VERITAS

VERITAS Storage Foundation™ 4.1

Release Notes

Solaris x64 Platform Edition

November 2005

Disclaimer

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

VERITAS Legal Notice

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

Third-Party Legal Notices

Sun Microsystems Trademarks

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

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

Contents

VERITAS Storage Foundation Release Notes

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

This product suite is activated by a license key, which must be obtained before installing the product. For information on obtaining a license key, see the *VERITAS Storage Foundation Installation Guide*.

Note For the latest information on updates, patches, and software issues regarding this release, see the following TechNote on the VERITAS Technical Support website: http://support.veritas.com/docs/278691

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

Product Descriptions

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

Storage Foundation Version	Products
Storage Foundation	VERITAS Volume Manager VERITAS File System

VERITAS Volume Manager

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

VxVM 4.1 Licenses

The following features are supported by a Full VERITAS Volume Manager 4.1 license:

- Concatenation
- Cross-platform Data Sharing (CDS)
- Device Discovery Layer (DDL)
- DRL logging for mirrors
- Dynamic Multipathing (DMP)
- Hot sparing
- Hot-relocation
- Intelligent Storage Provisioning
- Mirroring
- Mirroring plus striping
- Multiple disk groups
- Online data migration
- Online relayout
- RAID-5
- RAID-5 logging
- Smartsync

- Spanning
- Storage Expert
- Striping
- Striping plus mirroring
- Third-mirror break-off volume snapshots
- VERITAS Enterprise Administrator (VEA)
- Volume resizing

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

```
# vxdctl license
```

VERITAS File System

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

The primary purposes of a file system are to:

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

System Requirements

Solaris Operating System Requirements

The VERITAS Storage Foundation Standard 4.1 product for the x64 platform operates on the Solaris 10 (32- and 64-bit) Operating System.

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

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

Solaris Patch Requirements

Solaris Patches for Storage Foundation

Some required system patches may already be present in your operating system. You should check to see if your system already contains the patches needed. Use the command `showrev -p` to display the patches included on your system. For more information, see the `showrev`(1M) manual page.

If the patches shown in the required list are not already installed, go to http://sunsolve.sun.com to download them. You need to install the appropriate patches and then reboot.

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

Sun Patch Number (or higer)	Description	Notes
119375-03	SunSolve patch for SCSI disk driver	Obtain from http://sunsolve.sun.com
119131-09	SunSolve patch for Fibre Channel driver	Obtain from http://sunsolve.sun.com
118344-03	SunSolve patch for Fault Manager	Obtain from http://sunsolve.sun.com

Sun Patch Number (or higer)	Description	Notes
118844-19	Sunsolve patch which addresses kernel issues found since the initial Solaris 10 release.	You need to install patch 118844-19 or later before installing the VERITAS Storage Foundation. Obtain from http://sunsolve.sun.com

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

Installation Storage Foundation HA Configuration Requirements

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

Hardware Configuration Information

These are the minimum requirements.

Systems

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

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

Arrays

Array	Firmware Version
StorEdge 3310	3.25W
StorEdge 3510	4.11I

The array support libraries (ASLs) for these arrays are located on CD 1 as gzipped tar files (*ASL_package_name*`.tar.gz`).

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

Host Bus Adaptors

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

Other Devices

Device	Version
StorEdge Configuration Command Line Interface (sccli)	1.6.2

Software Limitations

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

VERITAS Volume Manager Software Limitations

Using the vxcdsconvert Utility

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

Local Zone Support

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

Temporary File Systems Used by DMP

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

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

VERITAS File System Software Limitations

No Support for Certain Manual Page options

The following manual pages options are not supported:

Manual Pages	Unsupported Options
`mkfs_vxfs`(1M)	The `version` option does not support disk layout Version 4 and 5.
`mount_vxfs`(1M)	`seconly`, `qio`, `nqio`, `qlog`, `ckpt`, `cluster`, `crw`

Unsupported Features

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

- Clustering feature of VERITAS Volume Manager
- Co-existence of DMP with any other multipathing driver (for example, MPxIO)
- Co-existence with any volume manager other than VERITAS Volume Manager
- Concurrent I/O (CIO)
- Data Management API (DMAPI)
- Disk layout Version 5 and previous versions
- Dynamic LUN expansion
- Dynamic Multipathing (DMP) support for active/passive arrays
- Extended Copy (ECopy)
- Extensible Firmware Interface (EFI) disks
- File Promotion
- Fileset Quotas
- FlashSnap™ features:
  - FastResync (fast mirror resynchronization)
  - Disk group split/join (DGSJ)
  - Storage Checkpoint feature of VERITAS File System
- Full-sized and space-optimized instant snapshots
- Solaris Local Zones
- Oracle Disk Manager (ODM)
- Root disk encapsulation and associated features
- Quality of Storage Service
- Quick I/O and cached Quick I/O features of VERITAS File System
- Quick Log
- SCSI III PGR
- Support for moving between disk group versions
- VERITAS File System Storage Checkpoints

