



VERITAS Volume Manager™

Installation Guide

Release 3.0.4

PLEASE READ THIS BEFORE INSTALLING THE SOFTWARE!

Solaris
April, 2000
P/N 100-001634

© 1999 - 2000 VERITAS® Software Corporation. All rights reserved.

TRADEMARKS

VERITAS is a registered trademark of VERITAS Software Corporation in the United States and other countries.

VERITAS Volume Manager, VERITAS File System, VERITAS NetBackup, 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 information on installing and initializing the VERITAS® Volume Manager and the Volume Manager Storage Administrator (VMSA) graphical user interface. This document also contains information on how to upgrade to the current release of the Volume Manager.

Contents

Topics covered in this document include:

- Getting Help
- Pre-Installation Instructions
 - Obtaining a License Key
 - SDS to VxVM Conversion
- Installing the Volume Manager
 - Installation Overview
- Loading the Software from CD-ROM
 - Disabling Hot-Relocation
- Initializing the Volume Manager
 - Quick Installation
 - Custom Installation
 - Rebooting After vxinstall
- Setting Up and Starting the Storage Administrator
 - Storage Administrator Installation Overview



-
- Installing the Storage Administrator Server and/or Client (UNIX)
 - Installing the Storage Administrator Client (Windows)
 - Customizing Security (Optional)
 - Setting Up the HTTP Server (Optional)
 - Starting the Storage Administrator Server
 - Starting the Storage Administrator Client
 - Running the Storage Administrator in Read-Only Mode
 - Enabling Cluster Support in Volume Manager (Optional)
 - Configuring Shared Disks
 - Disk Controller Reservation
 - Upgrading to VxVM Release 3.0.4
 - Upgrading to VxVM 3.0.4 and Solaris 2.5.1/2.6/7/8
 - Upgrading VxVM Only
 - Upgrading Solaris Only
 - Upgrading the Storage Administrator
 - Checking Volume Manager Processes
 - Deinstalling the Volume Manager
 - Shutting Down the Volume Manager
 - Removing the Volume Manager Package

Getting Help

If you have any comments or problems with the VERITAS products, contact the VERITAS Customer Support hotline:

US Customers: 1-800-342-0652

International Customers: +1 (650) 335-8555

Fax: (650) 335-8428

Electronic mail: support@veritas.com



Pre-Installation Instructions

Review the following items and make any appropriate changes *before* you install or upgrade the Volume Manager:

- Most of the commands used in the installation are in the `/sbin` or `/usr/sbin` directory. You should add these directories to your `PATH` environment variable.

Examples:

If you are using a Bourne Shell (`sh` or `ksh`), use the command:

```
$ PATH=/sbin:/usr/sbin:$PATH; export PATH
```

If you are using a C Shell (`csh` or `tcsh`), use the command:

```
% setenv PATH /sbin:/usr/sbin:${PATH}
```

- If any of the documentation, developer kit, or manual page packages are to be installed, the directory `/opt` should exist and be writable.
- In order to initialize the Volume Manager, you must obtain a license key from your customer support channel. When you contact your support representative for your key, be prepared to give them information about your machine. Refer to the section entitled “Obtaining a License Key” for details.
- Before initializing the Volume Manager, make sure you are aware of the contents of each disk. During the initialization process, you will have to decide the following:
 - Do you want to place all or just some of the disks on the system and/or controller under Volume Manager control?
 - Do you want to place the system’s root disk under Volume Manager control? (This is necessary in the case of a one-disk system.)
 - When you place disks under Volume Manager control, do you want to preserve data in any or all existing file systems and partitions through *encapsulation*?
 - For each disk, do you want to choose whether to encapsulate, initialize, or leave the disk alone (Custom Installation)? Alternatively, do you want to either encapsulate or initialize all disks on a controller together (Quick Installation)?



Note: The software initialization may require repeated reboots. You are usually asked for permission to reboot when the reboots are needed. If you choose not to reboot, the initialization will not complete. After the first reboot, the initialization process takes control of the system until the installation is complete.

- Any disk to be managed by the Volume Manager should have two free partitions and a small amount of free space (1024 sectors). The free space should be at the beginning or end of the disk and should not belong to a partition. This space is used for storing disk group configurations and a disk label that ensures that the Volume Manager can identify the disk, even if it is moved to a different address or controller. It is also used to keep track of disk configuration and to ensure correct recovery. The Volume Manager allocates approximately 1024 sectors (512K) from each disk for the disk group configurations and the disk label. This space is sufficient to handle normal disk group configurations for up to approximately 100 disks.

