

VERITAS Storage Foundation™ 4.0

Release Notes

Solaris

Maintenance Pack 2

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

Copyright

Copyright © 2005 VERITAS Software Corporation. All rights reserved. VERITAS, VERITAS Software, the VERITAS logo, VERITAS Storage Foundation, and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation in the USA and/or other countries. Other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650-527-8000
Fax 650-527-2908
<http://www.veritas.com>

Third-Party Copyrights

Data Encryption Standard (DES) Copyright

Copyright © 1990 Dennis Ferguson. All rights reserved.

Commercial use is permitted only if products that 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.

Contents

VERITAS Storage Foundation Products	1
Getting Help	3
Storage Foundation QuickStart	3
Storage Foundation Standard	3
Storage Foundation Standard HA	3
Storage Foundation Enterprise	4
Feature Options	4
Storage Foundation Enterprise HA	4
Feature Options	4
Installing Storage Foundation	5
Component Product Licensing	5
VERITAS Licensing Commands	6
Upgrading from VxVM and VxFS to Storage Foundation	6
File Change Log (FCL)	6
Storage Foundation Product Suite	6
Storage Foundation QuickStart	8
Licensable Features	10
End of Product Support	10
Documentation	10
Documentation for Storage Foundation Products	10
Cluster Server Documentation	11
Volume Replicator Documentation	12
FlashSnap Agent For Symmetrix Documentation	12



FlashSnap Documentation	12
Quality of Storage Service Documentation	12
Component Product Release Notes	13
Available Storage Foundation Patches for MP2	13
Component Product Release Notes and Readmes	14
Installing the Patches	15
Installing Using install_vp	15
Removing Patches	16
Storage Foundation Incidents Closed for 4.0 MP2	17
Volume Manager Closed Incidents	17
File System Closed Incidents	19
Volume Replicator Closed Incidents	21
VEA Closed Incidents	23
FlashSnap Agent for Symmetrix Closed Incidents	23
Open Incidents and Suggested Solutions	25



Important Release Information

VERITAS Storage Foundation Products

This document provides release information for VERITAS Storage Foundation 4.0 Maintenance Pack (MP) 2 for Solaris. Read this entire document before you install the product. The VERITAS Storage Foundation Release 4.0 MP2 product line contains:

- ◆ Storage Foundation QuickStart
- ◆ Storage Foundation Standard
- ◆ Storage Foundation Standard HA
- ◆ Storage Foundation Enterprise
- ◆ Storage Foundation Enterprise HA

All versions contain sets of VERITAS products that can be activated by a single license key, or features installed with the product packages and licensed separately.

The 4.0 MP2 release operates on the following Solaris operating systems:

Solaris 7 (32-bit and 64-bit)

Solaris 8 (32-bit and 64-bit)

Solaris 9 (32-bit and 64-bit)

Review this entire document before installing VERITAS Storage Foundation components. Also read the individual product release notes for important information, such as required patches and software issues, specific to those products (see “[Component Product Release Notes](#)” on page 13).

Topics in this guide include:

- ◆ [Getting Help](#)
- ◆ [Storage Foundation QuickStart](#)
- ◆ [Storage Foundation Standard](#)
- ◆ [Storage Foundation Standard HA](#)



- ◆ [Storage Foundation Enterprise](#)
- ◆ [Storage Foundation Enterprise HA](#)
- ◆ [Installing Storage Foundation](#)
 - ◆ [Component Product Licensing](#)
 - ◆ [VERITAS Licensing Commands](#)
- ◆ [Storage Foundation Product Suite](#)
- ◆ [Licensable Features](#)
- ◆ [End of Product Support](#)
- ◆ [Documentation](#)
 - ◆ [Documentation for Storage Foundation Products](#)
 - ◆ [Cluster Server Documentation](#)
 - ◆ [Volume Replicator Documentation](#)
 - ◆ [FlashSnap Documentation](#)
 - ◆ [FlashSnap Agent For Symmetrix Documentation](#)
 - ◆ [Quality of Storage Service Documentation](#)
 - ◆ [Component Product Release Notes](#)

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>.

Storage Foundation QuickStart

Storage Foundation QuickStart consists of the following products:

- ◆ VERITAS Volume Manager Base- VxVM Base
- ◆ VERITAS File System Base- VxFS Base

Storage Foundation Standard

Storage Foundation Standard consists of the following products:

- ◆ VERITAS Volume Manager (VxVM)
- ◆ VERITAS File System (VxFS)

The following feature is automatically licensed:

- ◆ VERITAS QuickLog

Storage Foundation Standard HA

Storage Foundation Standard consists of the following products:

- ◆ VERITAS Volume Manager (VxVM)
- ◆ VERITAS File System (VxFS)
- ◆ VERITAS Cluster Server (VCS)

VERITAS QuickLog is automatically licensed.



Storage Foundation Enterprise

Storage Foundation Enterprise consists of the following products and features:

- ◆ VERITAS Volume Manager (VxVM)
- ◆ VERITAS File System (VxFS)
- ◆ VERITAS FlashSnap
- ◆ Quality of Storage Service - with VxVM
- ◆ FlashSnap Agent for Symmetrix

VERITAS QuickLog is automatically licensed.

Feature Options

Storage Foundation Enterprise offers VERITAS Volume Replicator (VVR) - with VxVM as a feature option.

Storage Foundation Enterprise HA

Storage Foundation Enterprise HA consists of the following products and features:

- ◆ VERITAS Volume Manager - VxVM
- ◆ VERITAS File System - VxFS
- ◆ VERITAS Cluster Server (VCS)
- ◆ VERITAS FlashSnap
- ◆ Quality of Storage Service - with VxFS
- ◆ Global Cluster Option - with VCS
- ◆ FlashSnap Agent for Symmetrix

The following feature is automatically licensed:

- ◆ VERITAS QuickLog

Feature Options

Storage Foundation Enterprise HA offers VERITAS Volume Replicator (VVR) - with VxVM as a feature option.



Installing Storage Foundation

VERITAS 4.0 release products must be installed before updating them with Maintenance Pack 2 patches. If you have not yet installed a VERITAS 4.0 product, see the *VERITAS Storage Solutions Getting Started Guide* or *VERITAS Cluster File Solutions Getting Started Guide* for the location of release notes and installation instructions.

The VERITAS software discs have an automated installation and licensing procedure that lets you install packages using an Installation Menu instead of installing from the command line. The *Getting Started Guide*, included with the VERITAS software discs, provides complete information on using the Installation Menu. Review the *Getting Started Guide* before installing any of the Storage Foundation products.

Component Product Licensing

A Storage Foundation QuickStart key licenses a limited version of VxVM (VxVM Base) and VxFS (VxFS Base).

A Storage Foundation Standard key licenses VxVM, VxFS, and QuickLog.

A Storage Foundation Standard HA key licenses VxVM, VxFS, VCS, and QuickLog.

A Storage Foundation Enterprise key licenses VxFS, VxVM, VERITAS FlashSnap, FlashSnap Agent for Symmetrix, Quality of Storage Service, and the VERITAS QuickLog feature.

A Storage Foundation Enterprise HA key license VxFS, VxVM, VCS, VERITAS FlashSnap, FlashSnap Agent for Symmetrix, Quality of Storage Service, Global Cluster option, and the VERITAS QuickLog feature.

When you install using the VERITAS Installation menu, the following information is displayed on the Product Status Page:

- ◆ Products available for installation
- ◆ Products currently installed
- ◆ Products that are licensed
- ◆ Options for operations you can initiate

Select option **L** to enter the license key. You will not have to re-enter the key for other licensed products.

Note Some features require a separate license key (see “[Licensable Features](#)” on page 10).



VERITAS Licensing Commands

The VERITAS licensing commands are provided in the `VRTSVlic` package. You must install `VRTSVlic` for the licensing process to work. There are three licensing commands:

`vxlicinst`—Installs a VERITAS license key.

`vxlicrep`—View currently installed licenses.

`vxlictest`—Retrieves features that are encoded in a license key along with their descriptions.

You can review the descriptions and available options for these commands in the online manual pages installed with the `VRTSVlic` package.

