



VERITAS[®] Volume Manager
Release Notes
Release 3.0.2 - Binary

PLEASE READ THIS BEFORE INSTALLING THE SOFTWARE!

Solaris
October, 1999
P/N 100-001217

© 1998 - 1999 VERITAS® Software Corporation. All rights reserved.

TRADEMARKS

VERITAS, VxVM, VxVA, VxFS, and FirstWatch are registered trademarks of VERITAS Software Corporation in the United States and other countries.

VERITAS Volume Manager, VERITAS File System, VERITAS NetBackup, VERITAS HSM, VERITAS Media Librarian, CVM, VERITAS Quick I/O, VxSmartSync, and the VERITAS logo are trademarks of VERITAS Software Corporation.

Other products mentioned in this document are trademarks or registered trademarks of their respective holders.

Introduction

This document provides release information for the VERITAS Volume Manager™ (VxVM®) Release 3.0.2. This Release includes the Volume Manager Storage Administrator (VMSA) Release 3.0.3 graphical user interface.

Note: This version of Volume Manager supports VMSA version 3.0.3 only. If you have an older version of VMSA, you must install VMSA version 3.0.3.

The Volume Manager Storage Administrator consists of a server and a client. The Storage Administrator server must be run on a UNIX machine running Solaris Release 2.5.1 or higher. The Storage Administrator client can be run on any machine that supports the Java 1.1 Runtime Environment (including Solaris, HP-UX, Windows NT, Windows 98, or Windows 95).

This release of the Volume Manager supports and has been tested on:

- Solaris 2.5.1
- Solaris 2.6
- Solaris 7 (32- and 64-bit)

The Storage Administrator server and client have been tested on Solaris Release 2.5.1 and higher. The Storage Administrator client has also been tested on Windows NT, Windows 98, and Windows 95.

Note: Before you install the packages, read this entire document.



Contents

The following topics are covered in this document:

- Getting Help
- New Features and Changes
 - Volume Manager
 - Storage Administrator
 - Cluster Functionality (Optional)
- End of Support Statements
- Installing the Volume Manager
 - Installing Volume Manager Packages
 - Setting Up the Volume Manager
 - Setting Up the Storage Administrator
- Upgrading the Volume Manager
- Documentation
 - Displaying Documentation Online
 - Printing Documentation
- Software Problems Fixed In VxVM Release 3.0.2
 - Software Problems Fixed Prior to VxVM Release 3.0.2
- Software Problems Fixed in Storage Administrator Release 3.0.3
- Software Limitations and Problems in VxVM Release 3.0.2
 - Installation Issues
 - Upgrade Issues
 - Utility Issues
 - Device Issues
 - Hot-Relocation Issues
 - DMP Issues
 - Cluster Functionality Issues



-
- Miscellaneous Issues
 - Solaris Issues
 - Software Limitations and Problems in Storage Administrator Release 3.0.3
 - Encapsulating and Mirroring the Root Disk
 - Booting From DMP Devices
 - VxVM and Multi-Host Failover Configurations

Getting Help

For license information or information about VERITAS® service packages, contact VERITAS Customer Support.

US Customers: 1-800-342-0652

International Customers: +1 (650) 335-8555

Fax: (650) 335-8428

Electronic mail: support@veritas.com



New Features and Changes

Note: Volume Manager Release 3.0.2 and Storage Administrator Release 3.0.3 are Year 2000 compliant. For additional information, see the Year 2000 certification statement on the VERITAS Web site (www.veritas.com).

Volume Manager

VxVM incident (bug) fixes are listed in “Software Problems Fixed Prior to VxVM Release 3.0.2.”

New features with VxVM 3.0.2 include:

- Support for Solaris 2.5.1, Solaris 2.6, and Solaris 7 (32- and 64-bit)
- SDS to VxVM Conversion

If you have Solstice™ DiskSuite™ (SDS) configured on your system, VERITAS offers tools to assist you in converting your system to Volume Manager. VERITAS has a conversion tool which enables a *data conversion in-place* thereby allowing user data to remain unchanged while SDS structural data is replaced with VERITAS Volume Manager structural data. See the Volume Manager web page at www.veritas.com for more information.

New features with VxVM 3.0 and higher include:

- Striped-Mirror Volumes

This type of volume combines striping and mirroring, but the mirroring is done at column level (or smaller for stripe-mirror subdisks). In case of failure, this type of volume recovers faster and the tolerance for disk failure is greater.

- RAID-5 Snapshot

It is now possible to snapshot a RAID-5 volume.

- RAID-5 Subdisk Move

This is now done without the cost of data redundancy.



- **Online Relayout**

Volumes can now be changed to a different layout. This is done online and in place. Online relayout can be used to change the redundancy or performance characteristics of the storage. Data organization (RAID levels), the number of columns for RAID-5 and striped volumes, and stripe unit size can be changed.

- **Task Monitor**

The Volume Manager Task Monitor tracks the progress of system recovery by monitoring task creation, maintenance, and completion. The Task Monitor allows you to modify characteristics of tasks.

- **Disk Group Versioning**

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

Volume Manager Release	Disk Group Version	Supported Disk Group Versions
1.2	10	10
1.3	15	15
2.0	20	20
2.2	30	30
2.3	40	40
2.5	50	50
3.0	60	20-60



- Additional Array Co-existence Support for DMP

See “VERITAS Volume Manager Disk Array Configuration Updates” in the *VERITAS Volume Manager Hardware Application Note* for more information.

- Destroy Disk Group

The `vxchg` command now provides a `destroy` option that removes a disk group from the system and frees the disks in that disk group for use in other disk groups. Disk groups that are not needed anymore should be removed with the `vxchg destroy` command so that the disks can be used by other disk groups.

The `vxchg deport` command can still be used to make disks inaccessible. The Volume Manager will prevent disks in a deported disk group from being used in other disk groups.

- New Graphical User Interface

This release of the Volume Manager includes the new Java-based VERITAS Volume Manager Storage Administrator graphical user interface. The Storage Administrator replaces the Visual Administrator graphical user interface that was shipped with previous releases of the Volume Manager. Refer to the “Storage Administrator” section for information about Storage Administrator features and changes.



Storage Administrator

The VERITAS Volume Manager Storage Administrator is the graphical user interface for the Volume Manager. The Storage Administrator has the following features:

- **Ease of Use**

The Storage Administrator is a task-based user interface that provides access to tasks through menus or a task list. With the Storage Administrator, administrators can easily navigate and configure their systems. Administrators can use the Storage Administrator to browse through all of the objects on the system or view detailed information about a specific object.

- **Remote Administration**

With the Storage Administrator, administrators can perform Volume Manager administration remotely or locally. The Storage Administrator client runs on UNIX or Windows machines.

- **Java-Based Interface**

The Storage Administrator client is a pure Java-based interface. Administrators can run the Storage Administrator as a Java application or from a Web browser.

- **Scalability**

The Storage Administrator can handle systems containing a large number of disks. Administrators can view all of the objects on the system or focus on a specific object or set of objects.

- **Security**

The Storage Administrator can only be run by users with appropriate privileges. The administrator can restrict the use of the Storage Administrator to a specific set of users.



The following are minimum system recommendations for the Storage Administrator client:

Solaris:	SPARCstation 5 with 64M memory
Solaris x86:	100MHz Pentium (or similar) with 64M memory
HP-UX:	Hewlett-Packard D-class machine with 64M memory
Windows:	100MHz Pentium with 32M memory

The following changes have been made to the Storage Administrator since the Storage Administrator 1.0 release:

- The Storage Administrator supports the new Volume Manager features listed above.
- The Storage Administrator server has been renamed to `vmsa_server`.
- The Storage Administrator client has been renamed to `vmsa`.
- The Storage Administrator provides two new volume layouts:
 - Concatenated Pro – a layered concatenated volume that is mirrored.
 - Striped Pro – a layered striped volume that is mirrored.
- The Move Subdisk task allows you to specify a disk offset and select a move policy.
- The toolbar has changed. The new toolbar contains several new buttons.
- The Create menu has been removed from the menu bar. The create object tasks are now located under the Console > New menu.
- The Add Disk and New Disk Group tasks allow you to specify the Volume Manager disk name for a disk.
- When running as an application, the Storage Administrator now saves user preferences in `user's_home_directory/.vmsa/VMpreference.prf` on the machine where the client is running. If you want to keep your existing preferences, you must copy your old preferences file (`/var/opt/vmsa/user_name/preferences`) to the new location *before* you use this release of the Storage Administrator.



When running from a Web browser, the Storage Administrator still saves user preferences in `/var/opt/vmsa/user_name/preferences` on the machine where the server is running.

The following changes have been made to the Storage Administrator since the Storage Administrator 3.0 release:

- The Storage Administrator supports VERITAS QuickLog™, which improves file system performance. QuickLog is an optionally licensable product and is only available for the VERITAS File System™ (referred to as VxFS® or vxfs).

