VERITAS™

# VERITAS Storage Foundation™ 4.0
## Release Notes

**Solaris**

**Maintenance Pack 1**

July 2004

### 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



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

# Contents





# Important Release Information

## VERITAS Storage Foundation Products

This document provides release information for VERITAS Storage Foundation 4.0 Maintenance Pack (MP) 1 for Solaris. Read this entire document before you install the product. The VERITAS Storage Foundation Release 4.0 MP 1 Solaris 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 MP1 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 15).

Topics in this guide include:





# Getting Help

For assistance with any of the VERITAS products, contact VERITAS Technical Support:

- U.S. and Canadian Customers: 1-800-342-0652
- International: +1-650-527-8555
- Email: support@veritas.com

For license information:

- Phone: 1-800-634-4747 option 3
- Email: amercustomercare@veritas.com
- Fax: 1-650-527-0952

For software updates:

- Email: swupdate@veritas.com

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

http://www.veritas.com

For software updates and additional technical support information, such as TechNotes, product alerts, and hardware compatibility lists, visit the VERITAS Technical Support website at:

http://support.veritas.com



## The VRTSexplorer Diagnostic Program

The VRTSexplorer program is available to assist VERITAS Customer Support engineers in diagnosing technical problems associated with VERITAS products. You can download this program from the VERITAS FTP site or install it from a VERITAS software disc. For more information about the VRTSexplorer program, consult the README file located in the /support directory.

### Downloading VRTSexplorer from the Web

1. Use a web browser or the ftp program to download the VRTSexplorer program from the following URL:

   ```
   ftp://ftp.veritas.com/pub/support/vxexplore.tar.Z
   ```

   Save the file to a temporary directory, such as /tmp, as shown in the example session below.

2. Log in as root on the system that is experiencing the problem. Extract the contents of the downloaded file to the directory /tmp/VRTSexplorer:

   ```
   # cd /tmp
   # zcat vxexplore.tar.Z | tar xvf -
   ```

3. Run the VRTSexplorer program located in the VRTSexplorer directory:

   ```
   # /tmp/VRTSexplorer/VRTSexplorer
   ```

4. When VRTSexplorer prompts for a destination directory for the information that it collects, press Return to accept the default directory /tmp, or enter an alternative path name of your own choice. VRTSexplorer writes the results of its analysis to a compressed tar file named VRTSexplorer_*casenumber_hostname*.tar.z In the specified directory.

5. Use the file upload facility of your web browser, or the ftp program, to transfer the VRTSexplorer output file to the VERITAS Customer Support anonymous FTP site:

   ```
   ftp://ftp.veritas.com/incoming
   ```

6. Telephone VERITAS Technical Support at the number listed under "Getting Help" on page 3. Tell them that you have run VRTSexplorer and provide the name of the file that you transferred to the FTP site.

   Alternatively, if you have already been assigned a call ID number by Customer Support, send email to support@veritas.com and include your case ID number in the subject line.

## Installing VRTSexplorer from the VERITAS Software Disc

The VRTSspt package is included on the VERITAS software discs under the /support directory. To load the software from a software disc:

1. Log in as superuser.

2. Place the VERITAS Storage Foundation software disc into a CD-ROM drive connected to your system.

3. If Solaris volume management software is running on your system, when you insert the VERITAS software disc it is automatically mounted as /cdrom/cdrom0.

4. If Solaris volume management software is not available to mount the VERITAS software disc automatically, you must mount it manually. After inserting the disc, enter:

   ```
   # mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom
   ```

   where c0t6d0s2 is the default address for the CD-ROM drive.

5. Move to the support directory and install the VRTSspt package:

   ```
   # cd /cdrom/cdrom0/support
   # pkgadd -d . VRTSspt
   ```

6. The program is installed in the /opt/VRTSspt directory and takes approximately 500 KB of disk space. To run the program, enter:.

   ```
   # /opt/VRTSspt/VRTSexplorer/VRTSexplorer
   ```

7. Use the file upload facility of your web browser, or the ftp program, to transfer the VRTSexplorer output file to the VERITAS Customer Support anonymous FTP site:

   ```
   ftp://ftp.veritas.com/incoming
   ```

8. Telephone VERITAS Customer Support at the number listed under "Getting Help" on page 3. Tell them that you have run VRTSexplorer and provide the name of the file that you transferred to the FTP site.

   Alternatively, if you have already been assigned a call ID number by Customer Support, send email to support@veritas.com and include your case ID number in the subject line.