Upgrading from VxVM and VxFS to Storage Foundation

If you already have a Storage Foundation or Storage Foundation Enterprise license, you can upgrade VxVM and VxFS using the installer script. See “Using the VERITAS Installation Menu” in the *Getting Started Guide* for more information.

File Change Log (FCL)

FCL is not supported on file systems with asymmetric permissions or if writable clones are present.

Caution File Change Log is currently not supported, and VERITAS strongly cautions against using it in a production environment. Though FCL is not 100 percent complete, it can be used to begin developing new applications. FCL will be fully operational in the next VERITAS File System release. For more information, see TechNote 265313 available at: <http://support.veritas.com/docs/265313>.

Storage Foundation Product Suite

Storage Foundation QuickStart consists of VERITAS File System (Base) and VERITAS Volume Manager (Base).

Storage Foundation Standard consists of VERITAS File System and VERITAS Volume Manager. A Storage Foundation key licenses VxFS, VxVM, and the VERITAS QuickLog feature.



Storage Foundation Standard HA consists of VERITAS File System, VERITAS Volume Manager, and VERITAS Cluster Server. A Storage Foundation HA key licenses VxFS, VxVM, VCS, and the VERITAS QuickLog feature.

Storage Foundation Enterprise consists of VERITAS File System, VERITAS Volume Manager, VERITAS FlashSnap, FlashSnap Agent for Symmetrix, and Quality of Storage Service. A Storage Foundation Enterprise key licenses VxFS, VxVM, VERITAS FlashSnap, FlashSnap Agent for Symmetrix, Quality of Storage Service, and the VERITAS QuickLog feature. VVR is available as an option.

Storage Foundation Enterprise HA consists of VERITAS File System, VERITAS Volume Manager, VERITAS Cluster Server, VERITAS FlashSnap, FlashSnap Agent for Symmetrix, Quality of Storage Service, and Global Cluster option. A Storage Foundation HA key licenses VxFS, VxVM, VCS, VERITAS FlashSnap, FlashSnap Agent for Symmetrix, Quality of Storage Service, Global Cluster option, and the VERITAS QuickLog feature. VVR is available as an option.

VxFS is a quick-recovery, journaling file system that provides high performance and online management capabilities.

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

The QuickLog feature enhances file system performance for metadata intensive applications such as mail serving. Although QuickLog can improve file system performance, VxFS does not require QuickLog to operate effectively.



Storage Foundation QuickStart

Storage Foundation QuickStart is a limited feature set of VERITAS File System and VERITAS Volume Manager functionality. Storage Foundation QuickStart provides high performance and data integrity for environments with less stringent availability and management requirements. Storage Foundation QuickStart uses the same binaries as other Storage Foundation products, so upgrading to the full functionality of VERITAS Storage Foundation requires only the installation of a Storage Foundation license key—no product installation or reboot is needed. To enable Storage Foundation options, such as FlashSnap, VVR, and Quality of Storage Service, you must first upgrade to VERITAS Storage Foundation Standard or VERITAS Storage Foundation Enterprise.

The following features are available in Storage Foundation, but are not included with Storage Foundation QuickStart.

VERITAS Base File System Features (not available with Storage Foundation QuickStart)

- ◆ Online file system resize
- ◆ Online file system defragment
- ◆ Online disk layout upgrade
- ◆ Snapshot file systems
- ◆ Group quotas
- ◆ Forced unmount
- ◆ DMAPI (Data Management API, used with hierarchical storage management (HSM))
- ◆ Space reservation and setting extent sizes
- ◆ Caching advisories
- ◆ The `mount` command options `remount`, `mincache` and `convosync` (including direct and Discovered Direct I/O, and I/O error handling)
- ◆ Storage checkpoints (Storage Foundation Enterprise and Storage Foundation Enterprise HA only)
- ◆ Quicklog
- ◆ Multi-device support
- ◆ Support for file systems greater than 2 TB and not greater than 32 TB
- ◆ Support for file systems greater than 32 TB (Storage Foundation Enterprise and Storage Foundation HA only)
- ◆ Online intent log resize

- ◆ File change log
- ◆ Portable file systems

VERITAS Base Volume Manager Features (not available in Storage Foundation QuickStart)

- ◆ Striping
- ◆ Mirroring of user-data disks (QuickStart does include root mirroring)
- ◆ Striping + mirroring (RAID 0+1)
- ◆ Mirrored stripes (RAID 1+0)
- ◆ RAID-5
- ◆ RAID-5 with logging
- ◆ Log-based mirror recovery (dirty region logging)
- ◆ SmartSync (Oracle resilvering)
- ◆ Hot-sparing/hot relocation
- ◆ Dynamic multipathing (DMP)
- ◆ Online data migration
- ◆ Snapshots
- ◆ Task monitoring
- ◆ Online relayout
- ◆ Storage Expert
- ◆ SCS III PGR
- ◆ Import volumes from host with different OS
- ◆ VxMS plugin
- ◆ Volume sets
- ◆ Fast Mirror Resynchronization (Storage Foundation Enterprise and Storage Foundation Enterprise HA only)
- ◆ Disk group split/join (Storage Foundation Enterprise and Storage Foundation Enterprise HA only)
- ◆ Attribute-based allocation
- ◆ Dynamic LUN expansion
- ◆ Ecopy



Licensable Features

VERITAS Volume Replicator (VVR) is a 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.

VVR is installed with Storage Foundation Standard Enterprise and Storage Foundation Enterprise HA, but require a separate product license key.

End of Product Support

This is the last release to support the QuickLog statistic gathering functionality provided by the `qlogstat` command. Also, future releases of QuickLog will support only one VxFS file system per QuickLog device instead of the current 32 file systems per QuickLog device.

Documentation

Release Notes and Installation guides related to the VERITAS Storage Foundation are available on the *VERITAS Storage Solutions* Disc 1 under the `storage_foundation` directory. All other documents are on the *VERITAS Storage Solutions* Disc 3 under the `storage_foundation` directory.

After the installation procedure is complete, documents are available online under the `/opt/VRTSpackage_name/doc` directories. Documents are provided as Adobe Portable Document Format (PDF) files. To view or print PDF documents, you must have the Adobe Acrobat Reader installed.

Documentation for Storage Foundation Products

The following documentation is available with Storage Foundation Standard, Storage Foundation Standard HA, Storage Foundation Enterprise, Storage Foundation Enterprise HA:

The `VRTSfsdoc` package contains the following VERITAS File System documentation:

VERITAS File System Installation Guide (includes QuickLog information)
(`vxfs_ig.pdf`)

VERITAS File System Administrator's Guide (`vxfs_ag.pdf`)

The `VRTSvmdoc` package contains the following VERITAS Volume Manager documentation:

VERITAS Volume Manager Installation Guide (`vxvm_ig.pdf`)



VERITAS Volume Manager Administrator's Guide (vxvm_ag.pdf)

VERITAS Volume Manager User's Guide – VERITAS Enterprise Administrator (vxvm_ug.pdf)

VERITAS FlashSnap Point-In-Time Copy Solutions Administrator's Guide (pitc_ag.pdf)

VERITAS Volume Manager Troubleshooting Guide (vxvm_tshoot.pdf)

VERITAS Volume Manager Hardware Notes (vxvm_hwnotes.pdf)

VERITAS Volume Manager Intelligent Storage Provisioning Administrator's Guide (vxvm_ispag.pdf)

VERITAS Volume Manager Cross-platform Data Sharing Administrator's Guide (vxvm_cdsag.pdf)

Cluster Server Documentation

The `VRTSvcSDC` package contains the Cluster Server documentation. The following documentation is available only with Storage Foundation Standard HA, Storage Foundation Enterprise HA, and Storage Foundation QuickStart HA:

VERITAS Cluster Server User's Guide (vcs_ug.pdf)

VERITAS Cluster Server Installation Guide (vcs_ig.pdf)

VERITAS Cluster Server Bundled Agents Reference Guide (vcs_barg.pdf)

VERITAS Cluster Server Agent Developer's Guide (vcs_agd.pdf)

VERITAS Cluster Server Application Note: Sun Fire 12K/15K Dynamic Reconfiguration (vcs_appnote_f15kdr.pdf)

