

VERITAS Storage Foundation™ 4.1

Release Notes

Solaris x64 Platform Edition

Maintenance Pack 1

April 2006

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 © 2006 VERITAS Software Corporation. All rights reserved. VERITAS, the VERITAS Logo and VERITAS Storage Foundation 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 Street
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

Product Descriptions	2
Storage Foundation Basic	3
VERITAS Volume Manager	3
VERITAS File System	4
VERITAS Volume Replicator	4
VERITAS Cluster Server	5
New and Enhanced Features	5
VERITAS Storage Foundation	5
VERITAS Volume Manager	5
VERITAS File System	8
System Requirements	10
Solaris Operating System Requirements	10
Solaris Patch Requirements	11
VERITAS Patches	12
Installation Storage Foundation HA Configuration Requirements	12
Hardware Configuration Information	13
Software Limitations	15
VERITAS Volume Manager Software Limitations	15
VERITAS File System Software Limitations	16
VERITAS Storage Foundation for Databases Software Limitations	18
Unsupported Features	21
Hardware Issues	22
Software Issues	23



Installation Script Issue	23
License Upgrade Issue	23
Upgrading the VxVM Manual Page and Documentation Packages	23
Incorrect Reference to veritas_enabled Directory in the VERITAS Volume Manager Hardware Notes	25
VERITAS Volume Manager Software Issues	26
VERITAS File System Software Issues	39
VERITAS Storage Foundation for Databases Software Issues	39
Documentation Issues	40
No Longer Supported	40
Storage Foundation QuickStart	40
Available Documentation	41
Release Notes and Installation Guides	41
VERITAS Storage Foundation Guides	42
Online Manual Pages	44
Getting Help	45



VERITAS Storage Foundation Release Notes

This document provides release information about the VERITAS Storage Foundation 4.1 Maintenance Pack 1 (MP1) software for the Solaris 10 Operating System (OS) on the x64 platform.

This product suite is activated by a license key, which must be obtained 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/278691>

Review this entire document before installing the VERITAS Storage Foundation software.



Product Descriptions

This section lists the components of each of the Storage Foundation 4.1 MP1 product suites.

Storage Foundation Version	Products
Storage Foundation Basic	VERITAS Volume Manager VERITAS File System
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 for Oracle	VERITAS Volume Manager VERITAS File System
Storage Foundation for Sybase	VERITAS Volume Manager VERITAS File System



Storage Foundation Basic

For a product overview of Storage Foundation Basic products, see the *VERITAS Storage Foundation and High Availability Solutions Getting Started Guide*.

VERITAS Volume Manager

VERITAS Volume Manager (VxVM) is a storage management tool that manages physical disks as logical device volumes, removing the limitations of physical disk storage partitions.

VxVM 4.1 Licenses

The following features are supported by a Full VERITAS Volume Manager 4.1 license (new features in this release are also indicated):

- ◆ Clustering feature of VERITAS Volume Manager (new)
- ◆ Co-existence of dynamic multipathing (DMP) with other multipathing drivers such as MPxIO
- ◆ DMP for Active/Active and Active/Passive (new) arrays
- ◆ Concatenation
- ◆ Cross-platform Data Sharing (CDS)
- ◆ Device Discovery Layer (DDL)
- ◆ Dirty region logging (DRL) for mirrors
- ◆ Dynamic LUN expansion
- ◆ Extensible Firmware Interface (EFI) disks (new)
- ◆ FlashSnap™ features (new):
 - ◆ FastResync (fast mirror resynchronization)
 - ◆ Disk group split/join (DGSJ)
- ◆ Full-sized and space-optimized instant snapshots (new)
- ◆ Hot-relocation
- ◆ Hot sparing
- ◆ Intelligent Storage Provisioning
- ◆ Mirroring
- ◆ Mirroring plus striping



- ◆ Multiple disk groups
- ◆ Online data migration
- ◆ Online relayout
- ◆ RAID-5
- ◆ RAID-5 logging
- ◆ Smartsync
- ◆ Spanning
- ◆ Striping
- ◆ Striping plus mirroring
- ◆ Third-mirror break-off volume snapshots
- ◆ VERITAS Enterprise Administrator (VEA)
- ◆ Volume resizing

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

```
# vxctl license
```

VERITAS File System

VERITAS File System (VxFS) is a high performance file system with online management capabilities. File systems are a collection of directories organized into a structure that enable you to locate and store files. All information is eventually stored in a file system.

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 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, 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.

Clustering Feature

The cluster functionality of VxVM (CVM) allows up to 16 nodes in a cluster to simultaneously access and manage a set of disks under VxVM control (VM disks) using VERITAS Cluster Server (VCS). The same logical view of disk configuration and any changes to this is available on all the nodes. When the cluster functionality is enabled, all the nodes in the cluster can share VxVM objects.

