

This release of VERITAS Storage Foundation *for Sybase* includes the following new features and enhancements.

Unified Message Identifiers

Logging, error, and warning messages are now saved to a log file with a standard-format Unique Message Identifier.

New Command

The following VERITAS Storage Foundation *for Sybase* command and corresponding manual pages have been added:

Command Name	Description
edgetmsg2	Manages message log files.



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.

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 need to install the appropriate patches and then reboot.

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

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 need to 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 until you see "Sun StorEdge SAN 4.3." Click on the link to download. You have to be a registered user, and it 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 24 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 22, 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 `SUNW10f` if you plan to use the VERITAS Enterprise Administrator client. `SUNW10f` 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 VxFS

You must install the following patches for VERITAS File System before installing VERITAS Storage Foundation 4.1 for Oracle:

Operating System	Solaris Patch Number	Notes
Solaris 8	108528-14 (or later)	Patch 108528-14 requires patches 112396-02 (or later), 108987-09 (or later), 111293-04 (or later), 111310-01 (or later), and 111111-03 (or later). These patches must be installed in the order listed.
	108901-04	You must install Patch 108528-14 before installing Patch 108901-04.

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.

The latest status of patches 110934-10 and 113713-01 for use with specific VERITAS products is available at <http://support.veritas.com>.

VERITAS 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 manually install VERITAS patches

1. Make sure the same media disc you used to install VERITAS Storage Foundation is in your CD-ROM drive and mounted.
2. 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/* .
```

3. Decompress and then extract the patches.

```
# /cdrom/cdrom0/storage_foundation/scripts \
/install/gunzip *.gz
# tar xvf patch_id1.tar
```

Repeat the `tar` command above for each patch archive.

4. Install the patches individually using the `patchadd` command.

```
# patchadd patch_number
```

It is recommended that you install all the patches.

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

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

VxVM and VEA Patches

Please refer to the `patches` directory for information about the patches included in this release. You can find the `patches` directory at the same level as the `pkgs` and `scripts` directories on your software disc:

```
/cdrom/cdrom0/volume_manager/patches
```

It is highly recommended that you install all of the patches in this location on your system.



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 your system hangs, perform the following workaround

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

You need to 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 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 must be performed 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.



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 Limitations

- ◆ The VxDBA monitoring agent will fail to obtain Oracle tablespace information if the home directory of the Oracle DBA user (usually `oracle`) is different than the actual value of `ORACLE_HOME`. [295399]

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]

- ◆ 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 29.
- ◆ 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.

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 Limitation

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]



Fixed Issues

VERITAS Volume Manager Fixed Issues

Incident	Description
i102381	Using vxdiskadm to replace a failed disk requires two attempts.
i135566	VEA: Incorrect vxpool command
i140218	Volumes showing as DISABLED.
i137098	VEA: Setting a Comment on an ISP Volume.
i137625	VEA: Administering a Cache Volume Created on an ISP Volume
142289	VVR: V_PFLAG_NEED_RESYNC set erroneously.
146744	Solaris: vxswapctl fails with error when root is encapsulated
150064	Bogus DA record causes boot from VxVM root to fail in DBED/AC environment.
154913	Modifying CVM timeouts requires a reboot.
155139	When creating concat-mirror, users see warnings “No disk space matches specification”.
156011	Help and documentation for dmp_io_timeout should say seconds and not milliseconds.
156409	Adding a new cluster node requires cluster reboot.
156804	VVR VRAS: in.vxrsyncd and vxnetd should not call gethostname/gethostbyname during startup.
157058	DMP: enabling the controller might lead to corruption if dev_t is changed.
158823	vx dg move fails with VE_AGAIN due to device open/close assertion failure.
158911	FMR: Full resync required when a disk array forming one of the mirrors is unavailable across forced import of disk group.



Incident	Description
207290	Multiple problems with PGR support in DMP.
209541	MR3: Panic in voldco_merge_offset_lists due to incorrect handling in drl of volume shrink.
212397	Trespass initiated when vold issues MHIOCSTATUS ioctl through dmp_reserve_release.
212723	vxdumpinq and vxcheckasl can cause trespass on AP arrays.
217135	Panic on boot when using A/PG array with 64 LUNs
221295	vxclust error: Using CVM internal node id assignment scheme.
225179	Add an option to vxbootsetup to give user the choice to create partitions for non-system volumes.
228805	vxsync core dump due to uninitialized malloc memory.
229187	Panic in vol_subdisk_iogen.
232276	In case of message drop due to unreliable network protocol, CVM joiner does not respond to resends of JOIN_START message. This will lead to different view points on the membership among all nodes and endless reconfiguration. This scenario will be observed only when UDP is used and the workaround is to restart the cluster.
259150	Need to automatically mark the root volume to “nologging” (in /etc/vfstab) during root encapsulation.
262077	When using an Emulex HBA driver on a system running Solaris 10, it can take a long time for DMP to fail a path. We recommend that you do not use an Emulex HBA until this issue is fixed.
263010	When a master node aborts in the middle of the cluster path failover protocol and the slave that initiated the path failover protocol is in the process of taking over as master, vxreconfd will hang. In this situation, you should reboot the system.
267270	Sometimes CVM fails to recover from a previous reconfig and hangs. If this happens, you must reboot the system.
268063	Occasionally vxconfigd dumps core in the rec_touch() function when disks are rediscovered. The workaround is to turn off the ssb flag on the disk group by typing vxdg -g <diskgroup> set ssb=off.



Incident	Description
271206]	Occasionally CVM might hang when <code>ddpolicy=local</code> .
271460	Sol: encapsulation does not check for rootvol in the group.
271658	MPxIO does not currently work on Solaris 10.
274504	Error message shown after VRTS installation.
275542	Package removal on alternate disk may fail.
275944	<code>es_devfs.pl</code> should check if <code>vxesd</code> is running prior to exiting (error : syseventcnfd: process xx exited with status 3).
277867	Kernel Panic in <code>vol_vset_getvols_sz/bcopy</code> - buffer overflow.
278573	Incorrect handling of <code>VE_RESTART</code> leads to failures in <code>vxplex</code> , <code>vxsnap</code> operations when snapshots are created in parallel.
283674	Panic at <code>voldco_write_new_map()</code> .
none	Upgrading to Solaris 10 Using Live Upgrade With Encapsulated Root Disk May Fail [Sun Bug ID 5087504]
none	Restrictions on Using Fabric Devices
none	The Sun StorEdge Traffic Manager (SSTM) boot support feature that is available through SAN 4.3 is not supported with VxVM. [Sun Bug ID 4912232, 4909641, 4912667]

VERITAS File System Fixed Issues

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

Reorganizing Extents and Mapped Files

The reorganization of structural ilist extents or mapped files no longer fail due to allocations in space being removed from the file system.

In previous releases, VxFS was unable to move structural ilist extents or mapped regular files. Higher-level operations which required this ability failed. Prior to version 4.0, this affected only the file system shrink operation; however, with the introduction of multi-volume file systems this also affected volume shrink and removal.



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
33903	Previously, if a mount point was a symbolic link, the space monitoring agent would sometimes fail. This issue has been fixed.
34201	Previously, while restoring database files, Quick I/O files would not be recreated automatically. This issue has been fixed.
34295	Previously, you could not roll back a tablespace or datafile to a Storage Checkpoint using the GUI. This issue has been fixed.
34387	If 200 or more files resided in the <code>/dev/rdsk</code> directory, the <code>dbed_filestat</code> command would take a long time to display statistic information. This issue has been fixed.
75741	Error messages or informational messages that only allow you to click on OK to continue did not contain a mnemonic. This issue has been fixed.
85155	Previously, some Storage Checkpoint Capacity Planning messages were not internationalized. This issue has been fixed.
86264	Previously, using VERITAS Database FlashSnap on a database with volsets was not supported. This configuration caused the <code>dbed_vmchecksnap</code> command to fail because the volsets were not recognized. If you attempted to run <code>dbed_vmchecksnap</code> , the following error message would display: ERROR: The fastresync flag for Volume dbedvset on proddg is OFF, please set it ON This issue has been fixed.
86526	When Storage Foundation <i>for Oracle</i> removed a Storage Checkpoint based on the Storage Checkpoint allocation policy set in the monitoring agent configuration, the directory in which the Storage Checkpoint resided was not removed. This issue has been fixed.



Incident	Description
103787	<p>Previously, if you obtained datafile topology or statistics information on a snapshot database using the GUI, you would not be able to resynchronize the snapshot until you restarted the VEA Service. If you tried to resynchronize the snapshot, you would see an error message similar to the following:</p> <pre>DBED5616: ERROR: join diskgroup <i>diskgroup1</i> <i>diskgroup2</i> failed.</pre> <p>This issue has been fixed.</p>
129887	<p>Previously, VERITAS Storage Foundation for Oracle did not support Japanese datafiles and Japanese tablespace names. This issue has been fixed.</p> <p>If you are using an NLS_LANG other than English, you need to add it in your <code>/etc/vx/vxdba/ORACLE_SID/settings</code> file so that datafile and tablespace names display in the correct language in the GUI. For example, add <code>export NLS_LANG=JAPANESE_JAPAN.JA16EUC</code> in the <code>settings</code> file so that Japanese datafile and tablespace names display in the GUI.</p>

Incident	Description
221703	<p>Previously, Database FlashSnap did not work in an HA environment. This release provides Database FlashSnap support in an HA environment.</p> <p>The following limitations apply:</p> <ul style="list-style-type: none"> • In an HA environment, you must modify the default snapplan to use the virtual host name defined for the database resource group for the <code>PRIMARY_HOST</code> and/or the <code>SECONDARY_HOST</code> parameters and validate the snapplan before creating a snapshot by running the following command: <pre>dbed_vmchecksnap -S ORACLE_SID -H ORACLE_HOME \ -f SNAPPLAN -o validate</pre> • In an HA environment, the primary database must be down before you perform reverse resynchronization (<code>dbed_vmsnap -S ORACLE_SID -f SNAPPLAN -o reverse_resync_begin</code>). When VERITAS Cluster Server detects that the primary database is down, however, it starts the failover process and the VxDBA repository is unmounted and the <code>dbed_vmsnap</code> command is aborted. <p>To work around this issue:</p> <ol style="list-style-type: none"> 1. As root, temporarily freeze the VCS resource group for the database: <pre># hagr -freeze ResourceGroup</pre> 2. Shut down the primary database. 3. Run reverse resynchronization: <pre># dbed_vmsnap -S ORACLE_SID -f SNAPPLAN \ -o reverse_resync_begin</pre> 4. After reverse resynchronization changes are committed (<code>-o reverse_resync_commit</code>), verify that the database has been started in ARCHIVELOG mode. This information is provided in the status messages that appear after running committing reverse resynchronization changes. 5. Unfreeze the resource group: <pre># hagr -unfreeze ResourceGroup</pre>



Incident	Description
222898 266618	<p>The following problems occur when performing a multiple database parallel backup using flashsnap tools:</p> <ul style="list-style-type: none">♦ The <code>dbed_vmsnap -o resync</code> command can lead to an unrecoverable disk group.♦ The <code>dbed_vmchecksnap -o validate</code> command leads to discrepancy between the GUI and the <code>vxprint</code> output. <p>Database Flashsnap tools use the Volume Manager Provider to get volume and disk properties. During a disk group split/join and volume snapshot/resync, events are generated to notify the ISIS database about the changes in disk group/volume. Because of a large number of events, <code>vxsvc</code> gets delayed in processing them in a timely manner.</p> <p>Flashsnap tools like <code>dbed_vmchecksnap</code> and <code>dbed_vmsnap</code>, need real-time information.</p> <p>These two issues are resolved by relying on <code>vxconfigd</code> for volume and disk information via the <code>vxprint</code> utility.</p>
223123	<p>Previously, the length of a rescan interval was not consistent with the amount of time set for the interval. This issue has been corrected.</p>
236259	<p>Previously, the <code>dbed_vmchecksnap -o list -f filename</code> command would remove files from the repository, which would cause the <code>dbed_vmsnap -o resync</code> operation to fail. This issue has been fixed.</p>
236519	<p>Facilitates Database Flashsnap in creating multi-mirror snapshots. For further details, see “Creating Multi-mirror Snapshots” on page 44.</p>
260271	<p>Previously, if you converted a VxFS file to an Oracle Disk Manager file from the GUI, you could not convert it back to a VxFS file even if Cached Quick I/O was enabled. This issue has been corrected.</p>

Incident	Description
260671	<p>Previously, Storage Checkpoint information stored in the VxDBA repository was not carried over with snapshot volumes to the secondary host. This release provides support for mounting a Storage Checkpoint carried over from the snapshot volumes to a secondary host.</p> <p>On the secondary host, you can list the Storage Checkpoints carried over from the primary database using:</p> <pre>dbed_ckptdisplay -S ORACLE_SID -n</pre> <p>You can mount one of the listed Storage Checkpoints using:</p> <pre>dbed_ckptmount -S ORACLE_SID -c CKPT_NAME \ -m MOUNT_POINT</pre> <p>The following limitations apply:</p> <ul style="list-style-type: none"> Any mounted Storage Checkpoints must be unmounted before running the following commands: <pre>dbed_vmsnap -S ORACLE_SID -f SNAPPLAN \ -o reverse_resync_begin dbed_vmclonedb -o umount,new_sid=new_sid \ -f SNAPPLAN</pre> It is only possible to mount a Storage Checkpoint carried over with the snapshot volumes in a two-host configuration if the snapshot volumes were mounted with the <code>dbed_vmclonedb</code> command with the <code>-o mount</code> option without the use of <code>-r relocate_path</code>. Storage Checkpoints carried over with the snapshot volumes can be mounted before a clone database is created using <code>dbed_vmclonedb</code> with the <code>-o mount</code> option. After a clone database is created using <code>dbed_vmclonedb</code> with the <code>-o recoverdb</code> option, however, Storage Checkpoints are no longer present. After running <code>dbed_vmsnap -o reverse_resync_commit</code>, Storage Checkpoints created on the original database or on the clone database are deleted.
none	<p>You cannot run the VxDBA monitoring agent on a clone database created using a Storage Checkpoint or Database FlashSnap. The VxDBA monitoring now disables itself if you attempt to run it on a clone database created using a Storage Checkpoint or Database FlashSnap.</p>
none	<p>Remote installation and uninstallation of VERITAS Storage Foundation <i>for Oracle</i> is now supported.</p>



Creating Multi-mirror Snapshots

This procedure describes the process for creating multi-mirror snapshots using Database FlashSnap tools.

To make Database Snapshots highly available, the snapped snapshot volume should contain more than one mirror. This makes the snapshot volumes available even if one of the mirrors gets disabled. Snapshot volumes can be mounted and the entire database snapshot is usable even if one of the mirrors gets disabled. The multi-mirror snapshots are enabled via `SNAPSHOT_MIRROR=<n>` in the snapplan.

Note There are no changes to the Command Line usage or arguments for the FlashSnap tools.

Note Before taking the snapshot, make sure all tagged snapshot mirrors are in SNAPDONE state.

▼ To set up an Oracle database to take multi-mirror snapshots

1. Add the second mirror and DCO log. When allocating storage for the second mirror and DCO logs, it is important to be careful so that the snap volumes are splittable. If snap volumes are not splittable, `dbed_vmchecksnap` fails with appropriate errors.

Tag the SNAPDONE mirror with `dbed_flashsnap`.

Assume that the volume has `fastresync = on` and already has one SNAPDONE mirror and is tagged with `dbed_flashsnap`