Note that VERITAS QuickLog was previously known as VERITAS Accelerator. The Storage Administrator and its documentation still refer to QuickLog as the VERITAS Accelerator.

Cluster Functionality (Optional)

This Volume Manager release includes an *optional* cluster feature that enables VxVM to be used in a cluster environment. For information about the cluster functionality in the Volume Manager, refer to the *VERITAS Volume Manager Administrator's Reference Guide*.

With cluster support enabled, this release of VxVM supports up to four nodes per cluster. However, support for more than two nodes is currently only available if VxVM is used with a Sun StorEdge A3000/A5000.

Note: The new features introduced in Volume Manager 3.0 are available in private disk groups, but are not yet supported for shared disk groups.

The following VxVM features are now supported in a cluster environment:

- 32- and 64-bit Support
VxVM now provides 32- and 64-bit support and works with the Solaris 7 32- and 64-bit kernel in a cluster environment.
- Per-CPU Locking
Per-CPU reader/writer spinlocks have been implemented to eliminate global updates of variables. This improves scalability.



- Hot-Relocation

Hot-relocation is the ability of a system to automatically react to I/O failures on redundant (mirrored or RAID-5) VxVM objects and restore redundancy and access to those objects. The Volume Manager detects I/O failures on VxVM objects and relocates the affected subdisks to disks designated as *spare disks* and/or free space within the disk group. The Volume Manager then reconstructs the VxVM objects that existed before the failure and makes them redundant and accessible again.

Note: Hot-relocation is only performed for redundant (mirrored or RAID-5) subdisks on a failed disk. Non-redundant subdisks on a failed disk are not relocated, but the system administrator is notified of their failure.

The hot-relocation feature is enabled by default and it is recommended that you leave it on. However, you can disable hot-relocation by preventing the `vxrelocd` daemon from starting up during system startup. Refer to the *VERITAS Volume Manager Installation Guide* for details.

For more information on hot-relocation and `vxrelocd`, refer to the *VERITAS Volume Manager Getting Started Guide*.



End of Support Statements

The following software is no longer supported by VERITAS:

- VERITAS Volume Manager Release 1.3.x
- VERITAS Visual Administrator Release 1.3.x
- VERITAS Volume Manager Release 2.0.x
- VERITAS Volume Manager Release 2.1.x
- VERITAS Volume Manager Release 2.2.x
- VERITAS Volume Manager Release 2.3.x
- VERITAS Volume Manager Release 2.4.x
- Solaris 2.3 operating system
- Solaris 2.4 operating system
- Solaris 2.5 operating system
- VERITAS Volume Manager no longer supports the Sun-4c product line:
 - SPARCstation 1
 - SPARCstation 1+
 - SPARCstation 2
 - SPARCstation IPC
 - SPARCstation IPX
 - SPARCstation SLC

The following software is no longer available with the VERITAS Volume Manager:

- Volume Manager Visual Administrator (VxVA)



Installing the Volume Manager

Volume Manager installation consists of three parts:

1. Installing the packages onto the system.
2. Configuring and setting up the Volume Manager.
3. Setting up the Storage Administrator.

Installing Volume Manager Packages

Refer to the *VERITAS Volume Manager Installation Guide* for complete instructions on how to install VxVM using the `pkgadd` command. To install Volume Manager 3.0.2, use the instructions for installing Release 3.0.

The VERITAS CD-ROM contains the following packages:

- `VRTSvxvm`—Volume Manager software (driver and utilities)
- `VRTSvmdev`—Developer kit
- `VRTSvmman`—Manual Pages
- `VRTSvmdoc`—Volume Manager documentation
- `VRTSvmsa`—Storage Administrator software

Note: VxVM is a licensed product; you must obtain a license key before you install VxVM. To obtain a license key, complete a License Key Request Form and fax it to VERITAS Customer Support (see “Getting Help” for contact information). Refer to the section on obtaining a license key in the *VERITAS Volume Manager Installation Guide* for additional information.

Setting Up the Volume Manager

Refer to the *VERITAS Volume Manager Installation Guide* for information on how to initialize the Volume Manager (using `vxinstall`).

After the initialization is complete, you may be asked to reboot your machine to allow the kernel configuration to be updated. If necessary, you must reboot before any VxVM features are usable on your system.



Setting Up the Storage Administrator

Refer to the *VERITAS Volume Manager Installation Guide* for information on how to set up and start the Storage Administrator server and client.

Note: The Storage Administrator Release 3.x server is not backward compatible with Storage Administrator Release 1.x clients, so you must upgrade any existing client(s) to Release 3.x.

Upgrading the Volume Manager

Refer to the *VERITAS Volume Manager Installation Guide* for instructions on how to upgrade to Volume Manager Release 3.x and/or compatible releases of Solaris.



Documentation

The following documents accompany this Volume Manager release:

- *VERITAS Volume Manager Release Notes - Binary* (this document)
- *VERITAS Volume Manager Hardware Application Note*
- *VERITAS Volume Manager Installation Guide*
- *VERITAS Volume Manager Getting Started Guide*
- *VERITAS Volume Manager Administrator's Reference Guide*
- *VERITAS Volume Manager Command Line Interface Administrator's Guide*
- *VERITAS Volume Manager Storage Administrator Administrator's Guide*
- Online manual pages

The Storage Administrator provides online help files. To access the online help files, select the appropriate item from the Help menu or click Help in a dialog box.

Displaying Documentation Online

This product includes online documentation in Adobe Portable Document Format (PDF) and PostScript formats. You can view the documents online in either of these formats.

To view PDF documents, you must use the Adobe Acrobat Reader. You can use Acrobat reader as a stand-alone application, or as a plug-in to your web browser. However, VERITAS Software assumes no responsibility for the correct installation or use of Acrobat Reader. For more information on the latest versions of Acrobat Reader, or for help with installation problems, visit the Adobe web site at:

<http://www.adobe.com>

To view PostScript documents, you can use the Solaris Image Tool (`imagetool`) or any PostScript previewer.



The VERITAS Volume Manager guides are provided on the CD-ROM under the `pkgs/VRTSvmdoc` directory. If you have installed the `VRTSvmdoc` package, the documents are available in the following locations:

- *VERITAS Volume Manager Installation Guide*
`/opt/VRTSvxvm/docs/vminstall.ps`
`/opt/VRTSvxvm/docs/vminstall.pdf`
- *VERITAS Volume Manager Getting Started Guide*
`/opt/VRTSvxvm/docs/gsg.ps`
`/opt/VRTSvxvm/docs/gsg.pdf`
- *VERITAS Volume Manager Administrator's Reference Guide*
`/opt/VRTSvxvm/docs/ref.ps`
`/opt/VRTSvxvm/docs/ref.pdf`
- *VERITAS Volume Manager Command Line Interface Administrator's Guide*
`/opt/VRTSvxvm/docs/cli.ps`
`/opt/VRTSvxvm/docs/cli.pdf`
- *VERITAS Volume Manager Storage Administrator Administrator's Guide*
`/opt/VRTSvxvm/docs/vmsaguide.ps`
`/opt/VRTSvxvm/docs/vmsaguide.pdf`

Unformatted manual pages related to the VERITAS Volume Manager are located in the `VRTSvmmman` directory on the CD-ROM. If you have installed the `VRTSvmmman` package, the manual pages can be found in the `/opt/VRTSvxvm/man` directory. The Storage Administrator manual pages are in the `/opt/VRTSvmsa/man` directory. If you add these directories to your `MANPATH` environment variable, you can view these man pages with the `man(1)` command.



Printing Documentation

To print the documentation, you must have access to a PostScript printer. If you are not sure how to do this, or whether or not you have this functionality, consult your system administrator.

You can print the documents in the following ways:

- Use the print options in your PostScript previewer to print one or more pages.
- Use the print options in your Acrobat Reader viewer to print one or more pages.
- Print entire chapters using the `lp` command and your PostScript printer.



Software Problems Fixed In VxVM Release 3.0.2

The following problems have been addressed in this release of the Volume Manager.

- [28204] Changed device output message of the VxVM utility, `vxapslice`, when using a VxVM initialized disks that is in the Sun AP database. For example, the previous message displayed:

```
# /etc/vx/bin/vxapslice disk02
AP Database may be placed on partition:
/dev/ap/rdisk/mc0t1d1s0
```

and now the `vxapslice` command output displays:

```
/dev/rdisk/c0t1d1s0
```

- [28944] In VM 3.0.1, the VM diagnostic utilities `vxdevwalk` and `vxautoconfig` do not work on Solaris 7 and fail with the error messages:

```
vxvm:vxdevwalk: unsupported O/S, 5.7, error Bad address
vxvm:vxdevwalk: O/S 5.7 can not determine mode
```

and:

```
vxvm:vxautoconfig: unsupported O/S, 5.7, error Bad address
vxvm:vxautoconfig: O/S 5.7 can not determine mode
```

The VxVM diagnostic utilities, `vxautoconfig` and `vxdevwalk` are fixed in this release for Solaris 7. These diagnostic utilities are part of the set of diagnostic utilities called by the VxVM diagnostic script `/usr/lib/vxvm/diag.d/vxdmpdebug`.