- VERITAS File System Storage Checkpoints Rollback
- VERITAS Volume Replicator (VVR)
- VERITAS Storage Foundation Cluster File System
- `vxspcshow` command

Hardware Issues

Potential Data Corruption with the LSI Logic 1030 Controller

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

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

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

```
prtconf -vp
```

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

```
revision-id: 00000007
```

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

```
revision-id: 00000008
```

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

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

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

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

Software Issues

Installation Script Issue

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

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

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

VERITAS Volume Manager Software Issues

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

Disk Controller Firmware Upgrades

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

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

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

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

1. Disable the plex associated with the disk device:

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

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

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

   For the other controller on the HBA, enter:

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

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

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

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

5. Enable the plex associated with the device:

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

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

Utility Issues

Adding a Log and Mirror to a Volume

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

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

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

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

[13488]

Using vxdiskadm to Replace a Failed Disk

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

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

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

Replacement of the old_layout Attribute

The `vxdisksetup` command gives the error message `Attribute unrecognized` when the `old_layout` attribute is used to make a disk into a VxVM controlled disk. The `old_layout` attribute is no longer supported. Use the `noreserve` attribute in its place. [121258]

Using vxvol and vxmend with Layered Volumes

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

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

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

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

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

[134932]

Growing or Shrinking Layered Volumes

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

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

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

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

Maximum Size of a VxVM Volume

VxVM supports volume lengths up to 2^63-1 disk sectors when using VERITAS-specific `ioctl` calls. However, system calls such as `seek`, `lseek`, `read` and `write` are limited to a maximum offset that is determined by the operating system. For a system that supports large files, this is usually 2^63-1 bytes. Otherwise, the maximum offset value is usually 2^31-1 bytes (1 byte less than 2 terabytes). [141024]

Device Issues

Supported Hardware

See "Hardware Configuration Information" on page 5 for a list of supported systems, host bus adapters, disk arrays and other devices. This supplements the information given in the *VERITAS Volume Manager Hardware Notes*. Other array hardware and third-party multipathing driver coexistence are not supported in this release, and any mention of such should be treated as examples only.

Miscalculation of Disk Space

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

Removing a Disk Hangs VxVM

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

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

Adding Disks to VxVM on a Solaris x64 System

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

- On a Solaris SPARC system, the start of the Solaris partition (which may contain a primary boot executable and boot block in addition to the VTOC and any disk slices) is located in cylinder 0. The whole disk is accessed using the device `c#t#d#s2`.

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

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

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

Removing a Disk From VxVM Control

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

```
# vxdiskunsetup -C daname
```

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

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

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

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

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

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

Encapsulation of Disks with Insufficient Space for a Private Region

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

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

Hot-Relocation Issues

Impact of Hot-Relocation on Performance

Hot-relocation does not guarantee the same layout of data or performance after relocation. It is therefore possible that a single subdisk that existed before relocation may be split into two or more subdisks on separate disks after relocation (if there is not enough contiguous space on a single disk to accommodate that subdisk). [14894]

Disk Information in Notification Messages

When a disk failure occurs, the hot-relocation feature notifies the system administrator of the failure and any relocation attempts through electronic mail messages. The messages typically include information about the device offset and disk access name affected by the failure. However, if a disk fails completely or a disk is turned off, the disk access name and device offset information is not included in the mail messages. This is because VxVM no longer has access to this information. [14895]

DMP Issues

Usage of dmp_failed_io_threshold Parameter

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

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

- A portion of the I/O path below DMP fails in such a way that the error is not returned to DMP before 10 minutes have elapsed. Examples of failures that have been seen to cause this condition, include bad hardware (HBA, FCOT, GBIC, Switch Port, Array Controller) and HBA reconfigured to retry endlessly.

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

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

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

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

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

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

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

```
dmp_failed_io_threshold=57600
```

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

After changing the value, reboot the system.