```
# vxsnap -g dg_a addmir dg_a_vol1 alloc=dg_a03
# vxedit -g dg_a set putil2=dbed_flashsnap dg_a_vol1-03
```

2. Add `SNAPSHOT_MIRROR` keyword to the snapplan. Here is a sample snapplan.

```
SNAPSHOT_VERSION=4.1
PRIMARY_HOST=leopard
SECONDARY_HOST=leopard
PRIMARY_DG=PROD_dg
SNAPSHOT_DG=SNAP_PROD_dg
ORACLE_SID=PROD
ARCHIVELOG_DEST=/PROD_dg_a5
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=2
```

VERITAS Storage Foundation for DB2 Fixed Issues

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

Incident	Description
none	Remote uninstallation of VERITAS Storage Foundation <i>for DB2</i> is now supported.
75944	<p>If Oracle and DB2 databases reside on the same host, mapping and statistics from the GUI will work on the database used first. If the other database is used at the same time, the following error message displays:</p> <pre>ERROR:Internal error in vxstorage_stats</pre>
86731	<p>If a database instance is down and the <code>db2ed_ckptcreate_all</code> command is issued, the following error message displays:</p> <pre>Database on node 0 has already been quiesced by others. Please try again later.</pre> <p>The error message should say:</p> <pre>Instance Down</pre> <p>This problem will be fixed in the next release.</p>
86736	Previously, the GUI could not display database objects in a Chinese environment (LANG=zh). This issue has been fixed.
87000	Previously, in the VxDBA utility, the same message would display twice in the Configuring Monitoring Agent menu. This issue has been fixed.
209276	<p>When executing various options of the <code>vxdba</code> command, the database status may indicate Database Down as shown below.</p> <pre>VERITAS Storage Foundation for DB2 (DB2DATABASE 'PROD') Menu: Database Administration Database Status : Database Down # File Systems : (1) # Tablespaces : (3) # Containers : (4)</pre> <p>The Database Down status means that the database is inactive, and the instance is up. If the instance is down, all databases of that instance are inactive, and the database status is Instance Down.</p> <p>An explanation of the difference between the Database Down status and the Instance Down status has been added to the <i>VERITAS Storage Foundation for DB2 Administrator's Guide</i>.</p>
223123	Previously, the length of a rescan interval was not consistent with the amount of time set for the interval. This issue has been corrected.



Incident	Description
269858	<p>Previously, if you attempted to create a clone database from an online or offline Storage Checkpoint, the <code>db2ed_vmclonedb</code> command would fail and display the following error message:</p> <pre>SFDB2 db2ed_clonedb ERROR V-81-7211 Relocate failed.</pre> <p>This issue has been fixed.</p>
271299	<p>You cannot run the VxDBA monitoring agent on a clone database created using a Storage Checkpoint or Database FlashSnap. The VxDBA monitoring now disables itself if you attempt to run it on a clone database created using a Storage Checkpoint or Database FlashSnap.</p>
299826	<p>Previously, <code>db2ed_vmchecksnap -o validate</code> would not fail as expected if RAID-5 volumes were discovered. This issue has been corrected.</p>

VERITAS Storage Foundation for Sybase Fixed Issues

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

Incident	Description
none	<p>In the 4.0 release of the <i>VERITAS Storage Foundation for Sybase Administrator's Guide</i>, the procedure and example for creating a warm standby server were incorrect. They have been corrected.</p>
41994	<p>In earlier releases, running the command <code>qio_convertdbfiles -f</code> would produce the following error message if the file was a Quick I/O link file:</p> <pre>No such file or directory</pre> <p>This issue has been corrected.</p>

Software Issues

VERITAS Storage Foundation

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.

VERITAS Volume Manager Software Issues

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

Disk Controller Firmware Upgrades

This procedure enables you to upgrade disk controller firmware without performing a system reboot. The procedure is a workaround for [Sun Bug ID 4164338].

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

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

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

1. Disable the plex associated with the disk device:

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



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

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

For the other controller on the HBA, enter:

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

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

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

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

5. Enable the plex associated with the device:

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

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

6. Upgrade the disk group:

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

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

T3B Firmware Upgrade on Solaris 9

On Solaris 9 only, a T3B upgrade to firmware version 2.1 must follow the procedure below. Not using the procedure leads to disabled disk groups or an inability to mount file systems. (i95877)

1. Use the `umount` command to unmount related file systems