The boot (root) disk is a special case; if no other space is available, the Volume Manager will attempt to allocate space usually reserved for swap by shrinking the swap partition. This process is known as *swap relocation* and, if necessary, will happen automatically during root disk encapsulation.

Note: Although it is possible to put disks with no free space under Volume Manager control, this should only be used as a migration strategy because many VxVM capabilities based on disk identity are available only for disks with the required free space. Refer to the `vxdisk(1M)` manual page for information on the `nopriv` disk type.

- If you plan to encapsulate any disks, make sure that those disks:
 - have the required minimum free space (1024 sectors)
 - have an `s2` slice that represents the full disk – this `s2` slice cannot contain a file system
 - have two free partitions



- If you plan to encapsulate the boot disk, the swap partition should be tagged as `swap` so that it will be possible to dump to that partition later.

Before encapsulating your boot disk, the EEPROM variable `use-nvramrc?` must be set to `true` if you want to take advantage of the Volume Manager boot disk aliases to identify the mirror of the boot disk if a replacement is needed. If this variable is set to `false`, you will need to determine which disks are bootable yourself. You can set this variable to `true` as follows:

```
eeeprom use-nvramrc?=true
```

Note: If the path to an aliased boot device is different from the path in the `/device` directory, aliases may not function correctly.

- Before installing the software, confirm that there is sufficient space.

The following table shows the suggested *minimum* space allowances for each package.

Package	Contents	Size
VRTSvxvm	Driver and Utilities	19M
VRTSvmdev	Developer Kit	2M
VRTSvmman	Manual Pages	1M
VRTSvmdoc	Documentation	11M
VRTSvmsa	Storage Administrator Server & Client (UNIX)	23M

The following table shows the approximate breakdown of the space requirements for each package against each file system.

Package	/	/usr	/opt
VRTSvxvm	9M	10M	0
VRTSvmdev	0	0	2M
VRTSvmman	0	0	1M
VRTSvmdoc	0	0	11M
VRTSvmsa	0	0	23M



- The hot-relocation feature is enabled by default. The associated daemon, `vxrelocd`, is automatically started during system startup. If you want to disable hot-relocation for some reason, you can do so by preventing `vxrelocd` from running after you load the VxVM software. See the sections entitled “Disabling Hot-Relocation” and “Upgrading to VxVM 3.0.4 and Solaris 2.5.1/2.6/7/8” for details.
- The following software must be installed and running on the machine to be administered by the Volume Manager Storage Administrator:
 - Solaris 2.5.1 and higher.
 - VERITAS Volume Manager Release 3.0.4. (The `vxconfigd` daemon must be running on the machine to be administered.)
 - The Volume Manager Storage Administrator server.
 - An HTTP server. (This is only necessary if you plan to run the Volume Manager Storage Administrator from a Web browser.)
- If an earlier Volume Manager Storage Administrator package is installed on the machine, remove the old Storage Administrator package before you install the new Storage Administrator package. See “Upgrading the Storage Administrator” for details.
- If you plan to use the Volume Manager in a cluster environment, you need to do the following:
 - Obtain a license for the *optional* Volume Manager cluster feature from your Customer Support channel.
 - Before installing the Volume Manager, you should decide where to place the root disk group (`rootdg`) for each node in the cluster. A system using the Volume Manager has one or more disk groups, including `rootdg`. `rootdg` must exist and cannot be shared between systems. At least one disk must exist within `rootdg` while VxVM is running.
 - Before beginning the installation, you must decide on the layout of shared disk groups. There may be one or more shared disk groups. Determine how many you wish to use.
 - If you plan to use Dirty Region Logging with VxVM in a cluster, you should consider leaving a small amount of space on the disk for these logs. The log size is proportional to the volume size and the number of nodes (each log has one recovery map plus one active map per node).



For a two-gigabyte volume in a two-node cluster, a log size of three blocks (one block per map) should be sufficient. For every additional two gigabytes of volume size, the log size should then increase by approximately one block per map (so a four-gigabyte volume with two nodes would have a log size of six blocks) up to a maximum of 96 blocks. For larger volumes, DRL changes the log granularity to accommodate the increased size without exceeding the maximum log size. A four-node cluster requires larger logs.

See the section on Dirty Region Logging in the document entitled *Volume Manager Cluster Functionality* for more information about log sizes.

Obtaining a License Key

VxVM is a licensed product. When you purchase this product, it is necessary to have a license key issued by VERITAS Customer Support.

Use of this product requires a license key.

Note: You must obtain your license keys before installing VxVM 3.x. Regular license keys do not work. If you have a license from a previous version, you still need to obtain a license key.