VERITAS Cluster Server Application Note: Sun StorEdge 6800 Dynamic Reconfiguration (vcs_appnote_s6800dr.pdf)

VERITAS Cluster Server Application Note: Sun Enterprise 10000 Dynamic Reconfiguration (vcs_appnote_e10kdr.pdf)



Volume Replicator Documentation

The `VRTSvrdoc` package contains the following VERITAS Volume Replicator documentation:

VERITAS Volume Replicator Installation Guide (vvr_ig.pdf)

VERITAS Volume Replicator Administrator's Guide (vvr_ag.pdf)

VERITAS Volume Replicator Configuration Notes (vvr_config.pdf)

VERITAS Volume Replicator Web Console Administrator's Guide (vrw_ag.pdf)

VERITAS Cluster Server Agents for VERITAS Volume Replicator Configuration Guide (vcsvvr_cg.pdf)

FlashSnap Agent For Symmetrix Documentation

VERITAS FlashSnap Agent for Symmetrix 4.0 Administrator's Guide (vxfas_ag.pdf)

VERITAS FlashSnap Agent for Symmetrix 4.0 Installation Guide (vxfas_ig.pdf)

VERITAS Cluster Server Agents for Veritas Flashsnap Agent for Symmetrix 4.0 Installation and Configuration Guide (vxfas_vcsagent.pdf)

VERITAS FlashSnap Agent for Symmetrix 4.0 Release Notes (vxfas_notes.pdf)

FlashSnap Documentation

VERITAS FlashSnap Point-In-Time Copy Solutions Administrator's Guide (pitc_ag.pdf)

Quality of Storage Service Documentation

VERITAS File System Administrator's Guide (vxfs_ag.pdf)

Component Product Release Notes

Release notes for component products in all versions of the VERITAS Storage Foundation are located under the `storage_foundation/release_notes` directory of the VERITAS Storage Foundation disc or the `cluster_server/release_notes` 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 File System Release Notes (vxfs_notes.pdf)

VERITAS Volume Manager Release Notes (vxvm_notes.pdf)

VERITAS Volume Replicator Release Notes (vvr_notes.pdf)

VERITAS Cluster Server Release Notes (vcs_notes.pdf)

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

Available Storage Foundation Patches for MP2

With VERITAS Storage Foundation 4.0 MP2, the following patches are available in the `/patches` directory of the patch disc. Please refer to the VERITAS Cluster Server Release Notes for description of the VCS patches.

Operating System	Patch Number	Affected VERITAS Package
Solaris 7, 8, and 9	115217-04	VRTSvxvm
	116681-02	VRTSjavmm (language package)
	117094-03	VRTSvmpro
	117095-02	VRTSmuvmp (language package)
	116684-02	VRTSfspro
	117276-02	VRTSalloc
	117277-02	VRTSmualc
	116685-02	VRTSmufsp (language package)
	116686-02	VRTSjafsc



Operating System	Patch Number	Affected VERITAS Package
	115209-20	VRTSob
	115210-20	VRTSobgui
	115213-22	VRTSmuobg (language package)
	115212-21	VRTSmuob
	117093-02	VRTSfas
	117096-02	VRTSjafas (language package)
	117081-01	VRTSddlpr
	119288-01	VRTSmuofp (language package)
	118434-01	VRTSfppm
Solaris 7 only	116687-02	VRTSvxfs
Solaris 8 only	116688-02	VRTSvxfs
Solaris 9 only	116689-02	VRTSvxfs)

Component Product Release Notes and Readmes

In addition to reading these Release Notes before you install VERITAS Storage Foundation, it is also important that you read all the component product *Release Notes* and *Readmes*. Any product *Release Notes*, along with this document are on the product disc in the `storage_foundation/release_notes` directory. *Readmes* can be found in the `storage_foundation/patches` directory.



Installing the Patches

You can install the patches using the `install_vp` script or you can choose to use the `patchadd` command.

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

Installing Using `install_vp`

▼ To install the patches using `install_vp`

1. Stop the VERITAS Enterprise Administrator (VEA) Service before installing patches for `VRTSob`, `VRTSobgui`, `VRTSvmpro`, `VRTSfspro` and `VRTSorgui` (see the table in “[Available Storage Foundation Patches for MP2](#)” on page 13 for details):

- a. Verify the status of the VEA Service:

```
# /opt/VRTS/bin/vxsvcctl status
Current state of server : RUNNING
```

- b. Stop the VEA server:

```
# /opt/VRTS/bin/vxsvcctl stop
DBED: Successfully unloaded the Storage Foundation Provider
4.0 for Oracle
```

- c. Again verify the status of the VEA Service:

```
# /opt/VRTS/bin/vxsvcctl status
Current state of server : NOT RUNNING
```

2. Insert the patch disc into the CD-ROM drive. If you are using Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.
3. Install the patches using the `install_vp` command.

```
# ./install_vp
```

4. Reboot the system:

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



5. After installing the patches, restart the VEA Service.

- a. Verify the status of the VEA Service:

```
# /opt/VRTS/bin/vxsvcctl status
Current state of server : NOT RUNNING
```

- b. Start the VEA Service:

```
# /opt/VRTS/bin/vxsvcctl restart
```

- c. Verify the status of the VEA Service again:

```
# /opt/VRTS/bin/vxsvcctl status
Current state of server : RUNNING
```

Removing Patches

If you need to remove the patches for any reason, use the `patchrm` command.

▼ To remove the patches:

1. Log in as superuser (`root`).
2. Remove the necessary patches one at a time using the `patchrm` command.

```
# patchrm patch_number
# patchrm patch_number
...
```

See the table in “[Available Storage Foundation Patches for MP2](#)” on page 13 for a list of the patches that were required for your operating system.

3. Reboot the system:

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

Storage Foundation Incidents Closed for 4.0 MP2

The following Storage Foundation incidents were closed for 4.0 MP2.

[“ Volume Manager Closed Incidents ”](#)

[“ File System Closed Incidents ”](#)

[“ Volume Replicator Closed Incidents ”](#)

[“ VEA Closed Incidents ”](#)

[“ FlashSnap Agent for Symmetrix Closed Incidents ”](#)

Volume Manager Closed Incidents

The following table contains information about fixed issues in this release of VxVM. This table shows incidents that were escalations from either Sun or customers.

Incident	Description
146744	vxswapctl fails with error when root is encapsulated.
148168	vxconfigd core dumps due to duplicate dm records in configuration database.
153010	VxVM: Unable to boot system with encapsulated bootdisk due to missing bootdg link in /dev/vx/[r]dsk.
153137	unloading/loading of modules unsafe (req to use thread_join() instead of _fini in Sol9).
154434	Non-root user can get vxconfigd to exit.
155139	*4.00b2* When creating concat-mirror, see warnings "No disk space matches specification".
156409	Adding a new cluster node requires cluster reboot.
157740	Encapsulation with vxinstall does not work with MPXIO boot device.
157804	T3/T4 MPXIO failback disables all paths.
158823	vxdg move fails with VE_AGAIN due to device open/close assertion failure.
217093	return from mode sense cmd DAD_MODE_FORMAT may not be always reliable.
221295	vxclust error: Using CVM internal node id assignment scheme.



Incident	Description
224606	vxdumpadm displays incorrect path information.
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.
229703	vx dg split does not work on a DG created with ecopy enabled disks.
230827	VxVM can report two separate RDAC LUNs as one multipathed device.
260398	Data corruption because there were 2 volumes with the same dev minor.
263569	kernel stack overflow.
270748	Typo in vxdiskadm option to replace disk.
275942	es_devfs.pl should check if vxesd is running prior to exiting (error : syseventconfd: process xx exited with status 3.
278246	Implement B_FAILFAST support for VM 4.0 on Sol8.
289999	vxctl init on CVM slave should be disallowed since it modifies clusterid on shared storage.
293478	SOL 4.1: vxdumpinq returns non zero exit code on Successful inquiries.
293544	VM:VVR: FMR panic/hang on the system with 1+TB volume.
293862	DMP PGR not registering multiple primary keys.
294356	vxsvc daemon takes too much CPU time utilization
299814	vxconfigd dump core when the diskgroup go to SSB state.
304344	RFE: Support for option to set slow=iodelay option in /etc/default/vxrecover.
304618	No swapvol causes vxbootsetup to not partition some volumes.
304662	Both vxtranslog and vxcmdlog commands write to stderr, not stdout.
304665	NID protocol sends messages on public network for Sun Clusters.