9. To remove the VRTSspt package, enter:

   ```
   # pkgrm VRTSspt
   ```



# New Features

The following new features are included with all versions of Storage Foundation 4.0:

- Dynamic LUN Expansion (with VxVM)
- Hardware assisted copy (with VxVM)
- Cross-platform Data Sharing (CDS) (with VxVM)
- File Change Log (with VxFS)
- Extra-Big File Systems (with VxFS)
- Multi-Volume Support (with VxFS)

# 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 1 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 keys 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 12).

### 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 nor 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/VRTS`*`package_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 Administratrator'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 Administratrator'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 Adminstrator'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 Administratrator'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/VRTS`*`product_name`*`/doc` directory after the product installation so that they are available for future reference.

## Available Storage Foundation Patches for MP1

With VERITAS Storage Foundation 4.0 MP1, 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 | 115217-03 | VRTSvxvm |
| | 116681-01 | VRTSvxvm |
| | 117094-02 | VRTSvmpro |
| | 117095-01 | VRTSmuvmp (language package) |



| Operating System | Patch Number | Affected VERITAS Package |
|---|---|---|
| | 116684-01 | VRTSfspro |
| | 117276-01 | VRTSalloc |
| | 117277-01 | VRTSmualc |
| | 116685-01 | VRTSmufsp (language package) |
| | 116686-01 | VRTSjafsc |
| | 115209-12 | VRTSob |
| | 115210-12 | VRTSobgui |
| | 115213-14 | VRTSmuobg (language package) |
| | 115212-13 | VRTSmuob |
| | 117093-01 | VRTSfas |
| | 117096-01 | VRTSjafas (language package) |
| | 117504-01 | VRTSjacs (language package) |
| Solaris 7 only | 116687-01 | VRTSvxfs |
| | 117097-01 | VRTSodm |
| Solaris 8 only | 116688-01 | VRTSvxfs |
| Solaris 9 only | 116689-01 | 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 MP1" on page 15 for details):

   a. Verify the status of the VEA Service:

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

   b. Stop the VEA server:

   ```
   # /etc/init.d/isisd stop
   DBED: Successfully unloaded the Storage Foundation Provider
   4.0 for Oracle
   ```

   c. Again verify the status of the VEA Service:

   ```
   # /etc/init.d/isisd 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. After installing the patches, restart the VEA Service.

   a. Verify the status of the VEA Service:

   ```
   # /etc/init.d/isisd status
   Current state of server : NOT RUNNING
   ```

   b. Start the VEA Service:

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

   c. Verify the status of the VEA Service again:

   ```
   # /etc/init.d/isisd 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 MP1" on page 15 for a list of the patches that were required for your operating system.

# Storage Foundation Incidents Closed for 4.0 MP1

The following Storage Foundation incidents were closed for 4.0 MP1.





## Volume Manager Closed Incidents

| Incident Number | Description |
|---|---|
| e148283 | Unable to encapsulate T4 boot disk connected thru SAN switches. |
| e148283 | Unable to encapsulate a VLun root device on vxvm 3.5. |
| e153608 | Unwanted Console messages for Plexes / Volumes on CVM Master node. |
| e156996 | CVM: Need not to duplicate effort in setting mirror recovery mode. |
| e156996 | VM3.5: Default volume recovery is not happening after a disk failure. |
| e153427 | *4.0* scvxinstall gives errors when run on a node with MPxIO boot disk. |
| e155060 | *4.0b2* vxdiskadm truncates default disknames to 9 characters. |
| e155167 | *4.0b3* CD is not default disk format when initializing disk with vxdisksetup. |
| e155388 | *4.0b3* vxprint shows the volume in enabled state while in fact it is not. |
| e156395 | libvxvm_get_disks() problem resulting in sun cluster malfunction. |
| e156815 | vxio error messages are incomplete. |
| e144728 | Bad trap panic in dmp_get_paths_on_ctlr. |
| e157384 | vxdmp does not failback after p path is restored. |
| e157518 | License not autogenerated with photon, VM 3.5. |
| e157589 | vxplex error messages displayed on console. |
| 135566 | vxdiskadm - adddisk option failed to add disk into the shared DG. |
| e155821 | Failback does not work for active/passive failback arrays on SUNs T3/T4 MPXIO. |
| e206689 | Unable to add a new disk to a shared disk group using vxdiskadm. |
| e156292 | bootdg volumes on external fabric disks come up disabled after reboot. |
| e158566 | CVM cvmjoin state should be reset after call to volcvm_master in vxclust step2. |
| (e156292 | Upgrade from 3.5 to 4.0 caused the rootdg volumes to become disabled. |