See the *VERITAS Volume Manager Administrator's Guide* for details.



Co-existence of DMP with Third-Party Multipathing Drivers

The third-party driver (TPD) coexistence feature allows I/O that is controlled by third-party multipathing drivers to bypass DMP while retaining the monitoring capabilities of DMP. Provided that a suitable array support library (ASL) is available, devices that use TPDs can be discovered without requiring a specification file to be set up, a special command to be run, or the third-party multipathing driver to be modified.

See the *VERITAS Volume Manager Administrator's Guide* for details.

DMP Support for Active/Passive Arrays

The Dynamic Multipathing (DMP) feature supports Active/Passive arrays in this release.

An *Active/Passive* array (*A/P array*) allows access to its LUNs (*logical units*; real disks or virtual disks created using hardware) via the *primary* (active) path on a single controller (also known as an *access port* or a *storage processor*) during normal operation.

In *implicit failover mode* (or *autotrespass mode*), an A/P array automatically fails over by scheduling I/O to the *secondary* (passive) path on a separate controller if the primary path fails. This passive port is not used for I/O until the active port fails. In A/P arrays, path failover can occur for a single LUN if I/O fails on the primary path.

In Active/Passive arrays with *LUN group failover* (*A/PG arrays*), a group of LUNs that are connected through a controller is treated as a single failover entity. Unlike A/P arrays, failover occurs at the controller level, and not for individual LUNs. The primary and secondary controller are each connected to a separate group of LUNs. If a single LUN in the primary controller's LUN group fails, all LUNs in that group fail over to the secondary controller.

Active/Passive arrays in *explicit failover mode* (or *non-autotrespass mode*) are termed *A/PF arrays*. DMP issues the appropriate low-level command to make the LUNs fail over to the secondary path.

A/P-C, A/PF-C and A/PG-C arrays are variants of the A/P, AP/F and A/PG array types that support concurrent I/O and load balancing by having multiple primary paths into a controller. This functionality is provided by a controller with multiple ports, or by the insertion of a SAN hub or switch between an array and a controller. Failover to the secondary (passive) path occurs only if all the active primary paths fail.

See the *VERITAS Volume Manager Administrator's Guide* for details.

Dynamic LUN Expansion

The `vxdisk resize` command can be used to make VxVM aware of the new size of a virtual disk device that has been resized. This facility is provided to support dynamic LUN expansion by updating disk headers and other VxVM structures to match the new size. It does not resize the LUN itself. The LUN device must have a SCSI interface that is presented by a smart switch, smart array or RAID controller. Following a `vxdisk resize` operation to increase the length that is defined for a device, additional disk space on the device is available for allocation.

See the *VERITAS Volume Manager Hardware Notes* and the `vxdisk(1M)` manual page for more information.

Extensible Firmware Interface Disks

The Solaris 10 operation system provides support for disks with Extensible Firmware Interface (EFI) labels. VERITAS Volume Manager supports such disks with the limitation that they may not be formatted for use with the Cross-Platform Data Sharing (CDS) feature. Note that the `c#t#d#s2` device (corresponding to the `s2` slice) on EFI disks is just another partition, and does not represent the entire disk. The `c#t#d#` device can be used to access the entire disk.

See the Solaris operating system documentation for more information.

FlashSnap

VERITAS FlashSnap offers a flexible and efficient means of managing business critical data. It allows you to capture an online image of actively changing data at a given instant: a *point-in-time copy*. You can perform system backup, upgrade and other maintenance tasks on point-in-time copies while providing continuous availability of your critical data. If required, you can offload processing of the point-in-time copies onto another host to avoid contention for system resources on your production server.

Two kinds of point-in-time copy solution are supported by the FlashSnap license:

- ◆ Volume-level solutions are made possible by the Persistent FastResync and Disk Group Split/Join features of VERITAS Volume Manager. These features are suitable for implementing solutions where the I/O performance of the production server is critical.
- ◆ File system-level solutions use the Storage Checkpoint feature of VERITAS File System.

See the *VERITAS FlashSnap Point-in-Time-Copy Solutions Administrator's Guide* for details.



Full-Sized and Space-Optimized Instant Snapshots

Full-sized instant snapshots and space-optimized instant snapshots offer advantages over traditional third-mirror break-off snapshots such as immediate availability and easier configuration and administration. You can also use a third-mirror break-off usage model with full-sized snapshots, where this is necessary for write-intensive applications.

See the *VERITAS Volume Manager Administrator's Guide* for details.