Incident	Description
311250	vxdmpadm enable does not enable disks after upgrade to 4.0MP1 (115217-03.
311488	The entries in /dev/vx/[r]dsk are not cleared during reboot/system start-up after a temporary import of a dg in VxVM 4.0.
312026	CVM: kmsg sender and receiver threads should be interlocked for access of outbound queue for first sends.
313079	FMR: vxsnap Cannot execute /usr/sbin/vxassist: Bad address.
319284	ESD: Need to move RCM scripts on Solaris from /usr to /etc/ to support secure systems with /usr mounted read-only.
325018	Support to enable caching of WRITES in kernel caused due to UFS logging
344626	DMP PGR not registering keys on newly activated secondary paths during failover.

File System Closed Incidents

The following table contains information about fixed issues in this release of VxFS:

Incident	Description
i93459	The system froze in enospc.14a on the Solaris 9 operating system.
i140326	VxFS on Solaris, AIX, or Linux did not restrict the rm command in a sticky directory.
i142986	There was a panic in voldr1_unlog due to opening the device with OTYP_CHR instead of OTYP_LYR.
i145930	The vx_pageio1() call was not initializing ex.ex_dev.
i146224	The system panicked in the vx_cfs_inode_deinit() call during system boot-up while running KRT.
i146289	noise.fullfck failed with an I/O failure on a corrupt file system.
i146453	Extents pushed to the inode during a structural reorganization corrupts the file system.



Incident	Description
i146671	A large allocating write performed from a CFS secondary node could take a long time.
i147069	Mounting a filesystem without the <code>-o cluster</code> option, such as in exclusive mode on a shared volume that is not mounted on the other node, hits the <code>f:vx_remove_tran:2a</code> ted assert failure
i147078	The system panicked with <code>rw_exit</code> because the lock was not held.
i147121	The timed wait loop in <code>vx_do_cfs_frlock</code> used excessive cpu cycles.
i147407	If the size of the internal quotas file exceeded 2 GB, the <code>fsck</code> command failed validation on the primary ilist IFQUO inode.
i147926	Bad inode messages from <code>vx_inactive_tran</code> were seen at shutdown when using a VxFS file system as the <code>/tmp</code> directory.
i148573	Inodes were leaking due to inode table overflow.
i149244	Junk data could be readily found in non-sparse VxFS files after a system crash.
i149387	There was a panic in <code>vx_hsm_creaetandmkdir</code> .
i149908	Multiple DMAPI invisible reads on a file changed the file's atime.
i150160	There was a Solaris kernel panic when attempting to mount a VxFS file system as read only when one of the IFILT inodes was cleared due to either corruption or using the <code>fsck</code> command.
i150681	The <code>vxfsconvert</code> command failed with "ERROR: V-3-24720: passed EOF".
i150933	The <code>vx_vop_close()</code> call should always be called with <code>lastclose=1</code> .
i151064	Parallel mount at boot up time sometimes resulted in failures due to the VxFS mount dependency on <code>/opt</code> .
i152379	CFS set file permissions inconsistently.



Volume Replicator Closed Incidents

The following table contains information about fixed issues in this release of VVR:

Incident	Description
158253	The <code>vxconfigd</code> daemon on the Primary could hang during autosync of a large number of RLINKs. This occurred because readback memory was not freed in a timely manner, due to an uninitialized variable.
206566	In certain cases, deadlocks could occur because spinlocks were held while calling <code>untimeout</code> .
217203	The Primary panicked due to delaying in the interrupt context while running stress test cases.
258332	Incorrect handling of new error message types could cause replication to hang.
272651	The maximum length of the DCM has been increased to 256K by default; you can explicitly specify up to 2 MB.
275805	An attempt to resume replication on the Secondary caused the VVR Primary host to panic under certain circumstances.
279096	After Secondary log disk failure, the <code>secondary_log_err</code> flag on the Primary was not getting cleared even after the Secondary log disk was fixed.
286984	For an RVG containing multiple RLINKs, the VVR Primary host could panic under certain circumstances.
306656	The output of <code>vradm in syncrvg</code> and <code>vradm in syncvol</code> commands did not display volume size correctly for very large volumes ($\geq \sim 977G$).
306734	The <code>in.vxrsyncd</code> daemon did not start up correctly. This problem was fixed by removing the unnecessary calls to <code>gethostname/gethostbyname</code> .
306775	If differences were less than 1%, the output of the <code>vradm in syncrvg</code> and <code>vradm in syncvol</code> commands incorrectly displayed as 0% differences.
306852	After detaching, dissociating and reassociating the Secondary RLINK, the <code>vradm in migrate</code> command would leave the RDS in a Secondary-Secondary configuration.
306883	When the <code>vradm in</code> command was used without specifying a diskgroup, while the <code>vxconfigd</code> was disabled, the <code>vradmind</code> could dump core.



Incident	Description
306891	When adding a Secondary to an RDS that contains a large number of volumes, the <code>vradmin addsec</code> command could fail with the error "Lost connection to host <code>hostname</code> ".
306904	The message "Collect Stats for rlink..." was logged repeatedly, and could fill up the VRAS debug log. These unnecessary messages have been removed.
307031	When adding a Secondary to the RDS, the VVR VEA GUI incorrectly displayed a configuration error unless you refreshed the data.
307052	The <code>vradmin</code> could dump core on the Secondary during replication between the AIX and Solaris platforms.
307082	The <code>vradmin</code> command could fail with an incorrect message about network disconnection when the network connection was fine. This problem was fixed by stopping the stats ioctls while the command is in progress.
307090	When using VVR VEA GUI in the Ja locale, certain VRAS messages were displayed in English rather than in Japanese.
315248	In some cases, the <code>vradmin</code> could dump core when renaming a disk group.
322372	Allow better control of CPU usage in multi-threaded difference-based synchronization.



VEA Closed Incidents

The following table contains information about fixed issues in this release of VEA:

Incident	Description
i150138	VEA VxVM provider has missing relationship between disk groups and disks.
i150351	VEA VxVM provider to use thread safe time and date calls.
i151118	Relayout from striped-mirror to striped invoked incorrect VxVM command.
i151427	Snapshot refresh ignores disabling sync.
302852	vxsvc may spin after restarting vxconfigd.
330689	A previously renamed disk cannot be removed from a disk group.

FlashSnap Agent for Symmetrix Closed Incidents

The following table contains information about fixed issues in this release of VERITAS FlashSnap Agent for Symmetrix:

Incident	Description
316944	Support RAID-5 STD devices in Flashsnap Agent for Symmetrix operations.
317228	Attach/detach operations take a long time to complete when there are a large number of disks.
317229	If vxmake fails in a split operation, the BCV disk group that is created is not destroyed before the split is rolled back.
317230	If vxmake fails in a split operation, save the file that is given to vxmake for analysis.
317231	Trying to split a disk group with a large number of disks fails because of a correlation failure.
317232	vxassist and vxdg commands take a long time on a disk group with a large number of disks.



Incident	Description
317233	CorrelateVMDisk should not return an error if it fails to add the property PROP_VRTS_VXVM_DISKS on a virtual disk object because the property is not used.
317234	In vxsymrestore, import the disk group using the vxdg CLI instead of the VM provider operation.
337533	OnVMDiskAdded takes a long time to correlate a VM disk.



Open Incidents and Suggested Solutions

The following open incidents and suggested solutions are noted for the 4.0 MP2 release.