## File System Closed Incidents

| Incident Number | Description |
|---|---|
| 137488 | The behavior of -t option of the fsadm command is corrected. The time limit of the -t option is now a cumulative total time. |
| 140167 | The fstyp command now tries to interpret an existing device name as a clone pseudo device. The fs_is_clonedev() function now validates a split devname before calling the fs_open() function. Block and character device name conversions for the ",raw" postfix format were added. |
| 140326 | The sticky bit behavior is different between Solaris 8 UFS and VxFS 3.5, including MP1 and MP2.VxFS on Solaris will return EPERM, regardless of the type of file being removed. |
| 141372 | The fsck command exited if it encountered an IFEAU file that contained only values of zero while replaying a log. The fsck command should instead have proceeded to do a full fsck if the -y option was specified. In a recent customer escalation, the customer's file system's IFEAU and FAUS files were corrupted and contained only values of zero due to hardware issues. After running fsck -y on the file system, fsck exited during the log replay. |
| 141795 | The vxdump command failed with SIGSEGV when used on a mount point of an unmounted file system. The vxdump command succeeded when used on a raw device. |



## Volume Replicator Closed Incidents

| Incident Number | Description |
|---|---|
| e146194 | In this release, VVR allows the size of the DCM to be larger than 32 kilobytes so that there can be smaller region sizes for large volumes. This alleviates the problem where automatic synchronization or DCM replay would not complete for large volumes. |
| e156025 | If hosts with 32-bit versions of the operating system were running for more than 248 days, the RLINKs would not connect after they are disconnected. |
| e156268 | If the Primary SRL overflowed during checkpoint attach and DCM replay was initiated during this time, the Secondary SRL could potentially overflow. This could cause data corruption on the Secondary. |
| e157849 | When the SRL was dissociated and associated to the RVG without any change to the SRL, the checkpoint area of the SRL header was not being cleared. This issue caused VVR to erroneously interpret stale checkpoint data from the header. |
| e157977 | In certain situations, where there were frequent connects and disconnects of the RLINKs, VVR could cause a panic. |
| e158025 | In certain situations, when readbacks are in progress, VVR would write to invalid memory and cause a panic. |

## VEA Closed Incidents

| Incident Number | Description |
|---|---|
| 139665 | VEA unexpected error occurred in Search screen. |
| 140157 | Searching for objects not of a specified type behaves incorrectly. |
| 141085 | Snap reattach may lock GUI. |
| 141260 | Add api to allow for search help. |
| 141461 | Disks view rendering is too slow. |
| 139032 | Vxdg -n should not allow deport a DG in the name of bootdg/nodg. |

| Incident Number | Description |
|---|---|
| 139665 | VEA unexpected error occurred in Search screen. |
| 141997 | Java exception when user tries to do File->Open from the 'Disk View.' |
| 142000 | Java exception when user tries to click on 'Disk Groups.' |
| 142722 | Refresh op needs to register dgsj license update. |
| 140196 | GUI calls init disk for each disk sync. |
| 143604 | Provider does not ever unregister pool interface on disk. |
| 143454 | Relayout status monitor quits reporting progress. |
| 144053 | VxTF license check needs modifying. |
| 140819 | Controller display is blank and properties displays enclosure incorrectly. |
| 143094 | I18N VEA online help for link issues. |



# Open Incidents and Suggested Solutions

The following open incidents and suggested solutions are noted for the 4.0 MP1 release.

## Section 1 - Installing the VxVM 4.0 Patch 1

### 1.1 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.

### 1.2 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:vxdmp: 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.

### 1.3 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 2 - 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 3 - vxassist Relayout Considerations

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

## Section 4 - 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 5 - 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 6 - 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 7 - 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. Start VxVM:

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

5. Start the VxVM volumes:

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

6. Use the `mount` command to remount the file system.

## Section 8 - 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 9 - 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 &lt;diskgroup&gt; -p -F "%snap_rid"
&lt;plexname&gt
# vxprint -g &lt;diskgroup&gt; -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 &lt;original_volume&gt; &lt;plex1&gt; [&lt;plex2&gt;
... ]
```

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 &lt;original_volume&gt; 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 &lt;original_volume&gt;
```

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

```
# vxplex att &lt;original_volume&gt; &lt;staleplex1&gt;
[&lt;staleplex2&gt; ...]
```

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 10 - 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 be appear:

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



## Section 11 - 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 12 - Potential Solaris Patch Issues

**Issue 1 Do Not Use Solaris 2.8 Patch 110934-10 or Solaris 2.9 Patch 113713-01**

Solaris 2.8 patch 110934-10 and Solaris 2.9 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 2.8) or patch 113713-01 (Solaris 2.9) 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 2.8, patch 110934-10, refer to TechNote 252441:

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

For Solaris 2.9 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 2.8, 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 Superceded by 108993-XX**

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

Sun has superceded 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 13 - 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/rdsk/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 boththe 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 14 - 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.