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 both as direct mount and a loopback file system (LOFS) in local zones. Concurrent I/O and ODM are supported in local zones. See [“Non-Global Zone Support”](#) on page 15.

Software Developer's Kit

The VERITAS File System Software Developer's Kit (SDK) 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 MP1 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 `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.
- ◆ File Promotion
- ◆ Fileset Quotas
- ◆ FlashSnap™ feature (new):
 - ◆ Storage Checkpoint feature of VERITAS File System
- ◆ Solaris Local Zones
- ◆ Oracle Disk Manager (ODM)
- ◆ Quality of Storage Service
- ◆ Quick I/O and cached Quick I/O features of VERITAS File System
- ◆ VERITAS File System Storage Checkpoints
- ◆ VERITAS File System Storage Checkpoints Rollback
- ◆ VERITAS Storage Foundation Cluster File System

System Requirements

Solaris Operating System Requirements

The VERITAS Storage Foundation Standard 4.1 MP1 product for the x64 platform operates on the Solaris 10 (64-bit) Operating System.

The VERITAS Storage Foundation High Availability (HA) 4.1 MP1 product for the x64 platform operates on the Solaris 10 (64-bit) Operating System.

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

Solaris Patch Requirements

Solaris Patches for Storage Foundation

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, contact your vendor.

Sun Patch Number (or higher)	Description	Notes
119255-08	SunSolve patch for install and patch utility	Obtain from http://sunsolve.sun.com
119375-03	SunSolve patch for SCSI disk driver	Obtain from http://sunsolve.sun.com
119131-09	SunSolve patch for Fibre Channel driver	Obtain from http://sunsolve.sun.com
118344-03	SunSolve patch for Fault Manager	Obtain from http://sunsolve.sun.com
118844-27	Sunsolve patch which addresses kernel issues found since the initial Solaris 10 release.	You need to install patch 118844-27 or later before installing the VERITAS Storage Foundation. Obtain from http://sunsolve.sun.com

See the `patchadd(1M)`, `patchrm(1M)` and `showrev(1M)` manual pages for more information about administering patches.

Note VERITAS highly recommends that your system must be running the above Solaris OS patch numbers for Storage Foundation products to work properly.



VERITAS Patches

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

Patch ID	Product	Package
119753-xx	VERITAS File System Management Services Provider	VRTSfspro
120586-xx	VERITAS Volume Manager (Solaris 10)	VRTSvxvm
120586-xx	VERITAS Volume Manager documentation	VRTSvmdoc
120586-xx	VERITAS Volume Manager manual pages	VRTSvmman
120853-xx	VERITAS Intelligent Storage Provisioning	VRTSalloc
120854-xx	VERITAS Device Discovery Layer Services Provider	VRTSddlpr
120111-xx	VERITAS File System (Solaris 10)	VRTSvxfs
120112-xx	VERITAS File System manual pages and documentation	VRTSfsdoc
120112-xx	VERITAS File System manual pages and documentation	VRTSfsman
120112-xx	VERITAS File System manual pages and documentation	VRTSfsmnd
121750-xx	VERITAS Volume Manager Management Services Provider	VRTSvmpro

Installation Storage Foundation HA Configuration Requirements

For Storage Foundation HA, you must have `rsh` communication enabled between the systems. Configure the system where you run the installer program for `rsh` communication with all systems in the cluster, including itself.



Hardware Configuration Information

These are the minimum requirements.

Systems

System	Platform BIOS ID	Platform BIOS Version	Diagnostics Server Version	Operator Panel Firmware Version	SP Value-Add Software	SP Base Software
V20z AMD Opteron	255	V1.30.5	V2.1.0.16	V1.0.1.1	V2.1.0.16	V2.1.0.16
V40z AMD Opteron	239	V2.33.7.2*	V2.1.0.16	V1.0.1.1	V2.3.0.15*	V2.3.0.15*

* Download Network Share Volume (NSV) 2.3.0.11C from:
<http://www.sun.com/servers/entry/v20z/downloads.html>

Arrays

Array	Firmware Version
StorEdge 3310	3.25W or above
StorEdge 3510	4.11I or above

For additional supported arrays, see the VERITAS Technical Supported site, <http://support.veritas.com>. Select “Volume Manager” from the Product Family menu and “Volume Manager on Unix” from the Product menu. Then click on Knowledge Base Search and enter “asl” followed by the vendor or model name.



Host Bus Adaptors

HBA	Firmware Version
Ultra320 Dual PCI SCSI	1.3.24.0
PCI Dual-Port FiberChannel 2Gb	3.2.112

Other Devices

Device	Version
StorEdge Configuration Command Line Interface (sccli)	1.6.2



Software Limitations

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

VERITAS Volume Manager Software Limitations

Root Disk Encapsulation

Root disk encapsulation is not implemented in Volume Manager 4.1 MP1 for Solaris 10 on the x64 platform. If required, you can use Solaris Volume Manager (SVM) to manage the root partitions of your systems. For details, see the attachment to these release notes.