- [29080] Corrected various spelling errors in the `vxdiskadm` action and help menu's.
- [29499] A StorEdge A5000 would not always be recognized correctly. It could be detected as either an A5000 or as a generic, non-multipathed disk array. This has been fixed.
- [29622] Modified the VM DMP administrator utility, `vmdmpadm` and the DMP diagnostic utility so that all instances of host I/O controllers which are of the same type and which are configured identically are



displayed. To display all host I/O controllers using the VM `vxdmpadm` utility, use the `listctlr=all` option. See the manual page for `vxdmpadm(1M)`.

- [29683] Accelerated the VM disk encapsulation time within `vxinstall`.
- [30016] The VxVM `vxunroot` utility on certain Sun disk types sets the Sun host OpenBoot PROM (OBP) boot-device boot path to a value which does not translate to a bootable device path. This is because the Sun OBP boot-device boot path value for certain disk types (SSA, EIDE, etc.) is unique and different. Fixed the VxVM `vxunroot` utility to use the VxVM utility, `vxeeeprom`, to correctly determine and set the Sun OBP boot-device boot path for VxVM unrooted disks.
- [30309] During a VxVM upgrade with the root device VxVM encapsulated, upon system reboot, if the `/etc/rc2.d/S75savecore` start up script is uncommented, the following error message is displayed:
- ```
savecore: no dump device configured
savecore: open("swap"): No such file or directory
savecore: no dump device configured
savecore: no dump device configured
dumpadm: no swap devices could be configured as the
dump device
```
- The VxVM upgrade and VxVM startup scripts have been corrected to configure a dump device.
- [30389] Enhanced the `vxresize` utility by adding in mock run option; added more debug messages (debug mode option); added disk group output to usage messages; added error message for the case where a volume name exists for more than one imported diskgroup; added the verbose option; added return error codes from routines which call `vxassist` routines and filesystem routines; in general if `vxassist` related error - return 2, if filesystem related error - return 3, all other general errors - return 1, and return 0 on success; display `errno` from `open()` call of the volume argument.
- [30438] The VM `vxdmpadm listctlr all` command does not display the correct controller type for a Sun StorEdge A5X00 Array connected to a PCI interface host bus adapter card. The DMP device driver device/path search algorithms have been updated to handle this.



- [30456] Additional support was added to VxVM scripts and utilities to accept Solaris device node arguments with the *target\_id* designation in the device node name as optional.
- [30468] An error message within the `vxedit` utility incorrectly advised usage help naming `vxedit`, where the usage help message should have named `vxvol`. The `vxedit` error message has been corrected to give the proper advice on using `vxvol set len`.
- [31005] Fixed a kernel panic in the VxVM `vxio` driver routine, `vol_mv_atcpy_free()`, when resynchronizing a mirror volume and the nonstale plex gets read errors.
- [31075] Added Dynamic Multipathing support for the Hitachi 7700E Disk Array subsystem.
- [31271] Added support for 64-bit Dynamic Multipathing diagnostic macros for the VxVM diagnostic utility,  
`/usr/lib/vxvm/diag.d/vxdmpdebug`.
- [31501] Fixed a possible core dump in the VM/DMP diagnostic utility, `vxdmpdbprint`, in the case where a DMP path name is `NONAME`. Within the DMP utilities, a path name is `NONAME` if the path has been physically removed, and Solaris has reconfigured the I/O subsystem (that is, `drvconfig`, `disks`), and VM has rebuilt the volume and plex device nodes (that is, `vxdctl enable`). A connected path has the DMP name of the Solaris device node (for example, `c#t#d#s2`).
- [31601] Fixed a small system memory leak when using `vxassist` to create multi-gigabyte RAID-5 volumes.
- [31892] Cataloged and internationalized VM messages (I18N messages)
- [31951] Running the VM command `vxtask list` treats system-level errors (such as permission denied on the `vxtask` device, etc.) as if there were no tasks. The VM command, `vxtask`, has been changed to notice system-level errors and to print the appropriate error message.



---

## Software Problems Fixed Prior to VxVM Release 3.0.2

The following problems were addressed in prior releases of the Volume Manager:

### Miscellaneous Problems Fixed

- [13227] Updated the `vxassist(1M)` manual page.
- [13689] Updated the `vxmake(1M)` manual page.
- [14944] Enhanced the VxVM `vxprint` utility where with the `-e vol_pattern` option, the `vol_pattern` variable can be an array of items.
- [21319] `vxconfigd` would dump core and remain unstartable when a DRL was added to `rootvol`. This behavior has been corrected, though there is still no useful reason to add a DRL to `rootvol`.
- [23163] When using the low level VxVM utility `vxsd` to associate several subdisks to a plex, the `vxsd` utility could fail and coredump when associating the last subdisk. This has been fixed.
- [24086] A race condition in `vxconfigd` resulted in a disk group operation that occurred during certain phases of initialization of `rootdg`, causing `vxconfigd` to follow a null pointer and to core dump. The code was changed to recognize the situation where `rootdg` is not yet fully configured by `vxconfigd` and to avoid dereferencing the null.
- [24127] A manual page for the `vxdumpadm` utility was created.
- [24215] If the VxVM `/etc/vx/disks.exclude` or the `/etc/vx/cntrls.exclude` file was very large, with a large number of disk and disk controller entries, respectively (for example, hundreds of disks), the following error message could be displayed repeatedly during some `vxdiskadm` disk operations:  

```
Output line too long.
```

This has been resolved.
- [24253] VxVM hot-relocation has been enhanced to allow for the option of relocating failing subdisks only to VM disks designated as VxVM hot-spare disks. For administrators who want to VxVM hot relocate only to the VxVM hot-spare disk pool, use the following guidelines:



- Define and initialize some non-VM disk(s) as VxVM hot-spare disks
- In the VxVM system file `/etc/default/vxassist` add the line:

```
spare=only
```

NOTE: If there is no VxVM hot-spare disk storage available, or if there is insufficient free space from the VxVM hot-spare disk pool, and if the `(spare=only)` line in the `/etc/default/vxassist` file is present, then a VxVM hot-relocation process will fail (see `vxrelocd(1M)`).

For administrators who do not want the hot-relocation daemon to allocate space solely to the VxVM hot-spare disk pool, do not have the `(spare=only)` line in the VxVM system file `/etc/default/vxassist`.

There is a default method for selection of the destination VxVM disk for subdisk relocation. The default is to first relocate subdisks to the available VxVM hot-spare disk pool, then relocate to other VM disks (see `vxrelocd(1M)`). By default the VxVM system file `/etc/default/vxassist` will not have the `(spare=only)` line present.

- [24967] In cases where the VxVM low-level `vxmake` utility is used to create a DRL log and then attach the log to a mirror plex, this plex is attached to a volume object thereby creating a mirror volume. Then assume that the DRL log would not start. A fix to the VxVM utility `vxmake` fixed this. Using the more standard and user-friendly VxVM utility `vxassist` to create and attach DRL logs is the preferred method.
- [25686] Volume Manager, prior to the VxVM 3.0 release, could potentially hang if very large I/O requests (several megabytes) were sent to the VxVM driver by the file system. In some cases, a combination of tuning parameters given to VxVM and a file system such as VxFS can result in the file system passing to VxVM an I/O that requires VxVM to allocate more memory than it has access to. This situation was not detected due to a bug in VxVM code, which resulted in a system hang. See the *VERITAS Volume Manager Administrator's Reference Guide* for VxVM system tunables information.



---

This problem was fixed by changing the VxVM driver to detect very large buffers and break them up, and by changing the driver's memory management code to detect and prevent over-consumption of kernel memory.