[Section 1: Special Note for UFS Logging](#)

[Section 2: Installing the VxVM 4.0 Maintenance Patch 2](#)

[2.1 Localized VxVM 4.0](#)

[2.2 JBOD Array Considerations](#)

[2.3 vxnotify Message](#)

[Section 3: Duplicate Disk IDs](#)

[Section 4: vxassist Relay Layout Considerations](#)

[Section 5: Using vxunroot to Unencapsulate the Root Disk](#)

[Section 6: Suppressing a Path from DMP and VxVM in a Multipath Array](#)

[Section 7: Using Hitachi Arrays](#)

[Section 8: T3B Firmware Upgrade on Solaris 9](#)

[Section 9: Replacing a Failed Boot Disk](#)

[Section 10: Recovering from Master Node Failure During a Snapback Operation](#)

[Section 11: Suppressing Unwanted vxassist Authentication Messages When Accessing SAL](#)

[Section 12: SCSI-3 Fencing \(vxfen\) on Solaris 9](#)

[Section 13: Potential Solaris Patch Issues](#)

[Section 14: Troubleshooting a Duplicate Device Entry in vxdisk list](#)

[Section 15: Disk Connectivity Policy](#)

[15.1 Global Disk Detach Policy](#)

[15.2 Local Disk Detach Policy](#)

[Section 16: Known Issues with JNI HBAs](#)

[Section 17: Support for Hitachi Arrays](#)

[Section 18: Join and Master Failover Times](#)

[Section 19: cvm Timeout for SunCluster](#)

[Section 20: Localization Issues](#)

[Section 21: Using powervxvm With Volume Manager 4.0MP1 \(or later\)](#)

[Section 22: cvm 4.0 Upgrade Doesn't Upgrade Protocol Version](#)



[Section 23: vradm Print Commands](#)

[Section 24: Issues Regarding CDS \(Cross-Platform Data Sharing\)](#)

[Section 25: Install FAS Patch After Volume Manager Patch](#)

[Section 26: Multi-host Configurations With Sun StorEdge T3 or 6120/6320](#)

[Section 27: Booting From a Secondary Connected Device](#)

[Section 28: Documentation Issues](#)

Section 1: Special Note for UFS Logging

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

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

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

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

Section 2: Installing the VxVM 4.0 Maintenance Patch 2

2.1 Localized VxVM 4.0

This patch also supports localized releases of VxVM 4.0. When you install this patch, if previous 4.0 language packages are already installed, it will replace all the necessary localized packages including language packages and messages.

2.2 JBOD Array Considerations

If you have an array which was previously claimed as jbod, you may see the following warning message during reboot after the upgrade:

```
NOTICE: vxvm:vxdump: added disk array <disk-array-serial-no>  
vxvm:vxconfigd: WARNING: File /etc/vx/array.info not in proper  
format.  
Regenerating file /etc/vx/array.info.  
All attributes will be set to default values.
```



The system is coming up. Please wait.

This message informs you that your array configuration has changed and that VxVM is recreating the file.

2.3 vxnotify Message

When installing this patch, the following "vxnotify" message may appear on the console:

```
Adding patch 115217-03 on thor139 ..... Done 25 of 26
steps
VxVM vxnotify ERROR V-5-1-915 Reconnection to vxconfigd failed:
Configuration daemon is not accessible
VxVM vxnotify ERROR V-5-1-915 Reconnection to VxVM vxconfigd
failed:
Configuration daemon is not accessible
vxnotify ERROR V-5-1-915 Reconnection to vxconfigd failed:
Configuration daemon is not accessible
Adding patch 115217-03 on thor140 ..... Done 26 of 26
steps
```

This is a harmless message and it appears as a result of the vxconfigd process being shutdown and restarted. This vxnotify message appears during the temporary downtime of vxconfigd daemon.

Reference: Incident e158898

Section 3: Duplicate Disk IDs

When VxVM detects disks with duplicate disk IDs, VxVM attempts to select the appropriate disk (using array vendor-specific logic). If a disk cannot be selected, VxVM does not import any of the duplicated disks into a disk group.

In the rare case when VxVM cannot make the selection, you must choose which duplicate disk to use. An array with hardware mirroring capability is particularly susceptible to data corruption, if the wrong disk were to be selected.

The following examples describe situations where user intervention is needed.

Example 1:

When DMP is disabled to an array that has multiple paths, then each path to the array is claimed as a unique disk. VxVM detects duplicate disks whenever an array is configured as a unique disk for each accessible path and gives the following message:

```
vxvm:vxconfigd: NOTICE: Unable to resolve duplicate diskid.
```



When DMP is suppressed, VxVM can not know which path to select as the true path. You must decide which path to use. Decide which path to exclude and then either edit the file `/etc/vx/vxvm.exclude` or, if `vxconfigd` is running, use the `vxdiskadm` option 17 selection 1 (suppress all paths through a controller from VxVM view) or selection 2 (suppress a path from VxVM view).

The following example shows a `vxvm.exclude` file with paths `c6t0d0s2`,

```
c6t0d1s2, and c6t0d2s2 excluded from VxVM:
exclude_all 0 paths
c6t0d0s2 /pci@1f,4000/SUNW,ifp@2/ssd@w50060e8003275705,0
c6t0d1s2 /pci@1f,4000/SUNW,ifp@2/ssd@w50060e8003275705,1
c6t0d2s2 /pci@1f,4000/SUNW,ifp@2/ssd@w50060e8003275705,2
controllers
product
pathgroups
```

Example 2:

Some arrays such as EMC, HDS, and so on provide hardware mirroring. When a lun pair is split, depending on how the process is performed, you may get two disks with the same diskid. With duplicate disk IDs, VxVM gives the following message:

```
vxvm:vxconfigd: NOTICE: Unable to resolve duplicate diskid.
```

Check with your array vendor to make sure that your site uses the proper split procedure. If you know which luns to use, decide which path to exclude and then either edit the file `/etc/vx/vxvm.exclude` or, if `vxconfigd` is running, use the `vxdiskadm` option 17 selection 1 (suppress all paths through a controller from VxVM view) or selection 2 (suppress a path from VxVM view).

Example 3:

When you have disks duplicated using `dd` or any other disk copying utility, VxVM gives the following message:

```
vxvm:vxconfigd: NOTICE: Unable to resolve duplicate diskid.
```

Choose which set of duplicated disks that you want to exclude, and either edit the file `/etc/vx/vxvm.exclude` or, if `vxconfigd` is running, use the `vxdiskadm` option 17 selection 1 (suppress all paths through a controller from VxVM view) or selection 2 (suppress a path from VxVM view).

Section 4: vxassist Relay Layout Considerations

The `vxassist` relay layout operation requires all mirrors in the volume to have the same layout. If the volume contains mirrors with different layouts, then you need to relay layout the mirror plexes to the same layout before performing the volume relay layout operation.

Section 5: Using vxunroot to Unencapsulate the Root Disk

Only those encapsulated volumes that were derived from the pre-encapsulated root disk partitions can be unrooted. All other volumes created on the root disk after encapsulation must be either removed or moved to another disk before using vxunroot.

Section 6: Suppressing a Path from DMP and VxVM in a Multipath Array

(Ref. incident 108881) This applies only when running on Solaris 9.

Problem: If you have an array with multiple paths, and after suppressing one path from DMP, suppress that path from VxVM using the vxdiskadm option 17 and option 1. Then, if all rootdg disks are from that array you will receive errors of the following form, and vxconfigd will not start, so VxVM will not run:

```
vxvm:vxconfigd: NOTICE: Unable to resolve duplicate diskid
Please refer to release notes and admin guide for possible
action/solution.
Following are the disks with duplicate diskid:
Vendor: SUN Product: T300 - c1t1d2s2, c4t2d2s2...
Following are the disks with duplicate diskid:
Vendor: SUN Product: T300 - c1t1d3s2, c4t2d3s2
WARNING: vxvm:vxio: cannot log commit record for Diskgroup rootdg:
error 28
vxvm:vxconfigd: ERROR: enable failed: Error in disk group
configuration copies
Unexpected kernel error in configuration update; transactions are
disabled.
vxvm:vxconfigd: FATAL ERROR: Rootdg cannot be imported during boot
```

Suggested Solution:

If only one array is connected to one controller, perform the following steps:

1. Suppress path from DMP, using vxdiskadm option 17, then option 5.
2. Suppress path from VxVM, using vxdiskadm option 17, then option 1.

If more than one array is connected to one controller, perform the following steps:

1. Suppress path from DMP, using vxdiskadm option 17, then option 5
2. Suppress every path from VxVM belonging to the array, using vxdiskadm option 17, then option 2.