It is recommended that you exclude the root disk and other disks that are under SVM control from the view of VERITAS Volume Manager. To do this, select the item *Prevent multipathing/Suppress devices from VxVM's view* in the main menu of the `vxdiskadm` command. See “Disabling and Enabling Multipathing for Specific Devices” in the “Administering Dynamic Multipathing (DMP)” chapter of the *VERITAS Volume Manager Administrator's Guide* for details.

Using the `vxcdsconvert` Utility

The `vxcdsconvert` utility is used to make a disk or disk group compatible with the CDS feature. You must specify the `-o novolstop` option to this command on the Solaris x64 platform because the default disk layout is different from that on the Solaris SPARC platform.

Non-Global Zone Support

VERITAS Volume Manager returns an error if you attempt to run it in a non-global zone.

Temporary File Systems Used by DMP

Two temporary file systems are mounted for use by DMP. When mounted, output from the `df` command is similar to the following:

Filesystem	Kbytes	Used	avail	Capacity	Mounted on
...					
swap	902488	0	902488	0%	/dev/vx/rdmp
swap	902488	0	902488	0%	/dev/vx/dmp



VERITAS File System Software Limitations

Local Zone Support

The following features are supported with local zones:

- ◆ 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)
```

No Support for Certain Manual Page options

The following manual pages options are not supported:

Manual Pages	Unsupported Options
mkfs_vxfs(1M)	The <code>version</code> option does not support disk layout Version 4 and 5.
mount_vxfs(1M)	qlog



VERITAS Storage Foundation for Databases Software Limitations

Oracle and Sybase Software Limitations

No Support for Local Zones

The Standard, Standard HA, Enterprise, and Enterprise HA versions of VERITAS Storage Foundation *for Oracle* 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* 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* VxFS disk layouts Version 5 and earlier do not display Storage Checkpoint quotas in the GUI. [34432 and 34433]

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

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

VERITAS Storage Foundation for Oracle Software Limitations

Storage Checkpoint Limitations

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

Database Cloning Limitation

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

VEA and VxDBA Menu Utility Limitations

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

VxDBA Monitoring Agent Limitation

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

Database FlashSnap Limitations

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

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

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

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

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



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

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

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

Oracle Disk Manager Limitations

- ◆ If you want to run Oracle 10g on a local zone and use Oracle Disk Manager, the Oracle version should be 10.1.0.3 or higher. To enable Oracle Disk Manager file access from local zones with VERITAS File System, follow the procedure in “[Enabling ODM Access](#)” on page 16.
- ◆ Because Oracle Disk Manager uses the Quick I/O driver to perform asynchronous I/O, do not turn off the Quick I/O mount option, which is the default.
- ◆ Using Oracle Disk Manager with Cached Quick I/O enabled is not supported and could cause your system to panic [34281]. To avoid a system panic, ensure the following:
 - ◆ If you are using Oracle Disk Manager, do not enable Cached Quick I/O on your file system.
 - ◆ If you are converting from Quick I/O to Oracle Disk Manager, make sure you disable Cached Quick I/O.

Oracle Managed Files Limitation

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

Unsupported Features

The product documentation, manual pages and online help contain references to features that are not supported in the 4.1 MP1 release on the x64 platform, but which are supported on the SPARC platform. These features are:

- ◆ Data Management API (DMAPI)
- ◆ Disk layout Version 5 and previous versions
- ◆ Extended Copy (ECopy)
- ◆ Root disk encapsulation and associated features
- ◆ Quick Log
- ◆ Storage Expert
- ◆ Support for moving between disk group versions
- ◆ `vxspcshow` command



Hardware Issues

Potential Data Corruption with the LSI Logic 1030 Controller

Writing data to a raw disk device (including VxVM volumes) can result in data corruption. It has been asserted by LSI that the LSI 53C1020 revision B2 chip may cause mis-reads when split transactions occur while in PCI-X mode. Systems with the later LSI 53C1020 version C0 have been found to be error-free.

This is considered to be a corner case. It is unusual for applications to write directly to a raw drive using offset addressing. Typical applications are associated with block devices with defined boundaries.

To check for the LSI 53C1020 chip revision in your Solaris system, enter:

```
prtconf -vp
```

Review the output under the LSI Logic section for the revision ID. A revision ID of 00000007 indicates revision B2, which is the faulty version, for example:

```
revision-id: 00000007
```

A revision ID of 00000008 indicates revision C0, in which the bug is fixed, for example:

```
revision-id: 00000008
```

Sun Microsystems has a hardware workaround, which is a minor modification to the PCI riser PCB that does not prevent the use of the PCI slot. The modification forces the slot and embedded devices to run in 66MHz PCI Conventional mode instead of 66MHz PCI-X mode. The transaction efficiencies gained in PCI-X mode will be lost, but the slot and embedded devices will be fully functional, running at 66MHz clock speed.

For more information, see the following TechNote on the VERITAS Technical Support website:

<http://support.veritas.com/docs/279169>

Contact Sun Microsystems for workaround procedures. [Sun Bug ID 6269806]

Software Issues

Installation Script Issue

When installing Storage Foundation Standard HA on a local system, you may encounter a CPI error message indicating you cannot install on the system:

```
CPI ERROR V-9-10-1566 Cannot install Storage Foundation Standard HA
on system as its OS kernel level is permission denied-bit which is
not supported with this release.
```

Workaround: Enable the local system for rsh communication to itself. Add the system name to its `.rhosts` file.

License Upgrade Issue

If you are upgrading a system by adding a license for previously unlicensed features (for example, when upgrading from Storage Foundation Basic to Storage Foundation Enterprise), first use the `vxinstall` command to install the new license key, and then use the `vxctl license init` command to make VxVM enable the new license key.

Upgrading the VxVM Manual Page and Documentation Packages

The VxVM patch 120586-03 updates the `VRTSvxvm`, `VRTSvmman` and `VRTSvmdoc` packages. However, due to a limitation in the installation script, the `VRTSvmman` and `VRTSvmdoc` packages are not updated during an upgrade.

There are two possible workarounds depending on whether or not you have already upgraded the VERITAS Storage Foundation software.

▼ To upgrade the `VRTSvmman` and `VRTSvmdoc` packages before upgrading the VERITAS Storage Foundation software

1. Create a temporary directory for the packages (in this example, `/tmp/vmpkgs`).

```
# mkdir /tmp/vmpkgs
```

2. Change directory to the packages directory on the mounted VERITAS Storage Foundation disc (in this example, under `/cdrom/cdrom0`).

```
# cd /cdrom/cdrom0/storage_foundation/pkgs
```



3. Copy the VRTSvmman and VRTSvmdoc packages to the temporary directory.

```
# cp VRTSvmman.tar.gz VRTSvmdoc.tar.gz /tmp/vmpkgs
```

4. Change directory to the temporary directory, and uncompress the packages.

```
# cd /tmp/vmpkgs
# /usr/bin/gzcat VRTSvmman.tar.gz | tar xvf -
# /usr/bin/gzcat VRTSvmdoc.tar.gz | tar xvf -
```

5. Run the pkgadd command on the packages.

```
# pkgadd VRTSvmman VRTSvmdoc
```

6. Upgrade the VERITAS Storage Foundation software as described in the *VERITAS Storage Foundation Installation Guide*.

7. After the upgrade is complete, you can check the version of the VxVM documentation packages by using the showrev command:

```
# showrev -p | grep 120586
Patch: 120586-03 Obsoletes: Requires: Incompatibles: Packages:
VRTSvmdoc, VRTSvmman, VRTSvxvm
```

▼ **To upgrade the VRTSvmman and VRTSvmdoc packages after upgrading the VERITAS Storage Foundation software**

1. Upgrade the VERITAS Storage Foundation software as described in the *VERITAS Storage Foundation Installation Guide*.

2. After the upgrade is complete, you can check the version of the VxVM documentation packages by using the showrev command:

```
# showrev -p | grep 120586
atch: 120586-03 Obsoletes: Requires: Incompatibles: Packages:
VRTSvxvm
```

Only the VRTSvxvm package should be shown as having been patched.

3. Create a temporary directory for the patch file (in this example, /tmp/vmpatch).

```
# mkdir /tmp/vmpatch
```

4. Change directory to the patches directory on the mounted VERITAS Storage Foundation disc (in this example, under /cdrom/cdrom0).

```
# cd /cdrom/cdrom0/storage_foundation/patches
```

5. Copy the patch 120586-03 file to the temporary directory.

```
# cp 120586-03.tar.gz /tmp/vmpatch
```

6. Change directory to the temporary directory, and uncompress the patch file.

```
# cd /tmp/vmpatch
# /usr/bin/gzcat 120586-03.tar.gz | tar xvf -
```

7. Run the patchadd command as shown here.