- [25711] A RAID-5 volume with a log plex that alphabetically falls before its data plex (for instance, the log plex is named pl01 and the data plex is named pl02) could not properly start if the `force (-f)` option to `vxvol` is used. This has been corrected.
- [26195] `vxmake volume` uses the plex length as the default, it should use `contig_len` (especially for striped or raid5). Fixed this problem by setting `contig_len` as the default.
- [26301] Changed the VxVM utility `vxassist remove log alloc=<>` to make the policies of removing mirror logs consistent with the removing of mirror plexes of a mirror volume.
- [26612] Fixed the VxVM utility `vxencap`. When encapsulating a disk or a disk partition with a file system, and using the low-level VxVM `vxencap` utility, you must specify a *medianame* (see `vxencap(1M)`). If you do not specify the medianame, the result is a VxVM volume name that is partially nonprintable. Using the more standard and user-friendly VxVM utilities, `vxinstall` and `vxdiskadm` is the preferred method for encapsulating disks and disk partitions with file systems.
- [26650] `vxrecover` would not always completely start striped-mirror volumes.
- [26761] Volume Manager volumes on `non-rootdg` disk groups, where all the disks in that `non-rootdg` disk group become unavailable, and where the volumes of that `non-rootdg` disk group have a heavy I/O activity, might cause `vxconfigd` to core dump. This has been fixed.
- [26994/29091]  
RAID-5 volumes with logs would not be recovered quickly in log overflow cases.
- [27665] When using `vxassist` to create a volume comprised of the maximum available free disk space from a large number of disks, `vxassist` may core dump. This has been fixed.





- 
- [27745] The remove option to `vxassist` required specification of disk group for volumes outside of `rootdg`. This is contrary to most other `vxassist` options.
  - [27781] `vxsd` would allow the disassociation of a subdisk from a RAID-5 volume, thereby making the volume degraded without issuing a warning.
  - [27850] `fixunroot` scripts would not appropriately update state information.
  - [28234] Enhanced the VM kernel drivers to add in SIO queuing sanity checks.
  - [28506] During a hot-relocation, DRL logs were unconditionally removed, even if the subsequent allocation of the new log failed.
  - [28534] AP metadevices could not be encapsulated.
  - [28736] Root encapsulation would leave `/etc/vfstab` world writable.
  - [28860] An attempt to encapsulate any disk other than the root disk would erroneously attempt to use the root disk as the target of encapsulation. The VTOC of the target disk would be removed without the appropriate volumes being created, and incorrect volumes would be created on the root disk. The `vfstab` would also be incorrectly updated. The system would be potentially unbootable. This has been fixed.
  - [28941] Encapsulate all menu in `vxinstall` would not assign default disk names correctly.
  - [29046] Made packaging changes to support Jumpstart.
  - [29148] A Sun StorEdge array could be defined multiple times, once as array type Sun StorEdge array, and once as a SEAGATE array.
  - [29259] Errors during `vxinstall` would cause script to hang waiting for input, if `vxinstall` ran on other than system console.
  - [29863] When using `vxassist` to `relayout` a volume to a new volume that required the addition of new columns, allocation attributes such as specific disks to use would be ignored. This has been fixed.



---

## Cluster-Related Problems Fixed

The following cluster-related problems have been fixed since a previous release of the Volume Manager with cluster support:

- [25792] Shortly after a node abort was performed on the master, the master could panic.
- [25812] `vxedit` failed if a slave was joined. [SUNW #4108657]
- [25974] A layered open in the `ssd` driver caused panics in intensive stress test cases. This problem happened when the Sun StorEdge A5000 went offline and all the targets were offlined. This problem was observed during stress tests, but is not likely to occur under normal use. Sun has issued a patch for this problem. Contact your Sun sales representative for more information.
- [26125] Disks that are offline when a disk group is deported are marked as outcast to prevent them from being autoimported. If `vxconfigd` was stopped and restarted while a disk was marked as an outcast, a `vxconfigd` core dump could occur.
- [26134] If a master deported a disk group and then attempted to reuse a deported disk, the slave would continue to associate the disk with the old disk group and prevent a join. During an import, the slave now re-online disks before reading the headers.
- [26267] A node could go into an infinite loop.
- [26365] Memory leaks were observed when disk groups were deported in a cluster.
- [26377] The interlock code had problems responding to reconfigurations.
- [26402; 27659] The `vxstat` command did not work on 64-bit kernels.
- [27053] In situations where there were many cluster reconfigurations and heavy I/O, a `vxconfigd` core dump could occur.
- [27210] Aborting a node while it was attempting to take over as master could cause a panic if that node was performing I/O when the former master node aborted or crashed.
- [27407] A problem that caused a fail-fast has been fixed.



- 
- [27491] Importing a shared disk group on a new master could cause the master to lock up.
  - [27580] An attempt to make a shared disk group with a name that was not unique on all nodes failed after the disk group name and ID were written to disk. The disks that were being used to make the shared disk group had to be manually initialized before they could be used in another disk group.
  - [27611] A problem that caused a panic has been fixed.
  - [28155] `vx dg destroy` did not clear the autoimport flag.
  - [28211] The `vxconfigd` daemon did not clean up the shared object database after a node abort or stop.
  - [28234] Driver queue sanity checks have been added to the code.  
[SUNW #4202995]
  - [none] The disk hot-sparing mechanism did not work well for *partial* disk failures. The hot-relocation feature replaces hot-sparing and handles partial disk failures.



---

## Software Problems Fixed in Storage Administrator Release 3.0.3

The following problems have been addressed in this release of the Volume Manager Storage Administrator:

- [30982] After selecting an object via a Browse button, the Apply button in a dialog box was disabled in some situations. This problem has been fixed.

The following problems were addressed in Storage Administrator Release 3.0.2:

- [28463] The Task Properties window now uses scrollbars in text areas when needed.
- [28469] A wait cursor is now shown during updates to the Volume Layout Details window.
- [29305] If the mouse was clicked too fast, duplicate entries could appear in the main window tree.
- [29573] A problem where Java exceptions occurred when subdisks were collapsed in the Volume to Disk Mapping window was fixed.
- [29733] A space allocation problem occurred when a controller was specified for volume creation.
- [29917] An error where disconnected disks could not be replaced has been corrected.
- [29918] An error where disconnected disks were not properly displayed has been corrected.
- [27994] Removal of the last plex in a volume is no longer allowed.

The following problems were addressed in Storage Administrator Release 3.0.1:

- [27972] Previously, the amount of space displayed in the Free Space field of the Disk Group Properties window included space on disks marked as “reserved” disks. The free space no longer includes space set aside on reserved disks.
- [28006] The Storage Administrator now shows an alert for failing disks.



---

The following problems were addressed in Storage Administrator Release 3.0:

- [17410] The Storage Administrator did not provide a way to mount unmounted file systems that were not on volumes.
- [19316] The undocked toolbar did not always redraw correctly the first time it was undocked.
- [21066] The splitter cursor could be displayed incorrectly on Windows.
- [21082] Modal dialogs were not modal when running under the Motif Window Manager (mwm).
- [22489] Changing the disk group name in the Create Volume dialog box did not update the default volume name.
- [22601] Volume and file system resize calculations did not work properly in the Windows version of the Storage Administrator. The Add By and Subtract By fields of the resize dialog boxes are not supported in the Windows version of the Storage Administrator. To resize a volume or file system, you must specify the desired size.
- [22755] Java memory leaks could occur when the Storage Administrator client was run for an extended period of time.
- [23338] Double-clicking on objects in the Volume Layout Details window did not work when running the Storage Administrator with the Motif Window Manager (mwm).
- [23805] The % Complete column in the Task Request Monitor window did not get updated.
- [24012] A bug in Java 1.1.5 and 1.1.6 causes the main window to resize itself when an object is selected and the Selected menu renames itself. The Storage Administrator is now shipped with Java 1.1.7, which does not have this problem.
- [24103] For year 2000 compliance, the Storage Administrator now displays years as four digits. In previous releases, the Task Request Monitor and Task Properties windows displayed years as two digits. The dates are used for display purposes only.
- [24996] When the Storage Administrator client and server were in different time zones, the Storage Administrator's command log file logged the client's time instead of the server's time.



- 
- [25075] When a volume was removed, volumes with similar names could be removed from the `/etc/vfstab` file.
  - [25230] A Java error resulting from using the Mount File System dialog has been corrected.
  - [25302] A file system's mount point could not be changed after the file system was added to a volume.
  - [25304] When a striped layout was selected in the Create Volume dialog, clicking the Maxsize button could cause the number of columns and stripe unit size to change.
  - [none] The Storage Administrator did not support configuration of the Dynamic Multipathing (DMP) feature.



---

## Software Limitations and Problems in VxVM Release 3.0.2

The following problems and issues exist in this release of the Volume Manager:

### Installation Issues

[28202] If you are using Volume Manager on top of an AP metadvice and **do not** want to encapsulate any disks, be sure to add entries for all controller(s) corresponding to those disks in the `/etc/vx/cntrls.exclude` file before running `vxinstall`. This means that entries for the meta-controller and the physical controllers that correspond to those disks, should be added to the `/etc/vx/cntrls.exclude` file.

For example, assume that the disks on controllers `c2` and `c4` are represented by the meta-controller `mc2` by `A`, and are connected to an SSA. To exclude the SSA from being seen by the VxVM utilities, there must be entries for `mc2`, `c2`, and `c4` in the `/etc/vx/cntrls.exclude` file. If all of these are not added to the `/etc/vx/cntrls.exclude` file, some VxVM utilities might still display the disks on one of more of these controllers.