```
# umount mount_point
```

2. Stop all VxVM volumes:

```
# vxvol -g dg_name stopall
```

3. Stop VxVM:

```
# vxdctl stop
```

```
# vxiod -f set 0
```

4. Upgrade the T3B firmware to version 2.1.

5. Start VxVM:

```
# vxiod set 10
# vxconfigd -m disable
# vxdctl enable
```

6. Start the VxVM volumes:

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

7. Use the mount command to remount the file system, for example:

```
# mount -F vxfs /h/filesys
```

Installation and Upgrade Issues

Initializing Disks Previously Under VxVM Control

If you are planning to initialize disks, check to see if any of the disks were previously under VxVM control. If so, and if they were used on the same host system, the disk groups they represent are imported automatically during the installation process if the proper removal procedures were *not* followed. An attempt during installation to initialize or encapsulate disks that were previously under VxVM control fails. After installation, if you no longer want to use those disk groups, use the destroy option of the `vxdg` (1M) command to remove those disk groups. Alternately, you can use `vxdiskunsetup` (1M) to remove the disks from VxVM control. Be aware that these options can result in data loss if used incorrectly.

Patch Installation and Removal Instructions

For Solaris 8, 9 and 10 releases, refer to the manual pages for instructions on using the `patchadd` and `patchrm` scripts provided with Solaris.

The following example installs a patch to a stand-alone machine:

```
# patchadd 113701-01
```

The following example removes a patch from a stand-alone system:

```
# patchrm 113701-01
```

For additional examples, please see the appropriate manual pages.



Special Install and Uninstall Instructions

You need to use the `shutdown` command to reboot the system after patch installation or uninstallation:

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

Recognizing Simple Disks from Earlier Releases

In earlier releases of VxVM, some users minimized the allocation of disks to the disk group, `rootdg`, by associating `rootdg` with a small disk partition that was characterized as a simple disk. This procedure would have been achieved by using the command; `vxctl add disk` which is no longer supported in VxVM versions 4.0. and later. [137838]

If you created one of these simple disks, you will need to carry out a procedure similar to the one described in the following example.

Assuming that the simple disk is defined to be on `c1t21d0s7`, you would see the following entry in `/etc/vx/volboot`:

```
disk c1t21d0s7 simple privoffset=1
```

After upgrading to VxVM 4.1, you must reboot the system. After rebooting, execute the command, `vxdisk list`, and you will see that `c1t21d0s7` is not listed. This is because `vxconfigd` now ignores disk entries in `/etc/vx/volboot`.

▼ To retain access to data on a simple disk

1. Define a disk access record that will be created in `/etc/vx/darecs`

```
# vxdisk define c1t21d0s7 type=simple
```

2. Request that `vxconfigd` should extract information from this disk:

```
# vxctl enable
```

3. Discover the name of the disk's disk group:

```
# vxprint -th
```

4. Enable access to the disk's disk group; `rootdg` in this example:

```
# vxvol -g rootdg startall
```


Error Messages Output by the Upgrade Script

If a swap volume specified in `/etc/vfstab` is mirrored at the time that `upgrade_start` is run, the `upgrade_finish` 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:

```
xvm:vxvol: tutil0 field for plex plex_name changed unexpectedly
```

This message can be ignored.

For a system on which the root file system is contained on a mirrored volume, the `upgrade_start` 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 `upgrade_finish` can initially fail, claiming that the mirror on the boot disk is stale, as follows:

```
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.
```

Boot the system from one of the disks named. If the `eeeprom` option `use-nvramrc?` is set to `true`, boot the system by specifying `vx-diskname`.

Partitioning Disks on a Sun StorEdge 6120/6320 Array

If you plan to partition disks on a 6120/6320 array, it is important that you install VxVM and the 6120/6320 Array support Library packages in the correct order. See “Partitioning Disks on 6120/6320 Arrays” in the *VERITAS Volume Manager Hardware Notes* for more information. [115328]

Host Crashes During an Upgrade

In certain circumstances, the installation software cannot detect that a host has crashed during an upgrade installation. In some cases, the `rsh/ssh` request hangs for an indefinite period of time.

[12988]

Interruption of an Upgrade

If the installation software is interrupted on the local system during certain upgrade situations, VERITAS Volume Manager configurations may be lost after a reboot. If this happens, the entire VERITAS Volume Manager package must be reinstalled and a recover must be done manually by recreating the disks, disk groups, and volumes and restoring the data from backup. [13033]



“SIGHUP caught” Message on the Console

When running `vxinstall` on a system with a SENA array that is enabled with enclosure naming, you may see a message similar to the following:

```
newmac.veritas.com console login: Dec 31 00:00:00
syseventd[59]: SIGHUP caught
```

You can safely ignore this message. [Sun Bug ID 4955989, i138955]

Live Upgrade Only Supports Upgrades from Solaris 9 to 10

- ◆ To perform the Live Upgrade procedure you must have the Live Upgrade software from Sun Microsystems.
- ◆ Solaris Live Upgrade and Volume Manager Live Upgrade do not support the upgrade from Solaris 7 or 8 to 10. If the current OS level is Solaris 7 or 8, you must upgrade to Solaris 9 before using Live Upgrade to upgrade to Solaris 10.
- ◆ If you have already upgraded from Solaris 9 to Solaris 10, you must uninstall VxVM and reinstall it for the VxVM4.1 commands to be active even if VxVM 4.1 is already installed on Solaris 9). [268194]

Use the following procedure to upgrade Solaris 9 to Solaris 10

1. Upgrade Solaris using the `vxlustart` command.
2. **# `pkgrm -R /altroot.5.10 VRTSvxvm`**
3. **# `pkgadd -R /altroot.5.10 -d `pwd` VRTSvxvm`**
4. Complete the upgrade using the `vxlufinish` command.

Limitation of `upgrade_start`

The `upgrade_start` script processes the partitions: `/`, `swap`, `/usr`, `/var`, `/opt`, and `/export/home`.

With other partitions, including `/usr/openwin`, and any user-defined partitions (e.g., `"/abc"`) of the root disk, the script will encounter an error if the associated tag value in the disk's VTOC is not set to 0x0. It will display an error message similar to this:

```
The following conditions will prevent successful unencapsulation of
the requested file systems:
Device clt1d0s2 in volume opt for file system /opt does not have a
reasonable slice available for use in unencapsulating the volume.
```

You can free up a slice for this file system using the command

```
# /etc/vx/bin/vxpartm /dev/rdisk/c1t1d0s2 <slice-number>
```

Slice 0 is reserved for use by root file system slices, so free up a different slice number.

Please correct these problems (if possible) and try again later.

The partition tag values may be displayed with the `prtvtoc(1M)` command.

The tag value can be set using the `format(1M)` command. While configuring a partition, specify a tag string of "unassigned" for `/usr/openwin` and any user-defined partitions from the root disk. "unassigned" corresponds to the numerical tag value of 0x0. [309265]

Misleading Alerts Generated on a System Where VAIL Package is Installed

An alert with the text message "SymCLI command line tools are not installed properly" will be generated in each of the following two cases when SYMCLI is either absent or not installed properly on the host on which a VAIL package is installed.

Case 1. When host comes up after a reboot and SYMCLI is either absent or not installed properly.

Case 2. When a rescan of Symmetrix provider is initiated and SYMCLI is either found to be absent or found to be not installed properly but SYMCLI installation was proper before rescan of Symmetrix provider was initiated.

In either of Case 1 or Case 2 one should ignore the alert message on the host on which VAIL package is installed if there is no EMC Symmetrix array being managed on that host. [Sun Bug ID 6211778, 297830]

Utility Issues

Shrinking a Swap Volume

`vxassist` has no built-in protection to prevent you from shrinking the swap volume without first shrinking what the system sees as available swap space. If it is necessary to shrink the swap volume, the operation must be done in single user mode and the system must be rebooted immediately. Failing to take these precautions can result in unknown system behavior or lock-up. [i6154]

Mount Option Set During Root Disk Encapsulation

During the root disk encapsulation process, VxVM explicitly sets the mount option of the `/`, `/usr`, `/var`, `/opt`, `/export/home`, `/usr/openwin`, and swap partitions to "nologging".



When the root disk is unencapsulated, the previous UFS logging mount option is not restored. The previous UFS logging mount option may be "logging" or "nologging". It is not necessary to restore the mount option to its value prior to encapsulation. By not restoring it, these partitions will function with UFS logging disabled. However, to restore the UFS logging option, reset the mount option of the relevant root disk partitions by editing the `/etc/vfstab` file. To activate the settings, you must either reboot the system or issue a remount command. An example of the remount command is:

```
# mount -o remount,logging /dev/dsk/devicename mountpoint
```

The setting of the mount option also occurs during Live Upgrade and the use of the `upgrade_start` and `upgrade_finish` scripts. [284322]

Adding a Log and Mirror to a Volume

The `vxassist` command does not add a mirror and a log when processing a command such as the following:

```
# vxassist mirror volume layout=log ...
```

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

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

[i13488]

Replacement of the `old_layout` Attribute

The `vxdisksetup` command gives the error message `Attribute unrecognized` when the `old_layout` attribute is used to make a disk into a VxVM controlled disk. The `old_layout` attribute is no longer supported. Use the `noreserve` attribute in its place. [i121258]

Using `vxvol` and `vxmend` With Layered Volumes

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

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

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

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

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

[i134932]

Growing or Shrinking Layered Volumes

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

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

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

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

Maximum Size of a VxVM Volume

VxVM supports volume lengths up to $2^{63}-1$ disk sectors when using VERITAS-specific `ioctl` calls. However, system calls such as `seek`, `lseek`, `read` and `write` are limited to a maximum offset that is determined by the operating system. For a system that supports large files, this is usually $2^{63}-1$ bytes. Otherwise, the maximum offset value is usually $2^{31}-1$ bytes (1 byte less than 2 terabytes). The maximum size of a VERITAS File System (4.0 onward) file system that can be created on a VERITAS Volume Manager (4.0 onward) volume is documented in the *VERITAS File System 4.0 Release Notes*. [i141024]



Converting a Multipathed Disk

Under Solaris 10 when converting a multipathed disk that is smaller than 1TB from a VTOC label to an EFI label, you must issue a `format -e` command for each path. For example, if a node has two paths, `c1t2d0s2` and `c2t2d0s2`, you need to apply the `format -e` command to each of the two paths. [269566]

Startup Script Messages Not Seen on Console

With the introduction of SMF support in Solaris 10, startup script messages are no longer seen on the console. [269949]

These messages can be viewed (`cat` or `vi`) in SMF log files found at:

```
/var/svc/log
/etc/svc/volatile
```

The file names are based on the specific startup script:

```
#/var/svc/log: ls
system-vxvm-vxvm-startup2:default.log
system-vxvm-vxvm-sysboot:default.log
```

Also, other startup messages can be found in:

```
#/var/svc/log: ls
milestone-multi-user-server:default.log
milestone-multi-user:default.log
milestone-name-services:default.log
milestone-single-user:default.log
```

```
#/etc/svc/volatile
```

```
system-vxvm-vxvm-startup2:default.log
system-vxvm-vxvm-sysboot:default.log
```

Bad Disk Block Warning

When `vxio` detects a bad disk block on a disk, it will display a warning message indicating that an uncorrectable write error has been encountered. [272176]

Changing a Disk's I/O Policy

If you want to change a disk's I/O policy, the disk needs to be in a quiescent state, but when it is in this state, `vxconfigd` appears to hang. We recommend that you do not attempt to change a disk's I/O policy until this issue is fixed. [272263]

Incomplete Records in /etc/vx/disk.info

Sometimes a system that is not properly configured might result in `/etc/vx/disk.info` containing incomplete records. This situation might result in `vxconfigd` dumping core. If this occurs, you should remove `/etc/vx/disk.info` and restart by running the command:

```
# vxconfigd -kr reset
```

[272309]

Snapshots for a RAID-5 Volume

When taking a snapshot for a RAID-5 volume, it was observed that the checksum of the files on the snapshot and the original volume had a discrepancy. This discrepancy occurred because when the RAID-5 volume was created, it was not created as a multiple of 128 blocks. When creating a RAID-5 volume, make certain that its size is a multiple of 128 blocks. [272626]

Volume in Sync State

- ◆ When a volume is in SYNC state and the system crashes while the commands, `vxsnap prepare vol` or `vxvol set DRL=on` are running, the volume recovery operation will incorrectly use the DRL to recover the volume when the system is restarted.

We recommend that you do not use `vxsnap prepare vol` or `vxvol set DRL=on` when the system is in the SYNC state.

- ◆ If the volume is stopped, and DRL is enabled using `vxvol set DRL=on`, the volume recovery operation will incorrectly use the DRL to recover the volume.

We recommend that you do not use `vxvol set DRL=on` when the volume is in either a SYNC or a NEEDSYNC state.

- ◆ When the DCO has a BADLOG flag set, the recovery operation incorrectly uses the DRL to recover the volume.

If the BADLOG flag is set, you should remove the DCO before starting the volume. [273314]

Do Not Specify a Long Device Name in /etc/vx/

Do not specify a long device name in `/etc/vx/disks.exclude`. Some scripts like `vxdiskadm` fail with an error message if a long device name is specified in this file. Use `vxdiskadm(1M) option 17/18` to suppress/unsuppress devices from VxVM's view. [Sun Bug ID 6228464, 311275]



Encapsulating Boot File System That Spans Two Disks

If the boot file system spans two disks and the first encapsulated disk contains `/`, `/usr/`, or `/var`, an error message will occur when attempting to encapsulate the second disk. The encapsulation will complete without any issues. The error message(s) of the disks 'not found' will look similar to the following:

```
Continue with encapsulation? [y,n,q,?] (default: y)
/usr/lib/vxvm/voladm.d/bin/disk.encap: =__disk: not found VxVM
INFO
V-5-2-333 The disk device clt1d0 will be encapsulated and added to
the
disk group mydg with the disk name mydg02.
```

[Sun Bug ID 6230678, 310784]

Multiple Enclosures Generate Unexpected Names

If there are multiple enclosures of the same array types in the configuration, the enclosure names that are generated will be numbers appended to the enclosure type in the form of "xxx0, xxx01, xxx012". [229538]

Changing an Enclosure Name

`vxconfigd` will core dump if a user changes an enclosure name to a new name which is longer than the current one. The workaround is to rename the enclosure with a name consisting of the same number of characters or fewer. [Sun Bug ID 6230031, 311530]

Unable to Boot System Without `bootdg` Link to Boot Disk Group

When a system fails to boot or is not mounting with following errors:

```
ERROR: svc:/system/filesystem/root:default failed to mount /usr
(see 'svcs -x' for details)
[ system/filesystem/root:default failed fatally (see 'svcs -x' for
details) ]
Requesting System Maintenance Mode
(See /lib/svc/share/README for more information.)
Console login service(s) cannot run

Root password for system maintenance (control-d to bypass):
single-user privilege assigned to /dev/console.
Entering System Maintenance Mode

Feb 14 23:41:26 su: 'su root' succeeded for root on /dev/console
su: No shell /bin/ksh. Trying fallback shell /sbin/sh.
-sh: /bin/i386: not found
-sh: /usr/sbin/quota: not found
```



```
-sh: /bin/cat: not found
-sh: /bin/mail: not found
-sh: -o: bad option(s)
```

One possible cause for the error that the symbolic link of bootdg to the boot disk group under /dev/vx/dsk/ or /dev/vx/rdisk/ is missing.

▼ **The workaround for this error is as follows**

1. Make sure that your system does not have a link under /dev/vx/dsk and /dev/vx/rdisk

```
bootdg -> rootdg
```

2. Boot the system from cdrom or net.
3. Mount the / from CDROM. In this example cxtxdx is the boot disk.

```
# mount -F ufs -o nologging /dev/dsk/cxtxdxs0 /mnt
```

4. Create the link.

```
# cd /mnt/dev/vx/dsk
# ln -s rootdg bootdg
# cd /mnt/dev/vx/rdisk
# ln -s rootdg bootdg
```

If the boot disk group is called "rootdg", then

```
# ln -s rootdg bootdg

# cd
# umount /mnt
#init 0
```

5. Let the system boot by disk in question.

[Sun Bug ID 6230224]



Device Issues

Hitachi Arrays in Active/Active Mode

When Hitachi DF400 and DF500 arrays are configured as Active/Active mode, performance is degraded. [i73154]

Relayout of Volumes on the Root Disk

Do not run the `vxrelayout` and `vxassist` commands to relayout a volume that is part of root disk. This action may ruin the layout of the root disk so that you will not be able to boot from the disk. On an encapsulated root disk, a relayout would cause the upgrade to fail. [i103991]

Failure to Add a Disk from a T3 Array

On a T3 array, Volume Manager may get the following failure when trying to add a disk (typically from `vxinstall` or `vxdisksetup`):

```
vxvm:vxdisk: ERROR: Device XXXX: online failed
Device path not valid
```

This can happen in cases where the T3 disk was re-partitioned (or re-formatted) prior to one or more disks being added. [i105173]

SFCFS with I/O fencing not supported on HDS9200

If you attempt to boot a cluster with I/O fencing (PGR) enabled, HDS9200 disks will show up in error state on the slaves. This error does not appear if I/O fencing is disabled. [i131926]

Disks in V480 and V880 Internal Disk Enclosures

Fujitsu and Hitachi disks in V480 and V880 internal disk enclosures may not be automatically recognized as JBOD disks. This could potentially cause data corruption if multipathing is not configured correctly. After installing any Sun-qualified FC disks as FRU replacements, or before running `vxinstall` during installation of VxVM, use the procedure described in [“Adding Unsupported Disk Arrays to the JBOD Category”](#) on page 74 to add each such disk to the JBOD category. It is important that both the vendor ID and product ID are specified for each such disk to avoid conflicts with similar disks in other arrays. For Fujitsu disks, the number of characters in the serial number must also be specified. [Sun Bug ID 4900508, i133579]

Encapsulation of Disks with Insufficient Space for a Private Region

Disks with insufficient space (less than 2048 disk blocks) for the allocation of an on-disk database copy cannot be encapsulated. The database requires at least the same space as is allocated for other disks in the same disk group. The default size is 2048 blocks. To work around this, relocate the data on the last partition of the disk to a volume on a different disk, and free the space by reducing the partition size to 0.

The space for the database must be allocated from the beginning or the end of the disk, with the exception of the root disk. The root disk can be encapsulated by carving out space from the swap 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 swap partition.

Workaround: The problem of insufficient space on a disk to store private VxVM information has no workaround. VxVM requires at least a small region of private storage (2048 blocks) for proper disk identification.

Errors When Using JNI Cards

If the model number of your JNI card is one of FCE-1063, FCE2-1063, FCE-6410, FCE2-6410, or FCE2-6412, you may experience error messages of the form:

```
Oct 22 00:16:16 ds13un jnic: [ID 847178 kern.notice] jnic1: Memory
port parity error detected
Oct 22 00:16:16 ds13un jnic: [ID 229844 kern.notice] jnic1: Link Down
Oct 22 00:16:16 ds13un jnic: [ID 744007 kern.notice] jnic1: Target0:
Port
0000EF (WWN 500060E802778702:500060E802778702) offline.
Oct 22 00:16:18 ds13un jnic: [ID 709123 kern.notice] jnic1: Link Up
Oct 22 00:16:18 ds13un jnic: [ID 236572 kern.notice] jnic1: Target0:
Port
0000EF (WWN 500060E802778702:500060E802778702) online.
Oct 22 00:16:18 ds13un jnic: [ID 229844 kern.notice] jni
Contact JNI support for more information.
```

Workaround: Add the following parameter to the JNI configuration file (`jnic.conf`):

```
FcEnableContextSwitch = 1;
```

I/O Policy Defaults to Balanced for Active/Active Enclosures Only

The man page for `vxddmpadm` states that *balanced* is the default I/O policy for enclosures. This is true only for Active/Active enclosures. Other types of enclosures, such as *Single-active do not default to *balanced*. [Sun Bug ID 4949454, i137411]



Sun StorEdge Traffic Manager (SSTM)

The Sun StorEdge Traffic Manager (SSTM) boot support feature that is available through SAN 4.3 is not supported with VxVM 4.0. Booting from the fabric devices as well as the boot encapsulation of fabric devices under SSTM is not supported in VxVM 4.1. [Sun Bug ID 4912232, 4909641, 4912667].

3510 Array

When a 3510 disk greater than 512GB is initialized to be a CDS disk, sector count is miscalculated and some disk space will be lost. This problem is only seen on 3510 array. [272241]

DMX 1000Array

If you have installed Solaris 10, VxVM cannot boot from an encapsulated DMX 1000 array. [273850]

Hitachi 9990 Genesis Array

After installing Storage Foundation 4.1 errors like the following are displayed on the console.

```
d18b-root@[ /usr/sbin]>d18b-root@[ /usr/sbin]>get_geometry_info_common:
solaris disk label adj. failed for /dev/vx/rdump//GENESIS0_6 (err
22)get_geometry_info_common: solaris disk label adj. failed for
/dev/vx/rdump//GENESIS0_6 (err 22)get_geometry_info_common: solaris
disk label adj. failed for /dev/vx/rdump//GENESIS0_6 (err
22)get_geometry_info_common: solaris disk label adj. failed for
/dev/vx/rdump//GENESIS0_6 (err 22)get_geometry_info_common: solaris
disk label adj. failed for /dev/vx/rdump//GENESIS0_6 (err
22)get_geometry_info_common: solaris disk label adj. failed for
dev/vx/rdump//GENESIS0_6 (err 22)
```

This failure has been observed on the Hitachi 9990 (Genesis) arrays where the disk geometry data is being handled incorrectly by `vxconfigd`, resulting in the indicated message during `vxdctl enable` or `vxconfigd` startup. This message does not affect VxVM use of the array. [Sun Bug ID 6221005, 301931, 308975]

Error Messages Seen During vxconfigd

When an internal Intelligent Drive Electronics (IDE) device is claimed, VxVM attempts to obtain geometry data for the device using `scsi` commands, which results in these messages. No data is lost, and VxVM claims the device correctly. This message may be seen each time `vxconfigd` is restarted and an internal IDE device is found in the configuration. [Sun Bug ID 6222054, 308336]

S-VOL Devices on HDS with TrueCopy Enabled

When using HDS with True Copy enabled, the primary devices (P-VOL) and their mirrors (S-VOL devices) are both seen in `vxdisk list` output. The P-VOL devices are available for import but the S-VOL devices are not available for import. Do not try to use S-VOL devices even though they appear in the `vxdisk list` output. [300979]

Hot-Relocation Issues

Impact of Hot-Relocation on Performance

Except for `rootvol` and `swapvol`, hot-relocation does not guarantee the same layout of data or performance after relocation. It is therefore possible that a single subdisk that existed before relocation may be split into two or more subdisks on separate disks after relocation (if there is not enough contiguous space on a single disk to accommodate that subdisk). [i14894]

Disk Information in Notification Messages

When a disk failure occurs, the hot-relocation feature notifies the system administrator of the failure and any relocation attempts through electronic mail messages. The messages typically include information about the device offset and disk access name affected by the failure. However, if a disk fails completely or a disk is turned off, the disk access name and device offset information is not included in the mail messages. This is because VxVM no longer has access to this information. [i14895]

DMP Issues

Non-persistence of I/O Policy

The previous release of VxVM allowed DMP's I/O policy to be changed. The changes were not persistent across booting or reconfigurations. To make the settings persistent, you may have modified the `/etc/init.d/vxvm-recover` script in the earlier release. This release supports persistent I/O policy and you may remove the changes to the script [i140947]



Usage of `dmp_failed_io_threshold` Parameter

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

- ◆ A release of VxVM between VxVM 3.2 patch 2 and VxVM 4.1 is installed.
- ◆ VxVM volumes are not mirrored.
- ◆ There are multiple paths to the disk.
- ◆ A portion of the I/O path below DMP fails in such a way that the error is not returned to DMP before 10 minutes have elapsed. Examples of failures that have been seen to cause this condition, include bad hardware (HBA, FCOT, GBIC, Switch Port, Array Controller) and HBA reconfigured to retry endlessly.

Note It is the failure condition in the fourth type of configuration that triggers the problem. These failures are rare and are not seen during the normal operation of a healthy SAN.

In VxVM 3.2 patch 2, a new tunable parameter was introduced to prevent devices from experiencing certain failure conditions that would prevent a mirrored I/O from succeeding for an extended period of time. This tunable parameter set a threshold of 600 seconds (10 minutes) for an I/O error to be returned from the device. If the I/O takes longer than 10 minutes to return with an error, DMP assumes that the device is not working and passes the error up to VxVM without retrying the operation. This allows VxVM to use a mirror of the data to satisfy the request without further delay.

Note This is not a time-out. No DMP activity will occur after the 10 minutes has passed. DMP only checks the elapsed time of the I/O after it is returned by the lower layer. If the elapsed time is greater than `dmp_failed_io_threshold` seconds (default 600), the error will be returned to VxVM without retries. DMP will wait as long as it takes for the I/O to be returned.

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

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

To prevent this possibility in situations where mirrored volumes are not used, the threshold should be tuned to a sufficiently high value that is unlikely to be reached. In the following example, 16 hours is used.

To change the value of *dmp_failed_io_threshold*, modify the value in */kernel/drv/vxdmp.conf*:

```
dmp_failed_io_threshold=57600
```

where 57600 introduces a delay of 16 hours (16 x 60 x 60).

After changing the value, reboot the system.

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

Controller State Change

Sometimes *vxnotify* will not receive events relating to the change of controller states in DMP. [272891]

Bringing Back a Failed Path

In a previous release a technote said that if a path fails, a user should stop the restore daemon and in a cluster environment and use the *vxctl enable* command to bring back a failed path.

In VxVM 4.1 the restore daemon should not be stopped in a cluster and the *vxctl enable* command is not needed.

If *vxctl enable* is used during a failover/failback process, some of the disks which do not belong to any disk groups might appear in an error state. A subsequent *vxctl enable* will correct the error. [298317]

Cluster Functionality Issues

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

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

[Sun Bug ID 4087612; i20448]



RAID-5 Volumes

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

File Systems Supported in Cluster-Shareable Disk Groups

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

Reliability of Information About Cluster-Shareable Disk Groups

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

Messages Caused by Open Volume Devices

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

```
cannot assign minor #
```

This message is accompanied by the console message:

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

Data Integrity Issues

Disks with Write-Back Caches

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

Usage of `dmp_failed_io_threshold` Parameter

See “[DMP Issues](#)” on page 63 for information on this issue.

Snapshot and Snapback Issues

Using Snapshots as Root Disks

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

Warning Message when Taking a Snapshot of a CFS File System

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

```
vxio: WARNING: vxvm:vxio: Plex plex detached from volume vol
```

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

File System Check of a Snapshot

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

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

Mount Operation Can Cause Inconsistencies in Snapshots

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

Intelligent Storage Provisioning Issues

Creating Application Volumes

To create application volumes successfully, the appropriate licenses must be present on your system. For example, you need a full VERITAS Volume Manager and a VERITAS FlashSnap license to use the instant snapshot feature. Vendors of disk arrays may also provide capabilities that require special licenses for certain features of their hardware. [Sun Bug ID 4948093, i137185]



Miscellaneous Issues

Auto-import of Disk Groups

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

If it is important that a disk group *not* be auto-imported when the system is rebooted, the disk group should be imported temporarily when the intention is to deport the disk group (for example, in HA configurations). Use the `-t` flag to `vxchg import`. [i13741]

Volumes Not Started Following a Reboot

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

Workaround: Check the state of the volumes before starting the application, or place a sleep (`sleep sec`) before the last invocation of `vxrecover`. [i14450]

Forcibly Starting a Volume

The `vxrecover` command starts a volume only if it has at least one plex that is in the ACTIVE or CLEAN state and is not marked STALE, IOFAIL, REMOVED, or NODAREC. If such a plex is not found, VxVM assumes that the volume no longer contains valid up-to-date data, so the volume is not started automatically. A plex can be marked STALE or IOFAIL as a result of a disk failure or an I/O failure. In such cases, to force the volume to start, use the following command:

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

However, try to determine what caused the problem before you run this command. It is likely that the volume needs to be restored from backup, and it is also possible that the disk needs to be replaced. [i14915]

Failure of Memory Allocation

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

Using Long Device Paths with Sun Online:Backup

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.

Workaround: Use symbolic links to the longer `/dev/vx/dsk/volname` paths from a shorter path name.

Messages About VVR Licenses

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

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

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

Number of Columns in a RAID-5 ISP Volume

If an ISP volume is created with the RAID-5 capability, the parameters `ncols` and `nmaxcols` refer only to the number of data columns, and do not include the parity column. For this reason, the number of columns that are created in such a volume is always one more than the number specified. [Sun Bug ID 4976891]

Solaris Issues

Compatibility of Kernel Drivers

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:

```
WARNING: vxio: incompatible kernel version (5.X), expecting 5.X
```

If this message is displayed, the system must be booted for recovery (as explained in the *VERITAS Volume Manager Troubleshooting Guide*) and the correct kernel modules installed. To install the correct kernel module versions, `cd` to the `kernel/drv` directory of the mounted root file system. To list the VxVM kernel modules, use the following command:

```
# ls -l vxio* vxspec* vxdmp*
```



The release-specific versions of the kernel modules are stored as `module.OS_release`, where `OS` and `release` are the result of running the `uname -s` and `uname -r` commands on the system, respectively.

For example, on a misconfigured system running Solaris 2.6, the listing for `vxio*` may be similar to the following:

```
-rw-r--r-- 1 root other 1682424 ... vxio
-rw-r--r-- 1 root sys   1647664 ... vxio.SunOS_5.7
-rw-r--r-- 1 root sys   1661340 ... vxio.SunOS_5.8
-rw-r--r-- 1 root sys   1682424 ... vxio.SunOS_5.9
```

The size of the `vxio` kernel module that is in use matches the `vxio.SunOS_5.8` version. To correct the problem, copy the `SunOS_5.6` versions to the in-use module name:

```
# cp vxio.SunOS_5.6 vxio
```

Finally reboot the system. [i13312]

Encapsulation of Swap Partitions

During encapsulation, VxVM does not consider a partition to be a swap partition unless its partition tag (as shown by `prtvtoc`) is `swap` or `3`. Any partition used as a swap partition but not tagged as such is encapsulated as a file system. In the `vfstab` file, a note is made that the partition has been encapsulated, but the `vfstab` entry *is not* translated, and thus, the partition is not added as a swap area as part of the boot process. All partitions used as swap must be marked with the `swap` tag to be properly encapsulated. [i13388]

Protection of Block 0 on Disks

Since the disk label is stored in block 0 of the disk, block 0 must not be used (that is, no application should write any information in block 0). Special protection has been built into VxVM to protect block 0 from being overwritten.

Failure of the `dd` Command with Large Volumes

The `dd` command in Solaris uses only `lseek()` to seek to a particular offset in a file. It does not use `lseek()`. This causes `dd` to fail on volumes greater than 2 gigabytes.

Definition of Disk Slice 2

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`.

Messages Caused by Long Swap Volume Names

If multiple swap partitions are encapsulated on your disks, VxVM names them as `swapvol`, `swapvol1`, `swapvol2`, and so on. When the system is rebooted, the following error message is displayed:

```
/dev/vx/dsk/swapvol2 : Overlapping swap files are not allowed
```

However, the swap devices are correctly added with no ill effects on the system. To avoid seeing this message, shorten the names of swap volumes (other than `swapvol`) from `swapvoln` to `swapn`.

VEA Issues

Note Refer to the *VERITAS Storage Foundation Installation Guide* for information on how to set up and start the VEA server and client.

Accessing the Task Log

The task log accessed from the Log tree is not supported. At this time, entries are written to the log file in `/var/vx/isis/command.log`. [i76683, i97076]

Permitting Remote Access to the X Windows Server

The following X Windows system error may occur when starting VEA:

```
Xlib: connection to "hostname:0.0" refused by server
Xlib: Client is not authorized to connect to Server
```

Workaround: Allow access to the local X server by using the following command:

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

Disk Group Creation Failure with Duplicate Disk ID

VEA fails to create a disk group with a duplicate disk ID, and gives no other options. [Sun Bug ID 4923820]



Printing Errors from VEA on Windows 2000 Service Pack 2

When a user tries to print the volume layout view from VEA, the print is not clear.

Workaround: Upgrade the printer device driver to 0.3.1282.1 and install Service Pack 3. Upgrade to the latest version of VEA and print again. [286476]

Creating non-CDS disks

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

1. Create the file `/etc/default/vxdisk`, and add a line describing the desired non-CDS disk format as:
`format=simple`
or
`format=sliced`
2. Create the file `/etc/default/vxdg`, and add a line describing the CDS attribute as:
`cds=off`

To create CDS disk groups at a later stage, you will need to change the settings back again to `format=cdsdisk` and `cds=on`. For more information, see the *VERITAS Storage Foundation Cross-Platform Data Sharing Administrator's Guide*.
[Sun Bug ID 4913804]

Internationalization Issues

Some ISP Attributes Have Not Been Translated

The Intelligent Storage Provisioning (ISP) window for annotating a disk is not fully localized. In particular, auto-discovered attributes such as DiskGroup and Enclosure are not translated. [i139162]

Inaccuracies in ISP Attribute Fields

The ISP User Template Wizard shows two “attribute value” fields rather than one “attribute value” and one “attribute name” field. [i139762]

Upgrading Disk Group Versions

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

VxVM Release	Cluster Protocol Versions	Disk Group Version	Supported Disk Group Versions
2.0	n/a	20	20
2.2	n/a	30	30
2.3	n/a	40	40
2.5	n/a	50	50
3.0	n/a	60	20-40, 60
3.1	n/a	70	20-70
3.1.1	10, 20	80	20-80
3.2	30	90	20-90
3.5	40	90	20-90
4.0	50	110	20-110
4.1	60	120	20-120



You can use the following command to find out the version number of a disk group:

```
# vxdg list disk_group_name
```

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

To upgrade a disk group, use the following command:

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

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

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

To see the current cluster protocol version, type:

```
# vxdctl support
```

Adding Unsupported Disk Arrays to the JBOD Category

Caution The procedure in this section ensures that dynamic multipathing is set up correctly on an array that is not supported by VxVM. Otherwise, VxVM treats the independent paths to the disks as separate devices, which can result in data corruption.

▼ To add an unsupported disk array after VxVM has been installed

1. Use the following command to identify the vendor ID and product ID of the disks in the array:

```
# /etc/vx/diag.d/vxdmpinq device_name
```

where *device_name* is the device name of one of the disks in the array (for example, `/dev/vdsk/c1t20d0s2`). Note the values of the vendor ID (VID) and product ID (PID) in the output from this command. For Fujitsu disks, also note the number of characters in the serial number that is displayed. The following is sample output:

```
# /etc/vx/diag.d/vxdmpinq /dev/rdsk/c1t20d0s2
```

```
Vendor id (VID) : SEAGATE
Product id (PID): ST318404LSUN18G
Revision       : 8507
Serial Number   : 0025T0LA3H
```

In this example, the vendor ID is SEAGATE and the product ID is ST318404LSUN18G.

2. Enter the following command to add a new JBOD category:

```
# vxddladm addjbod vid=vendorid pid=productid
[length=serialno_length]
```

where *vendorid* and *productid* are the VID and PID values that you found from the previous step. For example, *vendorid* might be FUJITSU, IBM, or SEAGATE. For Fujitsu devices, you must also specify the number of characters in the serial number as the argument to the *length* argument (for example, 10).

Note In VxVM releases 3.5 and greater, a SEAGATE disk is added as a JBOD device by default.

Continuing the previous example, the command to define an array of disks of this type as a JBOD would be:

```
# vxddladm addjbod vid=SEAGATE pid=ST318404LSUN18G
```

3. Enter the following command to bring the array under VxVM control:

```
# vxdctl enable
```

4. To verify that the array is now supported, enter the following command:

```
# vxddladm listjbod
```

The following is sample output from this command for the example array:

VID	PID	Opcode	Page Code	Page Offset	SNO	length
=====	=====	=====	=====	=====	=====	=====
SEAGATE	ALL PIDs	18	-1	36	12	

5. To verify that the array is recognized, use the `vxdmpadm listenclosure` command as shown in the following sample output for the example array:

```
# vxdmpadm listenclosure all
```

ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS
=====	=====	=====	=====
OTHER_DISKS	OTHER_DISKS	OTHER_DISKS	CONNECTED
Disk	Disk	DISKS	CONNECTED

The enclosure name and type for the array are both shown as being set to Disk. You can use the `vxdisk list` command to display the disks in the array:

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
Disk_0	auto:none	-	-	online invalid
Disk_1	auto:none	-	-	online invalid
...				



6. To verify that the DMP paths are recognized, use the `vxddmpadm getdmpnode` command as shown in the following sample output for the example array:

```
# vxddmpadm getdmpnode enclosure=Disk
NAME      STATE      ENCLR-TYPE  PATHS  ENBL    DSBL    ENCLR-NAME
=====
Disk_0    ENABLED    Disk        2      2        0      Disk
Disk_1    ENABLED    Disk        2      2        0      Disk
...
```

This shows that there are two paths to the disks in the array.

For more information, enter the command `vxddladm help addjbod`, or see the `vxddladm(1M)` and `vxddmpadm(1M)` manual pages.

VERITAS File System Software Issues

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.

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.

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

The `fscdtask 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 the next 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	<p>The following is a known Oracle issue:</p> <p>If you are using Oracle Disk Manager on Oracle 10g, attempts to recreate the control file will fail. The SQL statement CREATE CONTROLFILE REUSE SET DATABASE DATABASE_NAME RESETLOGS results in the following error message:</p> <pre>CREATE CONTROLFILE REUSE SET DATABASE "DBNAME" RESETLOGS ARCHIVELOG * ERROR at line 1: ORA-01503: CREATE CONTROLFILE failed ORA-00200: controlfile could not be created ORA-00202: controlfile: '/tmp/tmp34/test/TEST/DBNAME' ORA-17505: ksfdrsz:1 Failed to resize file to size 150 blocks</pre> <p>To address this issue, request the patch for Oracle bug #3512248 from Oracle support and apply it. Oracle bug #3512248 is a bug ID used internally at Oracle as a reference. It cannot be searched on to obtain bug information on Oracle's support website.</p>
none	<p>Installing Oracle 10g on Solaris 10 will fail.</p> <p>To work around this installation issue, run the following command:</p> <pre># ./runInstaller -ignoreSysPrereqs</pre> <p>This prevents the Oracle 10g installer from checking the platform.</p>
none	<p>If there are multiple processors, starting or creating an Oracle 10g database on Solaris 10 may fail. To work around this issue:</p> <ol style="list-style-type: none">1. Add the following to <code>init.ora</code>: <code>_enable_NUMA_optimization=FALSE</code>2. Ignore the following error message, which appears when executing SQL*Plus: <pre>\$ sqlplus "/ as sysdba" Assertion failed: 0, file skgsn.c, line 309</pre>3. You may now start or create the database.



Incident	Description
none	<p>When database files or control files consist of VxFS files and non-VxFS files, the <code>dbed_vmchecksnap</code> command does not read the configuration properly. When you run the <code>dbed_vmchecksnap</code> command in a mixed environment, the following error messages are displayed:</p> <pre>DBED4421: ERROR: File ufs is not on a vxfs file system. DBED5643: ERROR: An error occurred while finding device from file list. /dev/vx/dsk/proddg/datavol /usr1/oracle/rw/DATA proddg vxfs rw,suid,delaylog,largefiles,qio,ioerror=mwdisable,dev2c80e9</pre>
none	<p>If you install the product packages using the installation script, but do not reboot, and then decide to remove the packages using the removal script, removal of the <code>VRTSodm</code> and <code>VRTSvxfs</code> packages may fail due to the device being busy. If this happens, exit out of the removal script and use the <code>pkgrm</code> command to uninstall the packages.</p>
301174	<p>Column headings in output for the <code>vxstorage_stat</code> and <code>dbed_analyzer</code> commands are not localized. If you are running in an environment other than English, please note the following information regarding the column headings.</p> <p>For <code>vxstorage_stats</code>:</p> <ul style="list-style-type: none">◆ NAME = Name◆ TY = Type◆ NSUB = Number of subsystems◆ AVG TIME = Averaged time◆ B_READ = Number of blocks read◆ B_WRITE = Number of blocks written◆ AVG_RD = Average number of reads◆ AVG_WR = Average number of writes◆ OFFSET = Location of the object (in bytes)◆ PROPERTIES = Description of the object <p>For <code>dbed_analyzer</code>:</p> <ul style="list-style-type: none">◆ DEVICE = Device name◆ TBSNAME = Tablespace name◆ DATAFILE = Datafile name◆ SIZE = Size in sectors

Incident	Description
311221	<p>In a single host configuration, you may occasionally experience snapshot disk group deport failure when multiple concurrent Database FlashSnap processes are running. In this situation, an error message similar to the following is displayed:</p> <pre>ERROR V-81-5612 Deport <SNAP-dg-name> failed</pre> <p>The cause is due to a delay in VEA when obtaining the snapshot disk group status. To work around this issue, allow one minute between taking the snapshot (dbed_vmsnap) and starting a clone database (dbed_vmclonedb).</p>
34106	<p>Previously, converting a Quick I/O file to a regular file will fail if you created a Quick I/O file for an Oracle datafile that has a smaller size. Oracle will generate an error if the size of the regular file is not a multiple of the 2 K database block size (db_block_size). The error message displayed is similar to the following:</p> <pre>*** 2002-08-22 15:11:02.850 *** SESSION ID:(2.1) 2002-08-22 15:11:02.847 ORA-01157: cannot identify/lock data file 51 - see DBWR trace file ORA-01110: data file 51: '/d05/oradata/rmdb/bill113_data.dbf' ORA-27046: file size is not a multiple of logical block size Additional information: 1</pre> <p>To work around this problem, extend the Quick I/O file such that the size is a multiple of both the file system block size and the database block size (db_block_size)</p> <pre>\$ qiomkfile -r 2147500032 rbstest01.dbf \$ ls .rbstest01.dbf rbstest01.dbf -rw-r--r-- 1 oracle dba 2147500032 Apr 3 20:27 .rbstest01.dbf lrwxrwxrwx 1 oracle dba 26 Apr 3 20:27 rbstest01.dbf -> .rbstest01.dbf::cdev:vxfs: \$ rm rbstest01.dbf \$ mv .rbstest01.dbf rbstest01.dbf \$ ls rbstest01.dbf -rw-r--r-- 1 oracle dba 2147500032 Apr 3 20:27 rbstest01.dbf</pre> <p>The database will now start as expected.</p>
41192	<p>Depending on what you have highlighted in the object tree in the GUI, some tool bar icons are truncated and some are not visible. If this happens, use the menu bar to access the operation you want to execute. This will be fixed in a future release.</p>



Incident	Description
86431	You must run the <code>dbed_update</code> command after upgrading to VERITAS Storage Foundation 4.1 <i>for Oracle</i> from a previous release. This will allow you to roll back to a Storage Checkpoint that was created prior to this release.
103598	<p>Two problems with the Storage Mapping feature have been identified. They are as follows:</p> <ul style="list-style-type: none">• LUN to Physical Disk Deep Mapping does not work when VAIL is performing a rescan. After the rescan is complete, Deep Mapping resumes working normally. By default, VAIL performs a rescan every 30 minutes and can be configured through VEA.• GAP does not provide an offset value on physical disks for the LUN. The information is displayed as zero (0) when you run the <code>vxstorage_stats</code> command and when you obtain this information from the GUI. The information should be displayed as NA (not available).
217068	The <code>dbed_clonedb</code> and <code>dbed_vmclonedb</code> commands do not support clone databases with double-byte SID names.
231602	<p>If a database contains a symbolic link to Quick I/O files, Oracle 10g produces the following error:</p> <pre>ORA-27094: raw volume can not be used</pre> <p>To avoid this error, do not use symbolic links to Quick I/O files.</p> <p>This issue applies to Oracle version 10.1.0.2. It has been reported to Oracle and will be fixed in a future Oracle release.</p>
268200	VERITAS Storage Foundation <i>for Oracle</i> does not support the <code>db_recovery_file_dest</code> and <code>db_create_file_dest</code> parameters for Oracle 10g. The work around is to use the <code>log_archive_dest_n</code> and <code>control_files</code> parameters respectively instead.
271180	In VERITAS Storage Foundation <i>for Oracle</i> , EMC SymCLI does not work on Solaris 10. As a result, <code>vxstorage_stats</code> and <code>dbed_analyzer</code> are not supported on Solaris 10 at this point.
276354	<p>If a snapplan is in a state where <code>dbed_vmchecksnap -o remove</code> will not remove it, you may remove it by executing the following command, which forces removal of the snapplan from the repository:</p> <pre>rm -r /etc/vx/vxdba/ORACLE_SID/snapplan_dir/snapplan/*</pre> <p>where <code>ORACLE_SID</code> is the database name and <code>snapplan</code> is the name of the snapplan file you want to remove.</p>

Incident	Description
294135	The <code>dbed_analyzer</code> command fails and displays an error message when attempting to analyze an entire database if the database spans more than 100 tablespaces and each datafile has more than 2000 extents.
308325	For a database created with a temporary tablespace using Oracle Managed Files (OMF), the <code>dbed_vmc1onedb</code> command will fail.

VERITAS Storage Foundation for DB2 Software Issues

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

Incident	Description
25272	<p>If you convert a tablespace from a Quick I/O file to a regular file after backing up the database, you will not be able to restore the tablespace from that backup. For example, if you take a backup of a database that has a DMS tablespace with Quick I/O files as containers, and later convert the Quick I/O files to regular files, restoring the database from that backup will fail.</p> <p>Workaround: Use the <code>qio_recreate</code> command to re-create the necessary Quick I/O files before you restore the database.</p>