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

Data Integrity Issues

Disks with Write-Back Caches

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

Intelligent Storage Provisioning Issues

Creating Application Volumes

To create application volumes successfully, the appropriate licenses must be present on your system. Vendors of disk arrays may also provide capabilities that require special licenses for certain features of their hardware. [Sun Bug ID 4948093, 137185]

Miscellaneous Issues

Auto-import of Disk Groups

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

If it is important that a disk group *not* be auto-imported when the system is rebooted, the disk group should be imported temporarily when the intention is to deport the disk group (for example, in HA configurations). Use the `-t` flag to `vxdg import`. [13741]

Volumes Not Started Following a Reboot

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

Workaround: Check the state of the volumes before starting the application, or place a sleep (`sleep` *sec*) before the last invocation of `vxrecover`. [14450]

Forcibly Starting a Volume

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

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

However, try to determine what caused the problem before you run this command. It is likely that the volume needs to be restored from backup, and it is also possible that the disk needs to be replaced. [14915]

Failure of Memory Allocation

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

Using Long Device Paths with Sun Online:Backup

The Sun Online:Backup™ facility does not accept the long device path names for volumes. A limitation of Online: Backup is that it does not accept device paths longer than 24 characters.

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

Number of Columns in a RAID-5 ISP Volume

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

VEA Issues

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

Accessing the Task Log

The task log accessed from the Log tree is not supported. At this time, entries are written to the log file in `/var/vx/isis/command.log`. [76683, 97076]

Setting a Comment on an ISP Volume

If you create a new ISP volume by right-clicking on a user template and selecting the New Volume menu item, a comment that you specify to the Create Volume Dialog is not set on the volume. To specify a comment for the newly created volume, select the volume, choose **Properties** from the popup menu, enter a comment in the **Comment** field and then click **OK**. [137098]

Disk Group Creation Failure with Duplicate Disk ID

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

Incorrect vxpool command

The VEA GUI may incorrectly show the `-p` option as an argument to the `vxpool list` command, although the command is not actually invoked.
[Sun Bug ID 4932404, IR 135566].

Internationalization Issues

Some ISP Attributes Have Not Been Translated

The Intelligent Storage Provisioning (ISP) window for annotating a disk is not fully localized. In particular, auto-discovered attributes such as DiskGroup and Enclosure are not translated. [139162]

Inaccuracies in ISP Attribute Fields

The ISP User Template Wizard shows two attribute value fields rather than one attribute value and one attribute name field. [139762]

Adding Unsupported Disk Arrays to the JBOD Category

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

Use this procedure when adding an unsupported disk array after VxVM has been installed:

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

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

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

   ```
   # /etc/vx/diag.d/vxdmpinq /dev/rdsk/c1t20d0s2
   Vendor id (VID)  : FUJITSU
   Product id (PID): FU318404LSUN18G
   Revision         : 8507
   Serial Number    : 0025T0LA3H
   ```

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

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

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

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

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

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

   ```
   # vxdctl enable
   ```

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

```
# vxddladm listjbod
```

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

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

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

```
# vxdmpadm listenclosure all
ENCLR_NAME      ENCLR_TYPE    ENCLR_SNO     STATUS
=========================================================
OTHER_DISKS     OTHER_DISKS   OTHER_DISKS   CONNECTED
Disk            Disk          DISKS         CONNECTED
```

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

```
# vxdisk list
DEVICE     TYPE        DISK      GROUP      STATUS
Disk_0     auto:none   -         -          online invalid
Disk_1     auto:none   -         -          online invalid
...
```

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

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

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

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

VERITAS File System Software Issues

VxFS NFS Server Runs out of Incore Inodes

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

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

API for Manipulating Disk Quotas

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

VERITAS Enterprise Administrator Issue

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

Available Documentation

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

Installing documentation and online manual pages is optional.

Release Notes and Installation Guides

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

Release notes for component products of the VERITAS Storage Foundation software are located under the `storage_foundation/release_notes` directory of the VERITAS software disc.

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

VERITAS Storage Foundation Guides

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

Guides in VERITAS Storage Foundation Documentation Set

Guide Title	Filename
VERITAS Storage Foundation Release Notes (this document)	`sf_notes.pdf`
VERITAS Storage Foundation Installation Guide	`sf_install.pdf`
VERITAS File System Administrator's Guide	`vxfs_admin.pdf`
VERITAS File System Programmer's Reference Guide	`vxfs_ref.pdf`
VERITAS Volume Manager Administrator's Guide	`vxvm_admin.pdf`
VERITAS Enterprise Administrator (VEA 500 series) Getting Started	`veax5_getting_started.pdf`
VERITAS Storage Foundation Intelligent Storage Provisioning Administrator's Guide	`sf_isp_admin.pdf`
VERITAS Storage Foundation Cross-Platform Data Sharing Administrator's Guide	`sf_cds_admin.pdf`
VERITAS Volume Manager Hardware Notes	`vxvm_hwnotes.pdf`
VERITAS Volume Manager Troubleshooting Guide	`vxvm_tshoot.pdf`

Online Manual Pages

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

Getting Help

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

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

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