### 14.1 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.

### 14.2 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 15 - 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 16 - Problem When Removing VxVM Package

(Ref incident 105230)

A panic may occur when removing the VxVM package (VRTSvxvm) using `pkgrm` on Solaris 9.

**Note** This problem is isolated to Solaris 9 only.

Suggested Solution: The problem can be avoided by running "`devfsadm stop`" before the `pkgrm` command:

```
/etc/init.d/devfsadm stop
pkgrm VRTSvxvm
```

## 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** Note: VxVM does not support SCSI3/PGR on HDS9200 in A/P mode on Solaris 2.9.

## Section 18 - Moving a Disk Device From One Physical Array Slot to Another

(Ref incident 137607)

Problem: If a disk device is moved (or swapped) from one physical array slot to another, then it is possible for two different devices to be incorrectly merged into a single DMP node - ultimately leading to possible data corruption.

**Note** This situation may occur even after the disks have been re-discovered (for example, by means of the `devfsadm -C` command) and re-scanned (for example, by using the `vxdctl enable` command).

For example; suppose the disk at slot 2 is powered off. The host will not then see any device in slot 2. Then:

```
# vxdisk list
DEVICE       TYPE      DISK         GROUP        STATUS
c1t1d0s2     sliced    -            -            error
c1t1d0s4     simple    c1t1d0s4     rootdg       online
c3t1d0s2     sliced    -            -            online
c3t3d0s2     sliced    -            -            online
c3t4d0s2     sliced    -            -            online
c3t5d0s2     sliced    -            -            online
c3t6d0s2     sliced    -            -            online
```

**Note** Note that c3t2d0 is missing. Also notice that DMP device c3t5d0 has only one entry.

Then:

```
# vxdmpadm getsubpaths dmpnodename=c3t5d0s2
NAME         STATE     PATH-TYPE  CTLR-NAME  ENCLR-TYPE   ENCLR-NAME
c3t5d0s2     ENABLED     -        c3          Disk          Disk
Suppose, now:
disk c3t5d0 at slot 5 is powered off and removed from slot 5
disk c3t2d0 at slot 2 is removed from slot 2
```

The two disks are swapped, so that the disk that was in slot 5 is now in slot 2, and the disk that was in slot 2 is now in slot 5. Both disks are powered on.

If rediscover and rescan are now performed, in an attempt to get DMP and the OS see the new disk arrangement, it is possible that DMP node c3t5d0s2 has become "confused," and has incorrectly multi-pathed BOTH the old device (c3t2d0s2) as well as the new device (c3t5d0s2) as the same device. Since these two disks are physically different, and should never be multi-pathed into the same DMP node, data corruption will occur if I/O goes through this DMP path.

```
# devfsadm -C
# vxdctl enable
# vxdisk list
DEVICE       TYPE      DISK         GROUP        STATUS
c1t1d0s2     sliced    -            -            error
c1t1d0s4     simple    c1t1d0s4     rootdg       online
c3t1d0s2     sliced    -            -            online
c3t3d0s2     sliced    -            -            online
c3t4d0s2     sliced    -            -            online
c3t5d0s2     sliced    -            -            onlin
c3t6d0s2     sliced    -            -            online

# vxdmpadm getsubpaths dmpnodename=c3t5d0s2
NAME         STATE     PATH-TYPE  CTLR-NAME  ENCLR-TYPE   ENCLR-NAME
=================================================================
c3t5d0s2     ENABLED     -        c3         Disk         Disk
c3t2d0s2     ENABLED     -        c3         Disk         Disk
# dd if=/dev/zero of=/dev/vx/rdmp/c3t5d0s4 bs=128k count=1000
1000+0 records in
1000+0 records out
# iostat
  device    bps     sps    msps
....
  c3t5d0   4076    15.9     1.0
....
  c3t2d0   4076    15.9     1.0
```

Notice that I/O is going through both devices, which is incorrect, and will lead to unexpected data corruption.