Section 7: Using Hitachi Arrays

(Ref. incident 100458)

If you are considering having your boot disk residing on an Hitachi array you should first contact Hitachi Data Systems for the latest information on supported boot disk configurations with Hitachi arrays.

Section 8: 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. The procedure is a result of VERITAS incident number 95877.

▼ To upgrade the T3B firmware:

1. Use the `mount` command to unmount related filesystems.

2. Stop all VxVM volumes:

```
# vxvol stop <vol_name>
```

3. Stop VxVM:

```
# vxdctl stop  
# vxiod -f set 0
```

4. Upgrade 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.



Section 9: Replacing a Failed Boot Disk

(Ref. incident 109757)

When using the `vxdiskadm` option #5 to replace a failed bootdisk (rootdisk) which is under VxVM control, you must select “yes” when the following message appears. Failure to do so will result in an unbootable system, even though the mirrors are complete.

```
The disk c0t0d0s2 was a previously encapsulated root disk.
Due to the disk layout that results from root disk encapsulation,
the preferred action is to reinitialize and reorganize this disk.
However, if you have any non-redundant data on this disk you should
not reorganize this disk, as the data will be lost.
Reorganize the disk [y,n,q,?] (default: n)
```

You must now select “y.”

Section 10: Recovering from Master Node Failure During a Snapback Operation

Crash Recovery for `vxassist -o resyncfromoriginal snapback snapvol`

If a default snapback operation (resynchronizing from the original volume) was in progress when a system crash occurred, the snapshot plexes are not associated with any volume when the system comes back up and the volumes are restarted.

▼ To reassociate snapshot plexes with a volume, perform the following steps:

1. Use the following commands to find out the original volume name of each snapshot plex:

```
# volrid=`vxprint -g <diskgroup> -p -F "%snap_rid" <plexname>`
# vxprint -g <diskgroup> -n -v -e v_rid=$volrid
```

2. Reattach all snapshot plexes that you discovered had the same original volume in step 1 to their original volume:

```
# vxplex att <original_volume> <plex1> [<plex2> ... ]
```

This results in a full resynchronization of these plexes.



Crash Recovery for vxassist -o resyncfromreplica snapback snapvol

If a snapback operation specifying a resynchronization from the replica snapshot volume was in progress when a system crash occurred, startup of the original volume fails with the following error message when the system comes back up:

```
vxvm:vxvol: ERROR: Volume <original_volume> has no CLEAN or  
non-volatile ACTIVE plexes
```

▼ **To reattach the plexes to a volume, perform the following steps:**

1. Dissociate all STALE plexes from the original volume:
vxplex dis <staleplex1> [<staleplex2> ...]
2. Convert all SNAPTMP plexes in the original volume to ACTIVE:
vxplex convert state=ACTIVE <tmpplex1> [<tmpplex2> ...]



3. Restart the original volume:

```
# vxvol start <original_volume>
```

4. Reattach the plexes that you dissociated in step 1:

```
# vxplex att <original_volume> <staleplex1> [<staleplex2> ...]
```

This results in a full resynchronization of these plexes from the original volume.

▼ To return to the original state, perform the following steps:

1. Remove the dangling snapshot volume (the volume without any snap and plexes).
2. Recreate the SNAPDONE plexes.

Section 11: Suppressing Unwanted vxassist Authentication Messages When Accessing SAL

In this release of Volume Manager, if SAL is installed on a host on which vxassist is run, warning messages may be output when vxassist tries to contact SAL. See “[Example 1](#)” and “[Example 2](#)”, below. In both cases, you can suppress communication between vxassist and SAL by adding the following line to the vxassist defaults file:

```
salcontact=no
```

The vxassist defaults file is usually “/etc/defaults/vxassist.” See the vxassist(1m) man page for more information.

Example 1

vxassist shows the warning message, “WARNING: SAL authentication failed”

This warning message occurs because SAL rejects the credentials supplied by vxassist. For example:

```
# vxassist make voltest 100m
vxvm:vxassist: WARNING: SAL authentication failed. Username "root"
not found in password file
```

If connection to SAL is desired then you need to set valid username and password, using the command vxspcshow. Refer to the vxspcshow man page for more information.

Example 2

If vxassist does recognize the version of SAL being used, or detects an error in the SAL output, the following message may appear:

```
"WARNING: Error while retrieving information from SAL".
```



Section 12: SCSI-3 Fencing (vxfen) on Solaris 9

(Ref. incident 111620)

You must have Solaris patch 113277-08 (or higher) installed if you intend to employ SCSI-3 fencing (vxfen) on Solaris 9.

Failure to install this patch may cause the master node to fail due to being fenced out, then the other nodes will not take over as master node. Shared disk groups will be inaccessible, and the CFS filesystems may be dismounted.

Section 13: Potential Solaris Patch Issues

Issue 1 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 them or remove them.

The following patch levels have been verified with VxVM 4.0:

110934-14

113713-11

To install the latest revision of a patch, use the `patchadd` command.

For example:

```
# patchadd 110934-14
```

To remove a patch, use the `patchrm` command. For example:

```
# patchrm 110934-10
```

Note The patch might not be removable, in which case a message similar to the following will be displayed.

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

```
It cannot be backed out.  
Patchrm is terminating."
```

For Solaris 8, patch 110934-10, refer to TechNote 252441:

<http://seer.support.veritas.com/docs/252441.htm>.

For Solaris 9 patch 113713-01, the suggested solution is to rename the space file before running patchadd, as follows:

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

then:

```
# patchadd <THIS_PATCH_ID>
```

For Solaris 8, you can use patch 110934-08 or lower. If you do not have or cannot obtain patch 110934-08, do not install patch 110934-10.

You can successfully install the VERITAS package 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>.

Issue 2 - Solaris Patch 108827-19 Superseded by 108993-XX

In some systems, you may not have required patch 108827-19 installed.

Sun has superseded 108827-19 with 108993-XX. In this case, you must install 108993-XX and override the patch warning.

Visit SunSolve (<http://sunsolve.sun.com>) for latest OS patch dependencies.



Section 14: Troubleshooting a Duplicate Device Entry in vxdisk list

(Ref. incidents 114479 and 101371)

Please follow this procedure if you encounter duplicate entry in "vxdisk list" or vxdisksetup gives "Duplicate DA" error as specified in Incident:101371 SunBug:4630477 and Incident:114479 SunBug:4769704.

```
# vxdisksetup -i c1t5d0
vxdisksetup: c1t5d0: Duplicate DA records encountered for this
device.
Refer to the troubleshooting guide to clear them.
# vxdisk list
c0t8d0s2      sliced      -           -           online
c1t2d0s2      sliced      c1t2d0s2   rootdg      online
c1t3d0s2      sliced      c1t3d0s2   ttdg        online
c1t5d0s2      sliced      c1t5d0s2   -           error
c1t5d0s2      sliced      c1t5d0s2   -           online/error <<<-- any
state is ok.
```

1. Remove c1t5d0s2 entries from vxvm control.

Run "vxdisk rm <da-name>" for all the duplicate entries. Since you don't know which one is the valid one, do it for all. There may be more than two duplicate entries.

```
vxdisk rm c1t5d0s2
vxdisk rm c1t5d0s2 <-- do it again to remove all the entries.
```

2. Remove the disk c1t5d0s2 using luxadm, a Solaris command.
3. Get A5K Array name and Slot number of the disk using luxadm.

```
# luxadm disp /dev/rdisk/c1t5d0s2
```

4. Remove device c1t5d0s2 using "luxadm remove_device" command. [luxadm remove_device enclr,slot-number]

```
# luxadm remove_device SAHYADRI,f5
```

5. Pull the disk out as per luxadm instructions.
6. Run command "devfsadm -C".
7. Run command "vxdctl enable".

Up to this point, we have removed the dev_t corresponding to the physical disk. Now we will remove all the stale dev_t's

8. Loop ac like the following example:

```
LOOP :
```

You will see one entry less, since we can have more than two duplicate entries.

```
# vxdisk list
c0t8d0s2      sliced      -           -           online
c1t2d0s2      sliced      c1t2d0s2   rootdg      online
c1t3d0s2      sliced      c1t3d0s2   ttdg        online
c1t5d0s2      sliced      c1t5d0s2   -           error
```

9. Remove *ALL* duplicate c1t5d0s2 entries from vxvm control.

```
vxdisk rm c1t5d0s2
```

10. Run command "luxadm -e offline <path to disk>" on *ALL THE PATHS* to the disk. This removes the stale dev_t.

A test machine uses two paths to the disk through controllers c1 and c2:

```
# luxadm -e offline /dev/dsk/c1t5d0s2
# luxadm -e offline /dev/dsk/c2t5d0s2
```

11. Run command "devfsadm -C".**12.** Run command "vxdctl enable" goto LOOP:

Continue this process until there are no more entries in vxdisk list of corresponding disk c1t5d0s2. Result:

```
# vxdisk list
<snip>
c0t8d0s2      sliced      -           -           online
c1t2d0s2      sliced      c1t2d0s2   rootdg      online
c1t3d0s2      sliced      c1t3d0s2   ttdg        online
```

Now both the OS device tree and VxVM are in a clean state corresponding to disk c1t5d0s2.