276354	<p>If a snapplan is in a state where <code>db2ed_vmchecksnap -o remove</code> will not remove it, you may remove it by executing the following command, which forces removal of the snapplan from the repository:</p> <pre>rm -r /etc/vx/vxdba/DB2.\$DB2INSTANCE.\$DB2DATABASE/NODE000/snapplan_dir/snapplan/*</pre> <p>where <code>\$DB2INSTANCE</code> is the instance name, <code>\$DB2DATABASE</code> is the database name, and <code>snapplan</code> is the name of the snapplan file you want to remove.</p>
296119	<p>There are no command options for the <code>qio_recreate_db2</code> command. However, <code>qio_recreate</code> has <code>-T</code> as an optional argument. You may use the <code>-T</code> option with <code>qio_recreate</code>, but not with <code>qio_recreate_db2</code>.</p>
303712	<p>The DB2 Database Partition Feature (DPF) does not fully support localization. As a result, the command output for <code>db2ed_checkconfig_all</code>, <code>db2ed_saveconfig_all</code>, and <code>db2ed_ckptcreate_all</code> may display garbage characters in the Japanese locale.</p>



Incident	Description
308441	<p>In the <i>VERITAS Storage Foundation for DB2 Administrator's Guide</i>, references to commands beginning with "dbed_" should begin with "db2ed_" instead.</p> <p>In addition, the occurrences of the word Oracle that appear on pages 89 and 437 are incorrect and may be disregarded.</p>
301174	<p>Column headings in output for the vxstorage_stat command are not localized. If you are running in an environment other than English, please note the following information regarding the column headings for vxstorage_stats:</p> <ul style="list-style-type: none">◆ NAME = Name◆ TY = Type◆ NSUB = Number of subsystems◆ AVG TIME = Averaged time◆ B_READ = Number of blocks read◆ B_WRITE = Number of blocks written◆ AVG_RD = Average number of reads◆ AVG_WR = Average number of writes◆ OFFSET = Location of the object (in bytes)◆ PROPERTIES = Description of the object
31604	<p>This release does not support spaces in container path names.</p>
34432	<p>VxFS disk layouts Version 5 and earlier do not display Storage Checkpoint quotas in the GUI.</p> <p>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:</p> <pre>DBED4646:ERROR: Unknown error code -30391108 (getquota) for DBED 46464</pre>
42112	<p>To use the Version Checkpoint feature in this release, no file system should be shared between two or more partitions. For example, to create file systems for creating a two-partition database on database directory /eedbhome/db2inst1, create at least two file systems as follows:</p> <pre># mount -F vxfs /dev/vx/dsk/\$DISKGROUP/NODE0000\ /eedbhome/db2inst1/NODE0000 # mount -F vxfs /dev/vx/dsk/\$DISKGROUP/NODE0001\ /eedbhome/db2inst1/NODE0001</pre>



Incident	Description
74780	In the DB2 GUI Rollback Wizard, the Rollback Buffer field format is not localized.
75222	In the Japanese version of this product, the Write Resume Time field for writable version checkpoints is always blank in the GUI.
76010	If you manually remove the VRTSjadb2 and VRTSjad2g Language Packages, the <code>install_lp</code> script may not detect they are missing if you run the script again. Workaround: Use the <code>pkgadd</code> command to manually install the packages.
86326	In the Display Container/File System Information on the VxDBA utility, the title and value in the output are misaligned. This occurs when the database container is created on a path with a long name.
86560	In the Display/Update Tablespace Information menu on VxDBA, container names are not aligned.
86650	When the DB2 instance is stopped, VxDBA displays an error message but does not indicate what the error is.
298101	If a disk is set up incorrectly, <code>db2ed_vmchecksnap -o validate</code> may give an incorrect disk name.

VERITAS Storage Foundation for Sybase Software Issues

There are no known issues in this release of VERITAS Storage Foundation *for Sybase*.



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.

Available Documentation

After the installation procedure is complete, documents are available online under the `/opt/VRTS/docs` directory. Documents are provided as Adobe Portable Document Format (PDF) files and in a searchable HTML-based format. To view or print PDF documents, you must have the Adobe Acrobat Reader installed.

Installing documentation and manual pages is optional.

VERITAS Documentation Disc

The VERITAS documentation disc provides searchable, HTML documentation for each product in this release. Printable PDF documents are also included on the disc.

All documentation is organized by product groups.

Release Notes and Installation Guides

Release notes and installation guides are not installed by any packages. VERITAS recommends that you copy them from the software disc to the `/opt/VRTS/docs` directory on your system after product installation so that they are available for future reference.

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 or the `cluster_server/release_notes` directory of the VERITAS Cluster Server disc.

Installation guides in all versions of the VERITAS Storage Foundation are located under the `storage_foundation/docs` directory of the VERITAS Storage Foundation disc or the `cluster_server/docs` directory of the VERITAS Cluster Server disc.

It is important that you read the relevant component product release notes before installing any version of VERITAS Storage Foundation:

- ◆ *VERITAS Volume Replicator Release Notes* (`vvr_notes.pdf`)
- ◆ *VERITAS Cluster Server Release Notes* (`vcs_notes.pdf`)

VERITAS Storage Foundation Guides

The following manuals, along with the online help, comprise the VERITAS Storage Foundation documentation set:

Guides in VERITAS Storage Foundation Documentation Set

Guide Title	Filename
<i>VERITAS Storage Foundation and High Availability Getting Started Guide</i>	
<i>VERITAS Storage Foundation Release Notes</i> (this document)	<code>sf_notes.pdf</code>
<i>VERITAS Storage Foundation Installation Guide</i>	<code>sf_install.pdf</code>



Guides in VERITAS Storage Foundation Documentation Set (continued)

Guide Title	Filename
<i>VERITAS Storage Foundation for Oracle Administrator's Guide</i>	<code>sf_ora_admin.pdf</code>
<i>VERITAS Storage Foundation for DB2 Administrator's Guide</i>	<code>sf_db2_admin.pdf</code>
<i>VERITAS Storage Foundation for Sybase Administrator's Guide</i>	<code>sf_syb_admin.pdf</code>
<i>VERITAS Array Integration Layer Configuration Guide</i>	<code>vail_config.pdf</code>
<i>VERITAS Volume Manager Administrator's Guide</i>	<code>vxvm_admin.pdf</code>
<i>VERITAS Storage Foundation Intelligent Storage Provisioning Administrator's Guide</i>	<code>sf_isp_admin.pdf</code>
<i>VERITAS Storage Foundation Cross-Platform Data Sharing Administrator's Guide</i>	<code>sf_cds_admin.pdf</code>
<i>VERITAS Enterprise Administrator (VEA 500 Series) Getting Started</i>	<code>vea5x_getting_started.pdf</code>
<i>VERITAS Volume Manager Troubleshooting Guide</i>	<code>vxvm_tshoot.pdf</code>
<i>VERITAS Volume Manager Hardware Notes</i>	<code>vxvm_hwnotes.pdf</code>
<i>VERITAS FlashSnap Point-In-Time Copy Solutions Administrator's Guide</i>	<code>flashsnap_admin.pdf</code>
<i>VERITAS File System Administrator's Guide</i>	<code>vxfs_admin.pdf</code>
<i>VERITAS File System Programmer's Reference Guide</i>	<code>vxfs_ref.pdf</code>

Note In this release, some of the above documents have new PDF file names.

VERITAS Cluster Server Documentation

The following VERITAS Cluster Server documentation is available with all VERITAS Storage Foundation HA product suites:

Guides in VERITAS Cluster Server Documentation Set

Guide Title	Filename
<i>VERITAS Cluster Server Release Notes</i>	<i>vcs_notes.pdf</i>
<i>VERITAS Cluster Server Installation Guide</i>	<i>vcs_install.pdf</i>
<i>VERITAS Cluster Server User's Guide</i>	<i>vcs_users.pdf</i>
<i>VERITAS Cluster Server Agent Developer's Guide</i>	<i>vcs_agent_dev.pdf</i>
<i>VERITAS Cluster Server Bundled Agents Reference Guide</i>	<i>vcs_bundled_agents.pdf</i>
<i>VERITAS Cluster Server SunFire 12K/15K Application Note</i>	<i>vcs_appnote_f15k.pdf</i>
<i>VERITAS Cluster Server SunFire 6800 Application Note</i>	<i>vcs_appnote_6800.pdf</i>
<i>VCS Enterprise Agent for Oracle Installation and Configuration Guide</i>	<i>vcs_oracle_install.pdf</i>
<i>VCS Enterprise Agent for DB2 Installation and Configuration Guide</i>	<i>vcs_db2_install.pdf</i>
<i>VCS Enterprise Agent for Sybase Installation and Configuration Guide</i>	<i>vcs_sybase_install.pdf</i>
<i>VCS Enterprise Agent for SunONE Installation and Configuration Guide</i>	<i>vcs_sunone_install.pdf</i>
<i>VCS Enterprise Agent for EMC SRDF Installation and Configuration Guide</i>	<i>vcs_srdf_install.pdf</i>
<i>VCS Enterprise Agent for IBM PPRC Installation and Configuration Guide</i>	<i>vcs_pprc_install.pdf</i>



VERITAS Volume Replicator Documentation

The following VERITAS Volume Replicator documentation is available with the VERITAS Volume Replicator option:

Guides in VERITAS Volume Replicator Documentation Set

Guide Title	Filename
<i>VERITAS Volume Replicator Release Notes</i>	<code>vvr_notes.pdf</code>
<i>VERITAS Volume Replicator Installation Guide</i>	<code>vvr_install.pdf</code>
<i>VERITAS Volume Replicator Administrator's Guide</i>	<code>vvr_admin.pdf</code>
<i>VERITAS Volume Replicator Planning and Tuning Guide</i>	<code>vvr_planning.pdf</code>
<i>VERITAS Volume Replicator Web Console Administrator's Guide</i>	<code>vvr_web_admin.pdf</code>
<i>VERITAS Volume Replicator Advisor User's Guide</i>	<code>vvr_advisor_users.pdf</code>
<i>VERITAS Cluster Server Agents for VERITAS Volume Replicator Configuration Guide</i>	<code>vvr_agents_config.pdf</code>

Manual Pages

The VERITAS manual pages are installed in the `/opt/VRTS/man` directory. This directory must be added to the `MANPATH` environment variable.

Installing documentation and manual pages is optional.

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>.