```
# patchadd -G -M . 120586-03
```

8. After the patch has been added, you can check the version of the VxVM documentation packages by using the showrev command:

```
# showrev -p | grep 120586
Patch: 120586-03 Obsoletes: Requires: Incompatibles: Packages:
VRTSvmdoc, VRTSvmman, VRTSvxvm
```

Should you wish to remove patch 120586-03, use the following command:

```
# patchrm -G 120586-03
```

Caution Do not omit the -G option. Otherwise error messages are displayed, and you will only be able to remove the patch if all three packages (VRTSvmdoc, VRTSvmman and VRTSvxvm) are removed.

Incorrect Reference to veritas_enabled Directory in the VERITAS Volume Manager Hardware Notes

The `veritas_enabled` directory was removed from the VERITAS software disc in the 4.1 MP1 release, so the procedure for "Adding an ASL Package" in the *VERITAS Volume Manager Hardware Notes* is no longer correct. To obtain an ASL package, you must access the Hardware Compatibility List (HCL) that is available on the VERITAS Support website, <http://support.veritas.com>. Select "Volume Manager" from the Product Family menu and "Volume Manager on Unix" from the Product menu. Then click on Knowledge Base Search and enter "asl" followed by the vendor or model name. Download the package, and then follow the procedure for adding the ASL package as detailed in the *VERITAS Volume Manager Hardware Notes*.



VERITAS Volume Manager Software Issues

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

VERITAS Enterprise Administrator Issue

Do not run the VEA server in the English (C) locale if the language pack is installed.

You can run the VEA client in the Japanese (ja) locale or English (C) locale if the VEA server is running in the Japanese (ja) locale. [398771]

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 `vxdmp`, `vxio` and `vxspec` (to guarantee data availability during the firmware upgrade procedure to a disk participating in a RAID mirror configuration).

For example, 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=first_cntlr
```

```
# /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.

Utility Issues

Reformatting a CDS disk can cause data corruption

On disks that are initialized by VxVM as CDS disks (the default format), the CDS information occupies the first sector of that disk, and there is no `fdisk` partition information. Attempting to create an `fdisk` partition (for example, by using the `fdisk` or `format` commands) erases the CDS information, and can cause data corruption.

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 ...
```

[13488]

Using `vxdiskadm` to Replace a Failed Disk

The `vxdiskadm` command requires two attempts to replace a failed disk. The first attempt can fail with a message of the form:

```
/usr/lib/vxvm/voladm.d/bin/disk.repl: test: argument expected
```

The command is not completed and the disk is not replaced. If you now rerun the command, using Option 5, the replacement successfully completes. [102381]

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 also 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
```

[134932]

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). [141024]

Do Not Specify a Long Device Name in `/etc/vx/disks.exclude`

Editing the `/etc/vx/disks.exclude` file to specify a long device name can cause scripts such as `vxdiskadm` to fail with an error message. Use `vxdiskadm(1M)` options 17 and 18 to suppress or unsuppress devices from VxVM's view. [Sun Bug ID 6228464, 311275]

Device Issues

Supported Hardware

See “[Hardware Configuration Information](#)” on page 13 for a list of supported systems, host bus adapters, disk arrays and other devices. This supplements the information given in the *VERITAS Volume Manager Hardware Notes*. Other array hardware and third-party multipathing driver coexistence are not supported in this release, and any mention of such should be treated as examples only.

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]

Miscalculation of Disk Space

If a LUN in a Sun SE3510 disk array is initialized to be a CDS disk, the sector count is miscalculated if the LUN is larger than 512GB. As a result, some disk space is lost. [272241]

Removing a Disk Hangs VxVM

VxVM hangs if a disk is removed from a Sun SE3510 disk array when active I/O is in progress. This is a known issue with RAID hardware configurations. A fix for the array firmware is currently being investigated.

To remove disk drives from such arrays without causing VxVM to hang, follow the instructions in the section “Removing and Replacing Disks” of the “Administering Disks” chapter in the *VERITAS Volume Manager 4.1 Administrator's Guide*. [400588]

Adding Disks to VxVM on a Solaris x64 System

The default disk layout on the Solaris x64 platform differs from that on the Solaris SPARC platform as follows:



- ◆ On a Solaris SPARC system, the start of the Solaris partition (which may contain a primary boot executable and boot block in addition to the VTOC and any disk slices) is located in cylinder 0. The whole disk is accessed using the device `c#t#d#s2`.
- ◆ On a Solaris x64 system, an FDISK partition (which may contain a master boot record (MBR)) is located in cylinder 0, and the start of the Solaris partition is located in cylinder 1. The device `c#t#d#s2` references the entire Solaris partition, but not the FDISK partition. The whole disk may be accessed using the device `c#t#d#p0`.

Before a disk with a `sun` partition label from a Solaris SPARC system can be used on a Solaris x64 system, it is necessary to use the `fdisk` command to rewrite its partition layout and VTOC, so destroying any data on the disk. However, a CDS disk group can be imported on a Solaris x64 system without needing to run the `fdisk` command. The layout of the partition table for CDS disks is the same on all supported platforms, and does not include an FDISK partition, or a Solaris partition and VTOC.

As on the Solaris SPARC platform, you can use the VERITAS Enterprise Administrator (VEA) GUI or the `vxdiskadm`, `vxdiskadd` or `vxdisk` commands to initialize a new disk with one of the following formats: `auto:cdsdisk`, `auto:simple`, `auto:sliced`, `nopriv`, `simple` or `sliced`.

Removing a Disk from VxVM Control

After removing a disk from its disk group, you can use the `vxdiskunsetup -C` command to clear the VxVM configuration on the disk:

```
# vxdiskunsetup -C daname
```

where *daname* is the disk access name (such as `c2t4d7`).

If the `vxdisk list` command shows that a disk is in the `error` state, use the following commands to reinitialize the disk with the default layout for a Solaris x64 system, and remove the disk from the VxVM configuration:

```
# fdisk -B -n /dev/rdisk/danamep0
# vxdisk rm danames2
# vxdisk scandisks
```

Note that the partition 0 device (for example, `c2t4d7p0`) is specified to the `fdisk` command, but the Solaris partition device (for example, `c2t4d7s2`) is specified to the `vxdisk rm` command.

The `vxdisk list` command should now show the disk's type as `auto:none` and its state as `online invalid`. If the disk is still not shown as being in the `online` state, use the following command to clear the first 512 blocks on the disk before rerunning the `fdisk` and `vxdisk` commands:

```
# dd if=/dev/zero of=/dev/rdisk/danamep0 count=512
```

Encapsulation of Data Disks with Insufficient Space for a Private Region

Disks with insufficient space (less than 2048 disk blocks) for the allocation of an on-disk private region cannot be encapsulated. The private region 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 private region must be allocated from the beginning or the end of the disk.

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.

Hot-Relocation Issues

Impact of Hot-Relocation on Performance

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). [14894]

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. [14895]

DMP Issues

Usage of the `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 the following conditions apply:

- ◆ VxVM volumes are not mirrored.
- ◆ There are multiple paths to the disks.



- ◆ A portion of the I/O path below DMP fails in such a way that the error is not returned to DMP before a long time has elapsed (perhaps 10 minutes or more). Examples of failures that have been seen to cause this condition include bad hardware (in an HBA, FCOT, GBIC, switch port or array controller) and HBAs that have been configured to retry endlessly.

Note This failure condition triggers the problem. Such failures are rarely seen during the normal operation of a healthy SAN.

The value of the `dmp_failed_io_threshold` kernel tunable parameter can be adjusted to prevent devices from experiencing certain failure conditions that would prevent a mirrored I/O request from succeeding for an extended period of time. This parameter sets a threshold time period for an I/O error to be returned from a device. If the I/O request takes longer than this time to return with an error, DMP assumes that the device is not working. DMP then passes the error up to VxVM without retrying the I/O request on another path, even if the delay in returning the I/O status is caused by a problem with the I/O path to the device rather than with the device itself. If the volume is mirrored, VxVM satisfies the I/O request by using another plex, and detaches the failed plex to prevent the volume from becoming hung.

Note The value of the `dmp_failed_io_threshold` does not represent a time-out. DMP waits as long as it takes for the error status of the I/O request to be returned. No DMP activity occurs if the threshold time period has passed and no status has been returned. DMP only checks the elapsed time of the I/O request after its status is returned by the lower layer.

If the volume is not mirrored, the error is 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. By default, a value of 57600 seconds (16 hours) is used.

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, `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.

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

```
dmp_failed_io_threshold=57600
```

After changing the value, reboot the system.

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

Snapshot and Snapback Issues

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`.

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 `vx dg import`. [13741]

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`. [14450]

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. [14915]

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.

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 actual number of columns that are created in such a volume is always one more than the number specified. [Sun Bug ID 4976891]



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`. [76683, 97076]

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].

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. [139162]

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. [139762]

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 cause data corruption.

Use this procedure when adding 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/vxdmping device_name
```

where *device_name* is the device name of one of the disks in the array (for example, */dev/vdsk/clt20d0s2*). 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/vxdmping /dev/rdisk/clt20d0s2
Vendor id (VID) : FUJITSU
Product id (PID): FU318404LSUN18G
Revision       : 8507
Serial Number  : 0025T0LA3H
```

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 Fujitsu devices, you must also specify the number of characters in the serial number as the argument to the *length* argument (for example, 10).

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

```
# vxddladm addjbod vid=FUJITSU pid=ST318404LSUN18G length=10
```

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 `vxddmpadm listenclosure` command as shown in the following sample output for the example array:

```
# vxddmpadm 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

VxFS NFS Server Runs out of Incore Inodes

The following issue is limited to NFS.v4 clients accessing a NFS.v4 export of a VxFS file system, when delegation is enabled. Certain NFS.v4 workloads accessing the VxFS NFS server can run out of incore inodes, because the NFS.v4 server on Solaris keeps a hold on VxFS inodes. This does not allow VxFS inode to be reused until they are released.