Suggested Solution: Just before changing or swapping the disk slots, multipathing should be disabled. The library definition and/or the jbod definitions should be removed.

The disks will then be added as single-pathed disks, and existing multi-pathing will be lost as a result of dmpnodes being split. Multipathing can be later re-enabled, and `vxdctl enable` will merge the correct dmpnodes.



▼ **To remove the definition:**

1. Exclude the library or remove the JBOD definition:

```
# vxddladm excludearray libname=<> or
# vxddladm rmjbod vid=<>
```

2. Scan the disks for VxVM:

```
# vxdisk scandisks or vxdctl enable
```

The purpose of steps 1 and 2 is to force all the dmpnodes from the array to be single-pathed.

3. Include the library or JBOD definition.

# vxddladm includearray libname=<> or # vxddladm addjbod vid=<>

4. Scan the disks for VxVM:

```
# vxdisk scandisks or
# vxdctl enable
```

The purpose of steps 3 and 4 is to re-group the paths from scratch to again form multi-pathed dmpnodes for the luns.

**Note** These steps will also fix the problem, if it has already occurred.

For example: If you have physically moved a particular XP1024 disk (corresponding DDL library libvxxp256.so) to a different array slot, then you need to carry out the following steps:

```
# vxddladm excludearray libname=libvxxp256.so
# vxdctl enable  (or vxdisk scandisks)
# vxddladm includearray libname=libvxxp256.so
# vxdctl enable  (or vxdisk scandisks)
```

**Caution** Once multi-pathing is disabled, the same disk is exposed as two different disks with the potential of the disk being initialized through the other path. Therefore, this state should not be continued for long. Also, the system should not be rebooted because the DG import would then fail due to duplicate disk ids.

## Section 19 - 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 20 - 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 21 - 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 144279

In SxRT 3.5 MP3 Japanese version, using EMC FAS or VxTF Toolkit "OVERRIDE_ID" is displayed instead of the text “Do you want to override array policies?”.

### Incident 144723

VxVmCommon_NOTASCII should be translated to say:

```
"Volume Manager takes ASCII characters to set the comment.
Please enter ASCII characters for the comment."
```

### Incident 144081

The resource string "VxVmCommon_COMMENTTOOLONG" is not included in localized resource file. The string gives the message:

```
"Value too long for comment field. The maximum length for comment
is 20."
```

### Incident 144258

In the Japanese localized VEA graphical user interface, when a user selects a disk group and right clicks to rename the group, the users does not see a mnemonic character inside. The user also sees a message:

```
"Mnemonic identified by key" 'DISK_GROUP_ID' not found in
'vrts.resource.ja.VmResourceBundle_ja.properties'.
```

### Incident 144271

In the Japanese localized version of VEA, the word COMMA is displayed instead of the “,” in the **VmDisableControllerDialog**.

**Incident 144274**

The Japanese localized VEA shows:

```
"VmRecoverDisk_Confirm"
```

instead of the text

```
"Are you sure you want to recover the disk <0>?"
```

**Incident 144277**

The Japanese localized VEA shows "VmDisableControllerDialog_LASTPATH_WARNING" instead of the text

```
"Device(s) {0} can be accessed only through this controller. Are
you sure you want to disable this controller<1>?
```

**Incident 144781**

In the French VEA, users may not see the disk name(s) when selecting a list of disks to create a new disk group in the **New Dynamic Disk Group Wizard** because the text area in the **Add Disk** page is very small.

If a disk containing the file system is initialized, the data on that disk will be overwritten. This issue also appears in the **Add Disk to Disk Group Wizard**.

**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 22 - Installation With vxinstall Does Not Work With MPXIO Boot Device

With fixes in IR#119997, VxVM is able to encapsulate the MPxIO boot disk and mirror to another MPxIO disk. However the boot encap and mirror process creates an alias at OBP using the `eeprom` program which may not work in some cases.