If you are using Volume Manager on top of an AP metadvice and you **do** want to encapsulate the disk using `vxinstall`, specify the two controller numbers that reside below the metadvice (for example, `c0` and `c1` for `mc1`) in the `/etc/vx/cntrls.exclude` file.

[32715] After installing VERITAS Volume Manager 3.0.2 on a fully-pathed Solaris 7 system, approximately 1800 lines (consisting of duplicated driver `forceload` statements) may be added to the `/etc/system` file, and the lines are interspersed with the contents of the `/etc/system` file from the previous installation. In most instances, the corruption occurs during installation of the `VRTSvxvm` package.

After removing previous versions of VERITAS Volume Manager, the `/etc/system` file may still contain performance settings that contain the following VERITAS keywords

`vxio`

`vxdump`

`vxspec`



## vxvm

Even if you comment out the lines containing these keywords, they can still cause the post-install script to corrupt the `/etc/system` file.

The workaround is to edit the `/etc/system` file prior to installing a new version of VERITAS Volume Manager and remove (instead of comment out) any lines containing these keywords.

- [none] Remove a VM disk from a disk group after evacuating any data on the disk. You can permanently remove this disk from Volume Manager control by removing the VM metadata partition from that disk. Use the VM low-level command, `vxdiskunsetup` as shown in this example:

```
/usr/lib/vxvm/bin/vxdiskunsetup c#t#d#s2
```

Note that this command permanently removes a disk from Volume Manager control and should be used with caution only by a system administrator who is trained and knowledgeable in the use of Volume Manager.

- [none] Note any of the disks you are planning to initialize that were previously under Volume Manager control. If so, and they were used on the same host system, and the proper deinstallation procedures were *not* followed, the disk groups they represent will be imported automatically during the installation process. An attempt to initialize or encapsulate disks during the installation that were previously under Volume Manager control will fail. After the installation, if you no longer desire to use those disk groups, the `vxdbg (1M)` command has a `destroy` option that will remove those disk groups. Alternately, you can use `vxdiskunsetup (1M)` to remove the disks from Volume Manager control. Be very careful when using these options, because they can result in data loss if used incorrectly.

## Upgrade Issues

- [none] If you have third-party multipathing software configured (such as Sun Enterprise Server Alternate Pathing), you must insert the following step in the upgrade sections of the *VERITAS Volume Manager Installation Guide* Release 3.0.1.





---

After completion of the upgrade `start_script` and before rebooting the system (for example, by using `/etc/shutdown`), you must deconfigure the third-party software multipathing upgrades before continuing with the VERITAS upgrade.

Once you have completed deconfiguring the Solaris and third-party multipathing software upgrades, return to the next step in the VERITAS upgrade procedure.

[none] If a swap volume specified in `/etc/vfstab` was mirrored at the time that `upgrade_start` was run, the `upgrade_finish` script starts a resynchronization of the volume. This can cause a message similar to the following to be printed when the command to reboot the system is issued:

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

This message can be ignored.

[none] For a system on which the root file system is contained on a mirrored volume, the `upgrade_start` script may choose a mirror on a disk other than the normal boot disk to perform the upgrade. If this occurs, the reboot after running `upgrade_finish` may initially fail, claiming that the mirror on the boot disk is stale:

```
vxvm:vxconfigd: Error: System boot disk does not have a valid rootvol plex.
```

```
Please boot from one of the following disks:
```

```
Disk: *diskname* Device: *device*
```

```
...
```

```
vxvm:vxconfigd: Error: System startup failed
```

```
The system is down.
```

The system should be booted from one of the disks named. If the `eeprom` option `use-nvramrc?` is set to `true`, the system can be booted by specifying `vx-diskname`. (See Chapter 1 in the *VERITAS Volume Manager Administrator's Reference Guide* for details on booting when boot plexes are stale.)



---

## Utility Issues

- [6154] It should be noted that there is no protection built into `vxassist` to prevent the user from shrinking the `swap` volume without first shrinking what the system sees as available swap space. If it is necessary to shrink the `swap` volume, this operation should be done in single user mode and the system should be rebooted immediately. Failing to take these precautions could result in unknown system behavior or lock-up.
- [11286] Using `vx dg free` with a non-existent `disk-media-name` does not print an appropriate error message. It simply prints a header.
- [13418] The `vx disksetup` utility allows the administrator to specify some region/partition configurations that are invalid. Specifically, overlapping private and public regions can be specified, but can cause failures or data corruption when the disk is actually used. The administrator should be sure to check that the partitioning of the disk does not cause overlapping public and private partitions when the default partitioning is overridden from the command line. This problem shall be addressed in a future release.
- [13488] 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. If a log and a mirror are to be added, it can be accomplished by adding the mirror and the log in two separate `vxassist` invocations:
- ```
vxassist mirror volume ...
vxassist addlog volume ...
```
- This problem will be addressed in a future release.
- [none] Due to the current implementation to handle the resize of layered volumes, it is recommended not to `grow` or `shrink` layered volumes (stripe-mirror, concat-mirror, etc.) while resynchronization is ongoing.



---

Internally, the Volume Manager converts the layout of layered volumes and updates the configuration database before it shrinks or grows their sizes. This causes any ongoing operation, such as the 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, the user has to use `relayout` 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.

Although Release 3.0.2 supports layered volumes, it is not recommended to create volumes with mixed layout types. For example, it is not recommended to add a mirror using the `vxassist mirror` command specifying `layout=mirror-stripe nmirror=1` to an existing volume with `stripe-mirror` layout.

[none] While doing `relayout` on a mirrored volume, `vxassist` keeps the volume as mirrored even if the layout attribute is specified as `stripe` or `nomirror`. For example,

```
vxassist make vol 1024 layout=mirror-stripe ncol=3
vxassist relayout vol layout=stripe ncol=2
```

The volume `vol` is converted to a 2-column volume, but it is still mirrored even if the layout attribute is specified as `stripe` and `nomirror`.

## Device Issues

[none] Disks with insufficient space (less than 1024 disk blocks) for the allocation of an on-disk database copy cannot be encapsulated. The database requires at least the same space as is allocated for other disks in the same disk group. This size defaults to 1024 blocks. A way to work around this is to 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 this database must be allocated from the beginning or the end of the disk, with the exception of the root disk. The root disk can be encapsulated by carving out space from the `swap` partition if there is no space at the beginning or at the end of the disk. This is done by creating a subdisk for the private partition in the space obtained from the `swap` partition.

**Workaround:** There is no workaround to the problem of not having space on a disk to store private VxVM information. VxVM requires at least a small region of private storage (1024 blocks) for proper disk identification.

[5316] The Volume Manager tracks disks using long unique identifiers that VxVM stores on each disk. VxVM expects each disk to have a different unique identifier, and does not effectively guard against the situation where two disks have the same unique identifier. Duplicate identifiers should only occur as a result of the administrator using `dd` or some other utility to perform physical copies of the contents of an entire disk.

**Workaround:** The only effective workaround is that the administrator should not do exact physical disk copying.

[8818] It should be possible to prevent any access of a disk by VxVM. For example, startup of VxVM could be severely impacted by a disk with errors that result in I/O operations that take a long time to fail. However, when VxVM starts up, it accesses every disk on the system, by reading its `VTOC` and possibly a few blocks from one partition. There is currently no mechanism to prevent this. A disk can be offlined persistently, but the offline state is only recognized *after* the probe of all disks.

## Hot-Relocation Issues

[14894] 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).



---

It is possible to prevent subdisks from being split during hot-relocation. To do this, change the following `vxassist` line in the file `/usr/lib/vxvm/bin/vxrelocd`:

```
vxassist -r -g $dg_name move $v_name !$dm_name \
spare=yes >$resultfile
```

to this:

```
vxassist -r -g $dg_name move $v_name !$dm_name \
spare=yes layout=nospan,contig >$resultfile
```

- [14895] When a disk failure occurs, the hot-relocation feature notifies the system administrator of the failure and any relocation attempts through electronic mail messages. These 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.

## DMP Issues

- [18387] If the system has the Sun AP driver already installed, then during installation of VxVM, the DMP functionality components are not installed because the Sun AP driver and DMP cannot co-exist.