To obtain a license key, complete a License Key Request Form and fax it to VERITAS Customer Support (see “Getting Help” for contact information).

To issue a new license key, VERITAS Customer Support requires the following information:

- the host ID
- the machine type

The host ID can be obtained by running the command:

```
hostid
```

The host machine type can be obtained by running the command:

```
uname -i
```

The host type is listed in the first line of output that follows a blank. For example, the host type of a SPARCserver 1000 is `SUNW,SPARCserver-1000`.



You will be prompted for the license key(s) when you run `vxinstall`.

Later upgrades of the licensing features may require the use of the command `vxlicense -c` to replace the old licenses with new ones.

You can view the currently installed licenses with this command:

```
vxlicense -p
```

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.



Installing the Volume Manager

This section describes how to install and set up the Volume Manager and the Storage Administrator. Volume Manager installation and setup consists of these parts:

1. Package installation.
2. Volume Manager configuration and setup.
3. Storage Administrator setup.
4. Shared disk setup (for cluster environments only).

Note: Only users with superuser privileges can install and initialize the Volume Manager and Storage Administrator.

For information on how to upgrade the Volume Manager and Solaris, refer to “Upgrading to VxVM Release 3.0.4.” For information on how to upgrade the Storage Administrator, see “Upgrading the Storage Administrator.”

Installation Overview

To install the Volume Manager and the Storage Administrator for the first time:

1. Install the software packages onto the system to be administered (as described in the section “Loading the Software from CD-ROM”).
2. Initialize VxVM using `vxinstall` (as described in the section “Initializing the Volume Manager”).
3. Set up and start the Storage Administrator (as described in “Setting Up and Starting the Storage Administrator”).
4. If you are planning to use the Volume Manager in a cluster environment, follow the instructions in “Enabling Cluster Support in Volume Manager (Optional).”



Loading the Software from CD-ROM

The VERITAS CD-ROM contains the following packages:

- `VRTSvxvm` – the Volume Manager software.
- `VRTSvmdoc` – online copies of the following VERITAS documents:
 - *VERITAS Volume Manager Getting Started Guide*
 - *VERITAS Volume Manager Command Line Interface Administrator's Guide*
 - *VERITAS Volume Manager Storage Administrator Administrator's Guide*
 - *VERITAS Volume Manager Administrator's Reference Guide*
 - *VERITAS Volume Manager Installation Guide*

If you do not want to be able to read or print these documents, you may omit the `VRTSvmdoc` package.

- `VRTSvmman` – the VxVM manual pages.
- `VRTSvmdev` – (optional) a VxVM developer's package (library and header files).
- `VRTSvmsa` – the Storage Administrator software.

Note: If an earlier Volume Manager Storage Administrator package is installed on the machine, remove the old Storage Administrator package before you install the new Storage Administrator package.

Load the software from the CD-ROM as follows:

1. Log in as superuser.
2. Insert the VERITAS CD-ROM into a CD-ROM drive connected to your system.
3. If the Solaris volume management software is running on your system, mount the CD-ROM as described in the section “Mounting the CD-ROM With Solaris Volume Management.” Otherwise, mount the CD-ROM as described in the section “Mounting the CD-ROM at the Command Line.”



▼ Mounting the CD-ROM With Solaris Volume Management

Once the CD-ROM is inserted, the Solaris volume management software automatically mounts the CD-ROM as `/cdrom/CD_name`.

▼ After the CD-ROM is mounted, enter the following to install the packages:

```
# pkgadd -d /cdrom/CD_name/OS_version/pkgs VRTSvxvm \
    VRTSvmsa VRTSvmdoc VRTSvmman VRTSvmdev
```

Note: List the `VRTSvxvm` package before the `VRTSvmsa` package because the Volume Manager must be installed before the Storage Administrator.

The system prints out a series of status messages as the installation progresses. No errors should be reported, and the process should take about five minutes.

▼ Mounting the CD-ROM at the Command Line

If the Solaris volume management software is not available to mount the CD-ROM automatically, enter the following to mount the CD-ROM manually, assuming your CD driver is on controller 0:

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

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

▼ After the CD-ROM is mounted, enter the following to install the packages:

```
# pkgadd -d /mnt/OS_version/pkgs VRTSvxvm VRTSvmsa \
    VRTSvmdoc VRTSvmman VRTSvmdev
```

Note: List the `VRTSvxvm` package before the `VRTSvmsa` package because the Volume Manager must be installed before the Storage Administrator.

The system prints out a series of status messages as the installation progresses. No errors should be reported, and the process should take about 5 minutes.