Workaround: In the `/etc/default/nfs` file, set `NFS_SERVER_DELEGATION=off` before bringing up NFS server daemon(s).

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 that 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.

VERITAS Enterprise Administrator Issue

The VEA server and client must run in the same language, either English(C) or Japanese (ja). If the VEA server and client are not running in the same language and the user attempts to create an action provider rule, VEA server dumps core. [398771]

VERITAS Storage Foundation for Databases Software Issues

VERITAS Storage Foundation for Oracle Software Issues

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

Incident	Description
none	ORAMAP-FMON might not start in Oracle 10gR2. To work around this issue, add the following entry to <code>init.ora</code> and try to start the ORAMAP: <code>_filemap_dir=\$ORACLE_HOME/rdbms/filemap</code>



Incident	Description
301174	The column heading TBSNAME in the output for the <code>dbed_analyzer</code> command is not localized. If you are running in an environment other than English, note that TBSNAME means “Tablespace Name.”
303238	<p>When installing the VRTSdbed package using JumpStart, you may see the following warning:</p> <pre>installing <a/opt/VRTS/man/man1m/qio_convertdbfiles.1m> with default mode of 644</pre> <p>It is safe to ignore this warning. The permission for this file is correctly set as 644.</p>
421816	The following table headings produced by <code>dbed_analyzer</code> are not aligned with the data shown below them: DATAFILE, DEVICE, SIZE(sectors), and TBSNAME. This can lead to data being aligned under the wrong heading.

Documentation Issues

- ◆ In the *Storage Foundation for Oracle 4.1 Database Administrators Guide*, on page 227 under the section “Establishing a Mandatory Archive Destination” the SQL*Plus example misplaces the second single-quote character. This single-quote should be located at the end of the command sequence, as shown below:

```
alter system set log_archive_dest_1 =
'LOCATION=/ora_mnt/oracle/oradata/PROD/archivelogs MANDATORY
[REOPEN] [scope=both];'
```
- ◆ In the *Storage Foundation 4.1 Installation Guide*, the package VRTSmuobg is incorrectly identified as being for “Windows client only.”

No Longer Supported

Storage Foundation QuickStart

The 4.1 MP1 release of VERITAS Storage Foundation is the last version to support Storage Foundation QuickStart.

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 on the Documentation CD. To view or print PDF documents, you must have the Adobe Acrobat Reader installed.

Installing documentation and online manual pages is optional.

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 of the VERITAS Storage Foundation software are located under the `storage_foundation/release_notes` directory of the VERITAS software disc.

The VERITAS Storage Foundation Installation Guides are located under the `storage_foundation/docs` directory of the VERITAS software disc.



VERITAS Storage Foundation Guides

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

Guide Title	Filename
<i>VERITAS Storage Foundation and High Availability Getting Started Guide</i>	getting_started.pdf
<i>VERITAS Storage Foundation Release Notes (this document)</i>	sf_notes.pdf
<i>VERITAS Storage Foundation Installation Guide</i>	sf_install.pdf
<i>VERITAS File System Administrator's Guide</i>	vxfs_admin.pdf
<i>VERITAS File System Programmer's Reference Guide</i>	vxfs_ref.pdf
<i>VERITAS Volume Manager Administrator's Guide</i>	vxvm_admin.pdf
<i>VERITAS Enterprise Administrator 500 Series Getting Started</i>	veax5_getting_started.pdf
<i>VERITAS Storage Foundation Intelligent Storage Provisioning Administrator's Guide</i>	sf_isp_admin.pdf
<i>VERITAS FlashSnap Point-in-Time-Copy Solutions Administrator's Guide</i>	flashsnap_admin.pdf
<i>VERITAS Storage Foundation Cross-Platform Data Sharing Administrator's Guide</i>	sf_cds_admin.pdf
<i>VERITAS Volume Manager Hardware Notes</i>	vxvm_hwnotes.pdf
<i>VERITAS Volume Manager Troubleshooting Guide</i>	vxvm_tshoot.pdf

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>

Online Manual Pages

The VERITAS online manual pages are installed in the `/opt/VRTS/man` directory. Add this directory to the `MANPATH` environment variable.

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>.