If the system has Volume Manager with DMP installed and you want to install the Sun AP driver, the following steps are required:

---

**Note:** Be sure to do these steps first:

1. `umount` all file systems created on Volume Manager volumes.
  2. Stop the Volume Manager (use `vxctl stop`).
- 

1. Remove the `vxdmp` driver from the `/kernel/drv` directory:

```
rm /kernel/drv/vxdmp
```

2. Edit `/etc/system`, and remove the line:

```
forceload: drv/vxdmp
```

3. Remove the Volume Manager DMP files:



```
rm -rf /dev/vx/dmp /dev/vx/rdmp
mv /kernel/drv/sparcv9/vxdmp \
/kernel/drv/sparcv9/vxdmp.SunOS.'uname -r'
```

4. Symbolically link /dev/vx/dmp to /dev/dsk:

```
ln -s /dev/dsk /dev/vx/dmp
```

5. Symbolically link /dev/vx/rdmp to /dev/rdisk:

```
ln -s /dev/rdisk /dev/vx/rdmp
```

6. Shut down the system to disable the DMP functionality:

```
/usr/sbin/shutdown
```

7. Reboot the system and install the Sun AP driver.

[none] Messages from the `Format` command can be ignored while you are running DMP.

## Cluster Functionality Issues

The following cluster-related issues exist for this release of the Volume Manager:

[none] The new features in Volume Manager Release 3.0 and higher are not supported in shared disk groups.

[none] If VxVM with cluster support is used with a Sun StorEdge A3000 or a Sun StorEdge A5000, up to four nodes are supported per cluster. With Solaris 7 and a Sun StorEdge A3000, only two nodes are currently supported per cluster. With a SPARCStorage Array or multipack, only two nodes are currently supported per cluster.

[none] Four-node support on Solaris 7 has not been fully tested yet.

[none] The cluster functionality in this release of VxVM has not been qualified on Sun SPARC PCI machines.



- 
- [none] It is possible to have private (non-shared) disk groups on physically shared disks. If these disks are on controllers that have been designated for fencing (i.e., reserved by Sun Cluster), the owner of the private disk group may not be able to access it when it is not in the cluster. For this reason, creating private disk groups on shared disks is not recommended unless the system administrator is fully aware of the consequences.
- [none] The Volume Manager does not currently support RAID-5 volumes in cluster-shareable disk groups. Creating and using RAID-5 volumes on shared disks may cause a system panic.
- [none] The use of file systems on volumes in cluster-shareable disk groups can cause system deadlocks. In particular, file systems on any type of volume in a shared disk group may lead to deadlocks during cluster reconfiguration events. `fsgen` volumes are not supported in shared disk groups; only `gen` volume types are supported.
- [none] When a node leaves the cluster due to clean shutdown or abort, the surviving nodes perform a cluster reconfiguration. If the leaving node attempts to rejoin before the cluster reconfiguration is complete, the outcome depends on whether the leaving node is a slave or master.
- If the leaving node is a slave, the attempt will fail with the error messages:
- ```
Resource temporarily unavailable
```
- and one of the following:
- ```
[vxclust] return from cluster_establish is configuration
daemon error -1
master has disconnected
```
- A retry at a later time should succeed.
- If the leaving node is a master, the attempt will generate disk-related error messages on both nodes and the remaining node will abort. The joining node will eventually join and may become master.
- [none] If `vxconfigd` is stopped on both the master and slave nodes and then restarted on the slaves first, Volume Manager output and GUI displays will not be reliable until `vxconfigd` has started on the



master and the slave has reconnected (which may take about 30 seconds). In particular, shared disk groups will be marked “disabled” and no information about them will be available. `vxconfigd` should therefore be started on the master first.

[none] When a node aborts from the cluster, open volume devices in shared disk groups on which I/O is not active are not removed until the volumes are closed. If this node later joins the cluster as the master while these volumes are still open, the presence of these volumes does not cause a problem. However, if the node tries to rejoin the cluster as a slave, this may fail with the error message:

```
cannot assign minor #
```

This is accompanied by the console message:

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

[none] To use the Volume Manager cluster functionality with a SPARCstorage Array, you must use firmware level 3.4 or higher.

[none] Dynamic Multipathing (DMP) enables the Volume Manager to use multiple host-to-disk paths in some multiported disk arrays. DMP enhances reliability by doing path failover in the event of the loss of one or more paths, and increases performance by doing I/O load balancing across multiple I/O paths.

VxVM does not currently support the Dynamic Multipathing feature in a shared write access disk environment because it is not possible to configure the disk arrays in a VxVM cluster to have multiple paths from a single host. It is therefore recommended that you disable DMP when VxVM is used in a cluster environment.

Disable DMP as follows:

1. Unmount all file systems created on Volume Manager volumes.
2. Stop the Volume Manager:

```
vxctl stop
```

3. Remove the `vxdmp` driver from the `/kernel/drv` directory:

```
rm /kernel/drv/vxdmp
```

```
mv /kernel/drv/sparcv9/vxdmp \
```





```
/kernel/drv/sparcv9/vxdmp.SunOS.'uname -r'
```

4. Edit `/etc/system`, and remove the following line:

```
forceload: drv/vxdmp
```

5. Remove the Volume Manager DMP files:

```
rm -rf /dev/vx/dmp /dev/vx/rdmp
```

6. Symbolically link `/dev/vx/dmp` to `/dev/dsk`:

```
ln -s /dev/dsk /dev/vx/dmp
```

7. Symbolically link `/dev/vx/rdmp` to `/dev/rdisk`:

```
ln -s /dev/rdisk /dev/vx/rdmp
```

8. Shut down the system to disable the DMP functionality:

```
/usr/sbin/shutdown
```

9. Reboot the system.

- [none] In the Sun Cluster, when a disk error occurs on a node, the disk is detached. This is the case even if the other node can access the disks successfully. As a result, an error in the path from a given node to a controller will result in the loss of all access to the disks on that controller.
- [20448] If a node leaves the cluster while a plex is being attached to a volume, the volume may remain in the SYNC state indefinitely. To avoid this, resynchronize the volume manually (with the command `vxvol -f resync volume`) after the plex attach completes. [Sun #4087612]
- [26175] Retries for failed I/Os result in hangs. This problem is not reproducible and may not exist now.
- [27090] In situations where there are many cluster reconfigurations and heavy I/O, a `vxconfigd` core dump may occur.
- [27177] A memory leak might occur in situations where there are many cluster reconfigurations and heavy I/O. This problem has only been seen in a debugging driver and when running specific invasive tests.



---

## Miscellaneous Issues

- [none] The Sun Online:Backup™ facility does not accept the long device path names for volumes. This is a limitation of Online: Backup which does not accept device paths longer than 24 characters. The simplest work-around for this problem is to use symbolic links to the longer `/dev/vx/dsk/volname` paths from a shorter pathname.
- [9936] RAID-5 volumes cannot currently be mirrored.
- [none] On machines with low memory (32 megabytes or less), under heavy I/O stress conditions against high memory usage volumes (i.e., RAID-5 volumes), we have encountered a situation where the system cannot allocate physical memory pages any more. For example, such a situation may result when exercising heavy I/O stress against RAID-5 volumes for 24 hours on a 32-megabyte machine.
- [13741] 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. It should be imported temporarily when the intention is to deport a diskgroup (for example, in HA configurations). Use the `-t` flag to `vxchg import`.
- [14450] During very fast boots on a system with many volumes, `vxconfigd` may not be able to autoimport 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:** A suggested workaround is to check the volumes before starting the application or place a sleep before the last `vxrecover`.
- [14909] If a disk fails after a snapshot is complete, the snapshot plex does not detect the failure or detach from the volume. This is because a snapshot plex is a write-only plex, so it notices I/O errors, but does not detach.



- [14915] 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, you can force the volume to start with the command:

```
vxvol -f start volname
```

However, you should 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.

- [25644] The way that Volume Manager handles minor numbers for volume devices can cause problems when upgrading from Volume Manager 2.x to 3.x. The problems are unneeded disk group remappings and the `vxconfigd` command can hang. The problem does not effect disk group versioning.
- [32576] Using an A5x00 array on a PCI bus as an encapsulated root (boot) disk or an alternate root disk is not supported at this time. If you encapsulate or mirror to an external root disk on an A5x00 array with a PCI bus, it may not be possible to boot from that disk.  
[Sun #4259045]



---

## Solaris Issues

Solaris issues involve these systems:

- Solaris 2.5.1, Solaris 2.6, and Solaris 7 (32- and 64-bit)

[6211] Driver close calls should take priority on SVR4 derivative operating systems. If this is not done, it is possible to receive a simultaneous last-close and first-open operation, swap the ordering of the arrival of operations to VxVM, and as a result, leave the volume device closed. (This can happen if VxVM detects an additional open of the device followed a last close).

[none] Since the disk label is stored in block 0 of the disk, block 0 must not be used (i.e., no application should write any information in block 0). Special protection has been built into VxVM to protect block 0 from being overwritten.