Problem - Rebooting with command: boot

```
Boot device: /scsi_vhci/ssd@g20000004cf0a32a5:a  File and args:
kadb
Evaluating: boot
Can't open boot device
```

Suggested Solution - Change the boot alias which is created as below:

```
vx-mpdg01 /scsi_vhci/ssd@g20000004cf0a33aa:a
```



to

```
vx-mpdg01
/pci@1f,0/pci@1/SUNW,qlc@1/fp@0,0/disk@w22000004cf0a33aa,0:a
```

To get the correct device string, follow these steps:

1. Before starting the boot encap procedure get the bootdisk information using "`df -k`" command.

```
# df -k

Filesystem            kbytes    used   avail capacity  Mounted on
/dev/dsk/c5t20000004CF0A33AAd0s0
                     5056862 1630241 3376053    33%    /
/proc                      0       0       0     0%    /proc
mnttab                     0       0       0     0%    /etc/mnttab
fd                         0       0       0     0%    /dev/fd
swap                 1039912      32 1039880     1%    /var/run
dmpfs                1039880       0 1039880     0%    /dev/vx/dmp
dmpfs                1039880       0 1039880     0%    /dev/vx/rdmp
swap                 1039952      72 1039880     1%    /tmp
```

2. Get the physical properties of the bootdisk using "`luxadm disp`" command.

```
# luxadm disp /dev/rdsk/c5t20000004CF0A33AAd0s2

DEVICE PROPERTIES for disk: /dev/rdsk/c5t20000004CF0A33AAd0s2
Status(Port A):       O.K.
Status(Port B):       O.K.
Vendor:               SEAGATE
Product ID:           ST318304FSUN18G
WWN(Node):            20000004cf0a33aa
WWN(Port A):          21000004cf0a33aa
WWN(Port B):          22000004cf0a33aa
Revision:             0726
Serial Num:           3EL0CHSJ0000
Unformatted capacity: 17274.674 MBytes
Read Cache:           Enabled
Minimum prefetch:   0x0
Maximum prefetch:   0x76
Location:            In slot 1 in the Rear of the enclosure named:
A5K#3
Device Type:          Disk device
Path(s):
/dev/rdsk/c5t20000004CF0A33AAd0s2
/devices/scsi_vhci/ssd@g20000004cf0a33aa:c,raw
```

```
Controller           /devices/pci@1f,0/pci@1/SUNW,qlc@1/fp@0,0
<<== NOTE_1
    Device Address              22000004cf0a33aa,0 <<=== NOTE_3
    Host controller port WWN    210000e08b061bde
    Class                       primary
    State                       ONLINE
   Controller           /devices/pci@1f,0/pci@1/SUNW,qlc@3/fp@0,0
<<== NOTE_2
    Device Address              21000004cf0a33aa,0
    Host controller port WWN    210000e08b081603
    Class                       primary
    State                       ONLINE
```

3. Get the physical controller info with out "`/devices`" from NOTE_1 or NOTE_2

```
Controller 1 : /pci@1f,0/pci@1/SUNW,qlc@1/fp@0,0
Controller 2 : /pci@1f,0/pci@1/SUNW,qlc@3/fp@0,0
```

4. Get the device address from NOTE_3 from the above example output.

```
22000004cf0a33aa,0
```

5. Prepare the bootdevice alias as below.

**Note** Use the physical controller information taken from the boot-device alias in the `eeprom` command.

   a. Get the controller info from step 2

   b. Append /disk@w to above 4.i string

   c. Append device address to above 4.ii string from step 4.

```
/pci@1f,0/pci@1/SUNW,qlc@1/fp@0,0/disk@w22000004cf0a33aa,0:a
```

6. Save the final string from step 5.

7. Complete the encapsulation process for the MPxIO boot. Rebooting with command: `boot`

```
Boot device: /scsi_vhci/ssd@g20000004cf0a32a5:a  File and args:
kadb
Evaluating: boot
Can't open boot device
```

8. Set up the alias for the boot-device with string from step 4 using below command.



```
ok nvalias vx-mpdg01
/pci@1f,0/pci@1/SUNW,qlc@1/fp@0,0/disk@w22000004cf0a33aa,0:a
```

**9.** Reset and reboot the machine.

## Section 23 - Using powervxvm With Volume Manager 4.0MP1

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 24 - 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 25 - vradmin Print Commands

Ref Incident e208339. The `vradmin` print commands display incorrect messages when the vradmin 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 26- 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 27 - 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 28 - 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 29 - 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.

## Special Install/De-Install Instructions:

You must use the shutdown command to reboot the system after patch installation or de-installation:

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

## Do Not Use Solaris 2.8 Patch 110934-10 or Solaris 2.9 Patch 113713-01

Solaris 2.8 patch 110934-10 and Solaris 2.9 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 2.8) or patch 113713-01 (Solaris 2.9) 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."
```

For Solaris 2.8, patch 110934-01, refer to TechNote 252441:

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

For Solaris 2.9 patch 113713-01, the suggested solution is to rename the space file before running `patchadd`, as follows:

When:

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

then:

```
# patchadd <THIS_PATCH_ID>
```

For Solaris 2.8, 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.