Disabling Hot-Relocation

The hot-relocation feature is enabled by default and it is recommended that you leave it on. However, if you choose to disable it for some reason, you can do so by preventing the `vxrelocd` daemon from starting up during system startup. This should be done after the VxVM packages have been installed.

To disable hot-relocation, comment out the `vxrelocd` line in the startup file `/etc/rc2.d/S95vxvm-recover`, as follows:

```
#vxrelocd root &
```

You should now do a:

```
reboot -- -r
```



Initializing the Volume Manager

This section describes how to initialize the Volume Manager using `vxinstall`. `vxinstall` places all specified disks under Volume Manager control. These disks are automatically placed in the default disk group, `rootdg`. If you want to place some of your disks in another disk group, you must use `vxinstall` to initialize at least one disk into `rootdg` and then use the Storage Administrator or `vxdiskadm` to place additional disks in other disk groups. Refer to the Volume Manager documentation set for information on how to create other disk groups for your disks.

Disks that are not initially placed under Volume Manager control by `vxinstall` can be added later using another Volume Manager interface (such as the Storage Administrator or `vxdiskadm`). However, you must place at least one disk under Volume Manager control with `vxinstall` before you can add disks using other utilities or interfaces.

Note: At least one disk must remain in `rootdg` at all times while VxVM is running.

Once you have completed the package installation, initialize the Volume Manager as follows:

1. Log in as superuser.
2. If you want to exclude one or more disks from Volume Manager control, add the names of those disks to the `/etc/vx/disks.exclude` file.

`vxinstall` ignores any disks listed in this file and leaves them untouched. The following is an example of the contents of the `disks.exclude` file (which you have to create):

```
c0t1d0
```

Similarly, you can exclude all disks on an entire controller by adding the name of that controller to the `/etc/vx/cntrls.exclude` file. The following is an example of the contents of the `cntrls.exclude` file:

```
c0  
c1
```

3. Enter the following to start `vxinstall`:

```
# vxinstall
```



`vxinstall` first runs the `vxlicense` command to initialize the Volume Manager license key file. `vxlicense` displays some licensing information and then prompts you for a key. You must obtain and enter a valid key in order to proceed with the initialization.

Note: The presence of certain hardware arrays (for example, A5000, SSA) automatically generates a key. `vxinstall` does not prompt for another key.

The `vxinstall` program then examines all controllers that it finds attached to the system and lists them:

```
Generating list of attached controllers....
```

```
Volume Manager Installation
Menu: VolumeManager/Install
```

The Volume Manager names disks on your system using the controller and disk number of the disk, substituting them into the following pattern:

If the Multipathing driver is installed on the system then for the disk devices with multiple access paths, the controller number represents a multipath pseudo controller number. For example, if a disk has 2 paths from controllers `c0` and `c1`, then the Volume Manager displays only one of them such as `c0` to represent both the controllers.

```
c<controller>t<disk>d<disk>
```

Some examples would be:

```
c0t0d0 - first controller, first target, first disk
c1t0d0 - second controller, second target, first disk
c1t1d0 - second controller, second target, first disk
```

The Volume Manager has detected the following controllers on your system:



```
c0: iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000
c1: iommu@f,e0000000/sbus@f,e0001000/dma@3,81000/esp@3,80000
c2: iommu@f,e0000000/sbus@f,e0001000/SUNW,soc@1,0/ \
    SUNW,pln@b0000125,52005088
```

Hit RETURN to continue.

Press Return to continue.

Next, vxinstall displays a brief introduction to the installation process:

Volume Manager Installation

Menu: VolumeManager/Install

You will now be asked if you wish to use Quick Installation or Custom Installation. Custom Installation allows you to select how the Volume Manager will handle the installation of each disk attached to your system.

Quick Installation examines each disk attached to your system and attempts to create volumes to cover all disk partitions that might be used for file systems or for other similar purposes.

If you do not wish to use some disks with the Volume Manager, or if you wish to reinitialize some disks, use the Custom Installation option. Otherwise, we suggest that you use the Quick Installation option.

Hit RETURN to continue.

Press Return to continue.

vxinstall then displays a menu with the following options:

```
1  Quick Installation
2  Custom Installation

?  Display help about menu
?? Display help about menuing system
q  Exit from menus
```



Select an operation to perform:

- 1 (Quick Installation) uses default options for installation.
 - 2 (Custom Installation) allows you to control whether and how each disk will be brought under Volume Manager control.
 - ? displays a help file describing the current operation or menu choices.
 - ?? displays general information about using vxinstall.
 - q exits from the current operation, or from the vxinstall program.
4. Select either 1 (Quick Installation) or 2 (Custom Installation) to proceed with the installation.
- If you choose Quick Installation, go to the section entitled “Quick Installation.”
 - If you choose Custom Installation, go to the section entitled “Custom Installation.”