[6914] Boot disks are to be replaced with disks of similar geometry. When replacing a boot disk using the “Remove a disk for replacement” or “Replace a failed or removed disk” menus from the `vxdiskadm` utility, ensure that the replacement device has the same disk geometry as the failed boot disk. For disks other than boot disks, this restriction does not apply. A better solution for boot disks is to move all volumes from the boot disk to an alternate disk and to then remove the old disk.

[none] The UNIX `dd` command uses only `lseek()` to seek to a particular offset in a file. It does not use `llseek()`. This causes `dd` to fail on volumes greater than 2 gigabytes.

[none] On Solaris, slice 2 of a disk is the full disk by default. When finding connected disks, VxVM checks slice 2 of a disk. Slice 2 on a disk must always be defined as the full disk slice with a tag of 0x05.

[none] If the PROM version is not at least version 2, the system is not suitable for the configuration of a bootable root volume; the following message appears on the console each time the machine is booted after adding the VxVM package:

```
VxVM: Root volumes are not supported on your PROM
version.
```

Any attempt to encapsulate the root disk fails on these machines.



[8948] A Solaris SCSI disk driver offlines a disk when it is not available and notifies the console. When the disk is connected back to the system, it does not automatically make the disk open. All the processes that had the disk open cannot read or write from or to the disk until the disk is opened by some other utility causing the disk to come online. You can repair this situation by issuing a command that opens a partition of the disk. For example:

```
: < /dev/rdisk/clt5d0s2
```

[none] If you have multiple swap partitions on your disks and they are encapsulated, the Volume Manager names them as `swapvol`, `swapvol1`, `swapvol2`, etc. While rebooting the system, you may see the following error message:

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

This is due to a problem with the `swapadd` scripts, which use a faulty `swap -l` output when device names are longer than a particular number of characters. The swap devices, however, are correctly added and there are no bad effects on the system. To avoid seeing this spurious message, rename the swap volumes other than `swapvol` to `swap1`, `swap2`, etc. instead of their current `swapvol[0-9]` names.

[13312] The versions of the kernel drivers for VxVM are incompatible with some versions of the Solaris operating system. Multiple kernel modules are installed and are properly maintained by the installation and upgrade software. It is possible, however, that a mismatch can occur (such as if the administrator moves the kernel driver files). If such a mismatch occurs, the VxVM kernel prints a warning message on the console similar to the following:

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

If this occurs, the system should be booted for recovery (as explained in Chapter 1 of the *VERITAS Volume Manager Administrator's Reference Guide*) and the correct kernel modules installed. To install the correct kernel module versions, `cd` to the `kernel/drv` directory of the mounted root file system. The VxVM kernel modules can be listed with the following command:



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

The release-specific versions of the kernel modules are stored as *module.OS\_release*, where *OS* and *release* are the result of running `uname -s` and `uname -r` on the system, respectively. For example, on a misconfigured system running Solaris 2.5.1, the listing might look like the following:

```
-rw-r--r-- 1 root other 417660 Apr 29 07:17 vxdmp
-rw-r--r-- 1 root sys 406316 Apr 28 15:46 vxdmp.SunOS_5.5.1
-rw-r--r-- 1 root sys 417660 Apr 28 16:23 vxdmp.SunOS_5.6
-rw-r--r-- 1 root sys 420064 Apr 28 16:26 vxdmp.SunOS_5.7
-rw-r--r-- 1 root sys 986 Apr 28 15:46 vxdmp.conf
-rw-r--r-- 1 root other 2108780 Apr 29 07:17 vxio
-rw-r--r-- 1 root sys 2088276 Apr 28 15:45 vxio.SunOS_5.5.1
-rw-r--r-- 1 root sys 2108780 Apr 28 16:23 vxio.SunOS_5.6
-rw-r--r-- 1 root sys 2118816 Apr 28 16:26 vxio.SunOS_5.7
-rw-r--r-- 1 root sys 958 Apr 28 15:45 vxio.conf
-rw-r--r-- 1 root other 15584 Apr 29 07:17 vxspec
-rw-r--r-- 1 root sys 15068 Apr 28 15:45 vxspecSunOS_5.5.1
-rw-r--r-- 1 root sys 15584 Apr 28 16:23 vxspec.SunOS_5.6
-rw-r--r-- 1 root sys 15976 Apr 28 16:26 vxspec.SunOS_5.7
-rw-r--r-- 1 root sys 1238 Apr 28 15:45 vxspec.conf
```

Note that the size of the kernel modules being used (those without suffixes) match the *driver.SunOS\_5.4* versions. To correct the problem, copy the *SunOS\_5.5* versions to the "in-use" module names:

```
cp vxio.SunOS_5.5 vxio
cp vxspec.SunOS_5.5 vxspec
```

The root file system should then be unmounted and the system can be rebooted.

[13388] During encapsulation, the Volume Manager does not consider a partition as a swap partition unless its partition tag (as shown by `prtvtoc`) is `swap` or `3`. Any partition used as a swap partition but not tagged as such is encapsulated as a file system. In the `vfstab`, a



note is made that the partition has been encapsulated, but the `vfstab` entry *is not* translated, and thus, the partition is not added as a swap area as part of the boot process.

All partitions used as swap must be marked with the `swap` tag if they are to be properly encapsulated.

- [24619] To install and test the Solaris-based Volume Manager 3.0.2 package, you **MUST FIRST** install the Sun patches according to the patch matrix shown in Table 1.

**Table 1** Solaris Operating System Patch Requirements

| Sun Operating System | Sun Array    | Sun Patch        |
|----------------------|--------------|------------------|
| Solaris 7            | N/A          | 106541           |
| Solaris 2.6          | SSA<br>A5X00 | 105223<br>105357 |
| Solaris 2.5.1        | SSA<br>A5X00 | 104708<br>105324 |

The Sun patches are available through the Sun web site [sunsolve.sun.com](http://sunsolve.sun.com).

Volume Manager 3.x `pkgadd` scripts for the `VRTSVXVM` have been modified so that it produces an informational message if the appropriate required Sun patches are not present on your system.



---

## Software Limitations and Problems in Storage Administrator Release 3.0.3

The following problems and issues exist in this release of the Volume Manager Storage Administrator:

[17772] The Volume Manager Storage Administrator does not support destroying deported disk groups.

**Workaround:** Import the disk group, then destroy it.

[18338] The Volume Manager Storage Administrator does not allow the user to specify volume usage types.

[21512] The Volume Manager Storage Administrator applet does not start up in HotJava when security is enabled.

**Workaround:** Set the CLASSPATH environment variable to:  
CLASSPATH=/opt/VRTSvmsa/vxvm/java

[21739] With Solaris 2.5.x, the Storage Administrator server (`vmsa_server`) needs to be restarted whenever `vxconfigd` is reset, stopped and restarted, or disabled and enabled.

[22456] The Maxsize operation always returns sizes in sectors.

[22482] The Volume Manager Storage Administrator server can hang on Solaris while waiting for `ncsd(1M)` to reply to a name service lookup. This can occur when the system is set up for DNS and a DNS server is not set up.

**Workaround:** Remove the `dns` entry from the `hosts: files dns` line in the `/etc/nsswitch.conf` file.

[22730] Online help is not supported for non-English locales. If the user's `$LANG` environment variable is set to a non-English value, the online help files can have problems printing and following their hyperlinks.

**Workaround:** Create a symbolic link:  
`cd /opt/VRTSvmsa/vxvm/java`  
`ln -s help help_locale`

where *locale* is the appropriate locale abbreviation.





- [22752] On Windows, attempts to print online help are silently ignored if there is no default printer on the system.
- [23730] The splitter cursor does not always go away. This can prevent the wait cursor from being displayed. This problem is caused by a Java bug.
- Workaround:** Move the cursor outside the main window and then back into the main window.
- [24701] When bringing up a dialog box such as a disk/space allocation dialog box, a message similar to the following may appear:
- Warning:
- ```
Name: scrollbar
Class: XmScrollBar
The scrollbar page increment is less than 1.
```
- This message can be ignored.
- [25080] The Storage Administrator does not run with HotJava 1.1.4.
- [25089] If you remove a Volume Manager Storage Administrator 1.x package and then install a Storage Administrator 3.x package, the new Storage Administrator server may not start properly.
- Workaround:** Before you remove the Storage Administrator 1.x package, run `server.sh -k` to stop the 1.x server.
- [25361] When the window manager's interactive placement setting is turned on, Storage Administrator windows may not be sized properly. The windows may be very small or very large.
- Workaround:** Turn off the interactive placement setting.
- Example:
- ```
Mwm*interactivePlacement: False
```
- [26304] Dragging the toolbar to a new location or moving the mouse button over the toolbar may result in an exception. This exception can be ignored.
- [26269] Menu shortcuts only work when input focus is in the tree or grid. To enable shortcuts, select an item in the tree or grid.