Follow the procedure to replace the failed disk or removed disk as described in the *Volume Manager Administration Guide* to replace a new disk in place of device c1t5d0s2.



Section 15: Disk Connectivity Policy

(Ref incident 118065)

With VERITAS Volume Manager (VxVM), it is possible to create a shared disk group on the master node of a cluster. This provides all nodes in the cluster with concurrent read and write access to the volumes within the shared disk group.

Only the master node can create a shared disk group. This has the following advantages and implications:

- ◆ All the nodes in the cluster see exactly the same configuration.
- ◆ Only the master node can change the configuration.
- ◆ Any changes on the master node are coordinated and propagated to the other (slave) nodes in the cluster.
- ◆ Any failures requiring a configuration change must be sent to the master so that they can be correctly resolved.
- ◆ As the master node resolves any failure, all the slave nodes are correctly updated; ensuring that all nodes have the same view of the configuration.

The practical implication of this approach is that any IO failure on any node results in the configuration of all nodes being changed.

However, in some cases, it is not desirable to have all nodes reacting in this way to an IO failure (this is known as the global disk detach policy). To address these cases, an alternative way of responding to IO failures was added for shared disk groups. Starting with VxVM 3.2 the local disk detach policy (more formally known as the connectivity policy) became available for disk groups version 70 and above.

Note When the master causes an IO failure on a disk, the disk is marked as failed and removed from the disk group. This is the same for both local detach policy and global detach policy. On all the nodes, the volumes on that disk will be unable to perform IO to an underlying physical device. When a slave causes an IO failure the disk is not removed from the disk group.

15.1 Global Disk Detach Policy

The global disk detach policy is the traditional and default policy for all nodes on the configuration.

In this case, if there is an IO failure on one node, the master node performs the normal IO recovery work to repair the failure, and the plex is detached cluster wide. All nodes in the cluster continue to perform IO functions.

Note The global detach policy must be used when Disk MultiPathing (DMP) is managing multi-pathing on Active/Passive arrays. This ensures that all nodes correctly coordinate use of the active path.

15.2 Local Disk Detach Policy

The local disk detach policy was first designed for failover applications in large clusters. If an application is affected by IO failures, the cluster framework is then able to move the application to a node that still had access to the volume.

The local disk detach policy is used particularly with mirrored volumes. For unmirrored or hardware mirrored) volumes there is no difference between the local and global detach policies.

In the case of local disk detach policy, if there is a write failure on a node, the master node performs the normal IO recovery work to repair the failure; but in addition, all the nodes are contacted to see if the disk is still acceptable. If the failure is not seen by all nodes, the local detach policy stops IO only from the node that had the failure.

Note For private disk groups, the local disk detach policy does not change the behavior of the disk group.



Section 16: Known Issues with JNI HBAs

If your JNI card has Model numbers FCE-1063, FCE2-1063, FCE-6410, FCE2-6410, or FCE2-6412, then 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."
```

Suggested Solution: Add the following parameter to the JNI configuration file (jnic.conf):

```
FcEnableContextSwitch = 1;
```

Section 17: Support for Hitachi Arrays

(Ref incident 129438)

It should be noted that only A/P mode is supported for the Hitachi arrays DF400/HDS5800 and DF500/HDS9200. If you have installed any of these arrays, you need to verify that they are configured for A/P mode.

Contact VERITAS Support if you need assistance in determining your Hitachi array mode.

If you need assistance in reconfiguring these arrays to A/P mode, contact your Hitachi Field Engineer.

Note VxVM does not support SCSI3/PGR on HDS9200 in A/P mode on Solaris 9.



Section 18: Join and Master Failover Times

(Ref incident 140965)

Problem: In some cases customers may experience long join times or long master failover times. Customers are notified by a message like this example:

```
"cluster_establish: timed out"
```

To enable VERITAS support to diagnose the cause, it may be necessary to turn on additional diagnostic information (see below) on nodes experiencing this problem. This information should then be provided to your VERITAS support engineer, or to VERITAS Customer Support.

To turn on debug messages go to:

```
/opt/VRTSvcs/bin/CVMCluster/online
```

Change line:

```
VXCLUSTADM=/etc/vx/bin/vxclustadm
```

to

```
VXCLUSTADM="/etc/vx/bin/vxclustadm -T"
```

Once the messages are turned on, you must offline the node and then return it online for the changes to take affect.

For more help contact your VERITAS support engineer or VERITAS Customer Support.

Section 19: cvm Timeout for SunCluster

(Ref incident 142776) Sun 5003523

Instantaneous timeouts for SunCluster may be caused by setting cvm timeout values too high. The maximum value for these timeouts must not be more than 2147.

The path to the file holding the timeout values is:

```
/opt/SUNWcvm/etc/cvm.conf
```

The timeout values for the following actions must not exceed 2147 seconds:

```
cvm.start_timeout
cvm.stop_timeout
cvm.abort_timeout
cvm.return_timeout
cvm.step1_timeout
cvm.step2_timeout
cvm.step3_timeout
cvm.step4_timeout
```

For more help contact your Sun support engineer or Sun Customer Support.



Section 20: Localization Issues

File Name Encoding

If you share files between different platforms with CDS, the encoding setting of the locale in the current session should be same as the encoding of file name in order to view the file name correctly. The VERITAS CDS function does not convert the file name to your current locale setting. This applies to non-English (non ASCII) file names including Japanese file names.

In the Japanese localized VEA in some circumstances the system messages may be incomplete or missing and interpreted as follows:

Incident 209529

In the Japanese localized version of VxVM, the text and buttons are not fully displayed in Web GUI when font size is not small enough.

Section 21: Using powervxvm With Volume Manager 4.0MP1 (or later)

The powervxvm script for EMC Symmetrix and Cx600 arrays was introduced to allow users to create simple disks in VxVM 3.5. However in the VxVM 4.0 release, the powervxvm script is not working as expected. Problems:

1. The powervxvm script must be modified to make "powervxvm define" work.
2. There are duplicated entries for da.
3. After rebooting the system, emcpower devices turns to error.

Suggested Solution for problem 1:

Use the powervxvm.40 script. Here is the difference from powervxvm.40 and powervxvm:

```
vm450e1:diff powervxvm.40 powervxvm
64c64
<     vxdisk $1 $device $3
---
>     vxdisk $1 $device
169c169
<     alldisks "define" "defining" "type=simple"
---
>     alldisks "define" "defining"
325c325
<         alldisks "define" "defining" "type=simple"
---
```

```

>          alldisks "define" "defining"
383c383
<          vxdisk -f define $device type=simple
---
>          vxdisk -f define $device

```

Suggested Solutions for problems 2 and 3:

When you are using the `powervxvm` script add a record using the `'vxdisk define'` command. This command directly creates a new da record in `vxconfigd`. Adding this record will not remove any duplicate entry as DDL is unaware of these records. This is an expected behavior.

The `emcpower` devices go into the error state because devices in `/dev/vx/[r]dmp` directory are no longer persistent.

To become active, the `'powervxvm setup'` commands must be run during each reboot after `tmpfs` is mounted on `/dev/vx/[r]dmp` directory but before `vxconfigd` is started.

Add `"/etc/powervxvm setup"` after mount `dmpfs` in `vxvm-startup2`.

Then reboot the system, all `emcpower` devices are online.

Section 22: cvm 4.0 Upgrade Doesn't Upgrade Protocol Version

When a single node in a cluster is upgraded to `cvm 4.0`, other nodes will have errors. The error message will appear when trying to create the shared diskgroup after the upgrade. The error message points to the diskgroup version:

```
VxVM vxdg ERROR V-5-1-585 Disk group sharedg: cannot create: Disk
group version doesn't support feature; see the vxdg upgrade
command.
```

The upgrade procedure should include the upgrade of the CVM protocol version. After an upgrade the CVM protocol version remains at 40 instead of 50.

A warning note has been added:

```
VxVM vxvm-startup2 WARNING V-5-2-0 CVM protocol version is not
up-to-date. vmesc3.veritas.com is running at CVM protocol version
40 while the highest available is 50. Refer to the release notes
for possible action/solution.
```

Solution: The 4.0 upgrade script should run `'vxdctl upgrade'`.

The upgrade to the newer protocol version can only be done when all nodes have been upgraded to the newer version of VxVM.



Section 23: vradmin Print Commands

Ref Incident e208339. The `vradmin print` commands display incorrect messages when the `vradmind` daemon and `vradmin` client are running in different locales.

Suggested Solution:

1. On the host on which VVR is installed, stop the `vradmind` daemon by issuing the command `/etc/init.d/vras-vradmind.sh stop`.
2. Set locale of the host to the required locale. Note that the locale must be `ja`, `ja_JP.PCK`, or `ja_JP.UTF-8`.
3. Restart the `vradmind` daemon by issuing the command `/etc/init.d/vras-vradmind.sh start`.
4. Make sure that the locale of the host on which the `vradmin` client is running is the same as that of the host on which `vradmind` daemon is running.

Section 24: Issues Regarding CDS (Cross-Platform Data Sharing)

Platforms Supporting CDS

Platform	Version
Solaris	4.0 MP1
Linux	4.0
AIX	4.0
HP	4.0

External Quota File

A CDS-converted file system does not acknowledge the external quota file, causing the file system to hit an assert (`f:vx_msgprint:ndebug`) if the mount option `-o quota` is used. This assert can also be hit if the converted file system is mounted and you attempt to enable quotas.

Suggested Solution:

Remove the external quota file prior to converting the file system. Convert, and then re-create the quota file on the platform.

CDS Validation Can Be Slow

CDS validation reads the metadata of all the inodes on a disk to determine which file system entities have exceeded the limits for the specified operating systems. This can be time-consuming, and because the file system is mounted, the usages can change while validation is in progress.

Disk Layout Upgrades

Following a disk layout upgrade, you must unmount the upgraded file system prior to running `fscdstask`. Otherwise you will receive an error message and the operation terminates. If this occurs, unmount, and then remount your file system.

Sharing file system quotas must be configured manually.

On the source:

1. Remove quotas and the `quotas.grp` file prior to unmounting the file system.
2. Run `fscdsconv` after unmounting the file system.

On the target:

1. Mount the file system without quotas.
2. Manually edit quotas and the `quotas.grp` file and enter the limits.
3. Enable quotas.

Handling ACLs with CDS

Because not all target platforms support ACLs, converting a file system with ACLs from the source to a target on which ACLs are not enabled results in ACLs not being supported on the target. If the file system is converted back to a target on which ACLs are supported, permission checks are enforced again.

Metasave is currently unsupported.



Section 25: Install FAS Patch After Volume Manager Patch

Incident 146435

The VERITAS Volume Manager (VxVM) patch 115217-03 (or later) must be installed before this VxFAS patch. If not installed in the correct order, the EMC configuration rule checking for Volume Manager operations on a Symmetrix device will not be enabled. You must reinstall this VxFAS patch any time VxVM software is upgraded to ensure that EMC configuration rule checking is enabled.

The installation order is important because the VRTSfas package component library, `/usr/lib/libarray.so`, provides the EMC configuration rule checking functionality. The VRTSvxvm package installs only a stub `libarray.so` file that does not perform EMC rule checking. When VRTSfas package is installed, it overwrites the stub `libarray.so` file with the real `libarray.so` file so that EMC configuration rule checking is enabled. When the VRTSfas package is removed, it restores the original (stub) `libarray.so` file.

To identify the correct `libarray.so` file, type the following command:

```
# strings /usr/lib/libarray.so | grep EMC
```

If any output is displayed, the VxFAS `libarray.so` is installed.

Section 26: Multi-host Configurations With Sun StorEdge T3 or 6120/6320

In multi-host configurations, the Sun StorEdge T3 or 6120/6320 must be configured to the MPxIO mode (explicit failover). In this mode, the Sun StorEdge Traffic Manager Software (STMS) `mpxio` driver on the host handles multipathing. VERITAS Volume Manager (VxVM) will not see multiple paths to any device on the array.

The `/kernel/drv/scsi_vhci.conf` file must have the MPxIO mode enabled (`mpxio-disable="no"`).

Use the command `"sys mp_support=mpxio"` on the Sun StorEdge T3 or 6120/6320 array to enable the MPxIO mode.

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

Section 27: Booting From a Secondary Connected Device

Incident e213970

Booting Volume Manager from a device which is connected only to the secondary controller in an A/P (active/passive) array is not supported.

Section 28: Documentation Issues

STORAGE FOUNDATION 4.0 INSTALLATION GUIDE

This section describes corrections to the *Storage Foundation 4.0 Installation Guide*.

Setting Up VERITAS Volume Manager

In “Modifying Connection Access” on page 49, the correct command to modify the Security key in the registry (step 3) is:

```
# /opt/VRTSob/bin/vxregctl /etc/vx/isis/Registry setvalue \
  Software/VERITAS/VxSvc/CurrentVersion/Security AccessGroups \
  REG_SZ veagrp
```

Note there is not a / between Current and Version in the pathname.

Preparing to Upgrade VERITAS Volume Manager

In the section “DMP Considerations,” the following text appears:

- ◆ If you are upgrading from VxVM 3.1 or an older version, DMP is automatically enabled in VxVM 4.0. You can use vxdiskadm to prevent DMP, if you choose.
- ◆ If you are upgrading from VxVM 3.1.1, and you had prevented or suppressed DMP, VxVM 4.0 retains your DMP setting. Your new installation will have DMP prevented or suppressed.

The correct information is as follows:

DMP is automatically enabled when you upgrade to VxVM 3.5 or higher. If you are upgrading from VxVM 3.1 or an older version, you must upgrade to VxVM 3.5 before upgrading to VxVM 4.0. You can use vxdiskadm to prevent DMP, if you choose.

Upgrading VxVM

In the “Upgrading VxVM” chapter, the following statement occurs in several places:

1. If you have not already obtained and installed a VxVM 4.0 license key, do so now. See “Product Licensing” on page 5 for details.

Note that you need not obtain a license key if you have an existing valid license key for a previous version of VxVM. As describe in the section “Product Licensing”, any existing valid VxVM license keys are accepted for backward compatibility when upgrading to VxVM 4.0. However, you may need to obtain additional licenses for features that are new in VxVM 4.0 or later.



Packaging Differences

In “Packaging Differences in the Sun Microsystems Distribution of VERITAS Products,” the following note should be removed:

- ◆ VERITAS Installation Menu and the install and uninstall scripts are not available.

The VERITAS Installation Menu and scripts are available if you purchased VERITAS Volume Manager from Sun Microsystems.

VERITAS VOLUME MANAGER 4.0 USER'S GUIDE

The following text appears in the section "Monitoring File System Capacity":

For systems other than vxfs, alert thresholds are not persistent (that is, the information is lost when the server restarts).

This should be changed to:

Space alert thresholds are not persistent (that is, the information is lost when the server restarts or the file system is unmounted).

A similar change should be applied to the equivalent text in the Volume Manager and File System online help.

VERITAS VOLUME MANAGER 4.0 USER'S GUIDE - VERITAS ENTERPRISE ADMINISTRATOR

The first page of Appendix B, Available ISP Definitions, states that the ISP objects are defined in `/etc/vx/alloc/alloc_capabilities.txt`.

This should be changed to:

`/etc/vx/alloc/configuration_database.txt`.