- 
- [27291] Under some circumstances, moving the mouse over a toolbar button may result in an exception. If this happens, the status area may not display the correct toolbar button description.
- [27348] When the Storage Administrator is run with the `f_vwm` window manager, window decorations appear above top of screen.
- [27349] When the Volume to Disk Mapping window is updated, the contents of the window may become inaccurate.
- Workaround:** Close the Volume to Disk Mapping window and then reopen it.
- [29621] When adding disks, the order in which the Volume Manager disk names are assigned may not match the order of the selected or specified devices. The disks are processed in the order in which they appear in the grid.
- [31029] The `VRTSvmsa` client-only installation output states that `VRTSvxvm` is a prerequisite. The `VRTSvxvm` package must be installed on the machine where you install the server portion of the `VRTSvmsa` package, but `VRTSvxvm` is *not* a prerequisite for the Storage Administrator client.
- [32052] The Storage Administrator does not support communication between the client and server across a firewall.
- [32599] The Storage Administrator may have trouble connecting to a host machine (server) if multiple host names are associated with a single IP address. The Storage Administrator displays the following message:
- Summary:  
There is no such server (*host1*)
- Detail:  
java.net.UnknownHostException: Unknown host:  
[*host2*:32839]; nested exception is:  
java.net.UnknownHostException: *host2*
- In this example, the administrator specified *host1* for the VMSA connection, but the server host machine was identified as *host2*.



---

In some cases, this may be a problem with the way DNS is set up. A DNS reverse lookup (by IP address) may return a host name that differs from the host name provided to VMSA at startup (*host1*). You may have to make appropriate changes to DNS so that the names are consistent. In other cases, this happens because *host1* is not the first host in the list of hosts for the associated IP address in the `/etc/hosts` file.

**Workaround:** Make sure *host1* shows up as the first host in the list of names for the address of *host1* in `/etc/hosts`.

- [none] To prevent core dumps on NCD terminals, the `jre/lib/font.properties` file has been renamed to `font.properties-`. On some machines, this has a negative effect on the appearance of default fonts for online help and other text areas. If you *do not* use NCD terminals, you can improve the appearance of these fonts by renaming `font.properties-` to `font.properties`.
- [none] The Volume Manager Storage Administrator does not support statistics and analysis.
- [none] The following X Window System error may occur when starting the Storage Administrator:
- ```
Xlib: connection to "hostname:0.0" refused by server
Xlib: Client is not authorized to connect to Server
java.lang.InternalError: Can't connect to X11 window
server using hostname:0.0'as the value of the DISPLAY
variable.
    at sun.awt.motif.MToolkit.<init>(MToolkit.java:48)
    at java.awt.Toolkit.getDefaultToolkit(Toolkit.java:244)
```
- Workaround:** Type `xhost + [hostname]` to allow X server access.
- [none] If a volume with an unmounted file system is resized, the file system may not be resized.
- [none] If you unpack and uncompress the Apache tar file included with this release, package remove scripts will not remove the resulting files. However, package remove scripts will remove the Apache tar file.



-
- [none] To run the Storage Administrator from Netscape, the UniversalConnect privilege must be granted. Once UniversalConnect is granted, be very careful when browsing the Internet. You will always be prompted by Netscape Communicator when an unknown applet attempts to enable privileges. Any codebase principal you run across should be considered suspect, as there is no way to verify the identity of its creator or even the integrity of the bits as they travel from the host machine to your own.
- [none] Although you can run the Storage Administrator from a Web browser, it is recommended that you run it as an application instead. Due to the nature of Web browsers, running the Storage Administrator from a Web browser can slow down response time significantly.
- {none} If you encounter a JVM (Java Virtual Machine) system failure while running some versions of Solaris, the probable cause is that JVM is affecting the library used by VERITAS Volume Manager for making Java calls. VERITAS Volume Manager currently uses the green threads that ship with JVM. In some instances, using these threads has caused JVM to fail.

Workaround: Use the native threads supplied through Java. To change to native threads, edit the `vmsa` client script located in `/opt/VRTSvmsa/bin` and change the line containing the `TH_TYPE=` to: `TH_TYPE=-native`.

Note: Solaris 2.5.1 cannot run under native threads. Do not switch to native threads if you encounter this problem on Solaris 2.5.1.



Encapsulating and Mirroring the Root Disk

If you plan to mirror the root disk (which contains the root file system) so that an alternate root disk exists for booting purposes, you should place the root disk under Volume Manager control through encapsulation. The root disk can be encapsulated either during the `vxinstall` process (while installing VxVM), from the `vxdiskadm` menus (after VxVM is installed), or from the Storage Administrator. Once encapsulated, the root disk can be mirrored using `vxdiskadm`.

Booting From DMP Devices

When the root disk is placed under Volume Manager control, it is automatically accessed as a Dynamic Multipathing (DMP) device with one path if it is a single disk, or with more paths if the disk is part of a multiported disk array. By encapsulating the root disk, the system reliability is enhanced against loss of one or more of the existing physical paths to a disk. For more information, refer to the *VERITAS Volume Manager Getting Started Guide*.

VxVM and Multi-Host Failover Configurations

VxVM disk groups can be “imported” (made available) from only one host at any given time. When a host imports a disk group, the volumes and configuration of that disk group becomes accessible to the host. If the administrator or system software wants to use the same disk group from another host, the host that already has the disk group imported (*importing host*) must “deport” (give up access to) the disk group. Once deported, the disk group can be imported by another host.

If two hosts are allowed to access a disk group concurrently, the configuration of the disk group, and possibly the contents of volumes, could be corrupted. Similar corruption would also occur if a file system or database on a raw disk partition were accessed concurrently by two hosts, so this is not a problem limited to VxVM.

When a host imports a disk group, an import lock is written on all disks in that disk group. The import lock is cleared when the host deports the disk group. The presence of the import lock prevents other hosts from importing the disk group until the importing host has deported the disk group. Specifically, when a host imports a disk group, the import normally fails if any disks within the



disk group appear to be locked by another host. This allows automatic reimporting of disk groups after a reboot (called *autoimporting*) and prevents imports by another host (even while the first host is shut down). If the importing host is shut down without deporting the disk group, the disk group can only be imported by another host by clearing the host ID lock first (discussed later).

Note: The import lock contains a host ID (with VxVM, this is the host name) reference to identify the importing host and enforce the lock. Problems can therefore arise if two hosts have the same host ID. Since VxVM uses the host name as the host ID (by default), it is advisable to change the host name of one machine if another machine shares its host name. This can be followed with the command `vxdctl hostid new_hostname`.

The import locking scheme works well in an environment where disk groups are not normally shifted from one system to another. However, consider a setup where two hosts, Node A and Node B, can access the drives of a disk group. The disk group is first imported by Node A, but the administrator wishes to access the disk group from Node B if Node A crashes. This kind of scenario (commonly called *failover*) can be used to provide manual high availability to data (where the failure of one node does not prevent access to data). Failover can be combined with a “high availability” monitor to provide automatic high availability to data: when Node B detects that Node A has crashed or shut down, Node B imports (fails over) the disk group to provide access to the volumes.

VxVM can support failover, but it relies on the administrator or on an external high availability monitor to ensure that the first system is really shut down or unavailable before the disk group is imported to another system. For details on how to clear locks and force an import, refer to the `vxdg(1M)` manual page and the section on moving disk groups between systems in Chapter 3 of the *VERITAS Volume Manager Administrator's Reference Guide*.



CAUTION! If `vx dg import` is used with `-C` (clears locks) and/or `-f` (forces import) to import a disk group that is still in use from another host, disk group configuration corruption is very likely to occur. Volume content corruption is also likely if a file system or database is started on the imported volumes before the other host crashes or shuts down.

If this kind of corruption occurs, it is likely that you have to rebuild your configuration from scratch and reload all volumes in the disk group from a backup. There are typically a large number of configuration copies for each disk group, but corruption nearly always affects all configuration copies, so redundancy does not help in this case.

Disk group configuration corruption usually shows up as missing or duplicate records in the configuration databases. This can result in a wide variety of `vxconfigd` error messages, including errors such as:

```
Association not resolved
Association count is incorrect
Duplicate record in configuration
Configuration records are inconsistent
```

These errors are typically reported in association with specific disk group configuration copies, but usually apply to all copies. The following is usually displayed along with the error:

```
Disk group has no valid configuration copies
```

See Appendix A of the *VERITAS Volume Manager Administrator's Reference Guide* for more information on VxVM error messages.

If you use the VERITAS FirstWatch® product, all disk group failover issues can be managed correctly. FirstWatch includes a high availability monitor and includes failover scripts for VxVM, VxFS®, and for several popular databases.

The `-t` option to `vx dg` prevents automatic reimports on reboot and is necessary when used with a host monitor (such as FirstWatch) that controls imports itself, rather than relying on automatic imports by VxVM.