Note: Quick Installation will bring every disk attached to your system under Volume Manager control, unless you have chosen to exclude those disks (through the `disks.exclude` or `cntrls.exclude` file, described earlier). If you have any disks on your system that you do not want placed under Volume Manager control, either list them in the `disks.exclude` file or choose the Custom Installation option (which allows you to leave certain disks alone).



Quick Installation

The Quick Installation option is the easiest to use. For each controller, Quick Installation gives you the option of either initializing or encapsulating all disks. If you want to encapsulate some disks on a given controller and initialize others, use Custom Installation instead.

Quick Installation allows you to initialize disks on a controller, even if valid partitions are found on those disks. However, you should avoid selecting the initialize option for any disks containing valid data.

With encapsulation, the `/etc/vfstab` file is updated to ensure that file systems previously mounted on disk partitions will be mounted on volumes instead.

To use the Quick Installation option, do the following:

1. Select menu item 1 (Quick Installation) from the `vxinstall` main menu.

`vxinstall` proceeds to ask you a series of questions for each controller and disk it finds connected to your system. When a default response is displayed in parentheses, you can simply press Return to accept that default. At any of the `vxinstall` prompts, you can select `q` to completely abandon the initialization and then start over.

Note: All disks are encapsulated or initialized (according to your instructions) at the end of the `vxinstall` procedure. If you quit `vxinstall` before it enters its final phase of actually initializing or encapsulating the disks, all disks will be left as they were before `vxinstall` started.

2. At the following prompt, indicate whether to encapsulate your boot disk (this example assumes your boot disk is named `c0t3d0`):

The `c0t3d0` disk is your Boot Disk. You can not add it as a new disk. If you encapsulate it, you will make your root file system and other system areas on the Boot Disk into volumes. This is required if you wish to mirror your root file system or system swap area.

Encapsulate Boot Disk [y,n,q,?] (default: n)

- If you enter `n`, your root and other file systems will be left unaltered.



- If you enter `y`, `vxinstall` asks you to enter a disk name for the boot disk and press Return to continue. `vxinstall` encapsulates your root file system as a volume, along with your swap device and all other disk partitions found on your boot disk. `/usr`, `/opt`, `/var`, and any other file systems on your boot disk are also encapsulated.

`vxinstall` now goes through each controller and asks you how to handle the disks contained on that controller. `vxinstall` begins this process for each controller by identifying the controller and generating a list of its disks:

```
Generating list of attached disks on c1....
```

```
The Volume Manager has detected the following disks on controller c1:
```

```
c1t0d0
```

```
c1t1d0
```

```
c1t2d0
```

```
c1t3d0
```

```
c1t4d0
```

```
Hit RETURN to continue.
```

If any disks are listed in the `/etc/vx/disks.exclude` file, they are listed here separately as excluded disks.

Press Return to continue.

3. At the following prompt, indicate whether you want to initialize all of the disks on the controller:

```
Volume Manager Quick Installation For Controller c1 Menu:  
VolumeManager/Install/QuickInstall/c1
```

```
Initialize all disks on this controller? (destroys data on these  
disks) [y,n,q,?] (default: n)
```

- If you enter `n`, `vxinstall` assumes that you want to encapsulate all disks on the controller instead and displays the following:

```
Volume Manager will now try to encapsulate all the disk on this  
controller.
```

```
Disks not having valid partitions will be initialized.
```

```
Hit RETURN to continue.
```



Press Return to continue.

- If you enter `y`, `vxinstall` prompts you for confirmation and warns you that this will destroy any existing data on the disks:

```
Are you sure (destroys data on these disks)
[y,n,q,?] (default: n)
```

CAUTION! Be careful when using this option as it may result in loss of data.

Enter `y` to continue with the disk initialization. The following message appears:

```
Volume Manager will now initialize all the disks on this
controller including disks having valid partitions.
```

```
Hit RETURN to continue.
```

Press Return to continue.

4. At the following prompt, indicate whether you want default disk names to be assigned to all of the disks on the controller:

```
Use default disk names for these disks? [y,n,q,?] (default: y)
```

- If you enter `y`, `vxinstall` automatically assigns and lists default disk names for each disk:

```
The clt0d0 disk will be given disk name disk01
The clt1d0 disk will be given disk name disk02
The clt2d0 disk will be given disk name disk03
The clt3d0 disk will be given disk name disk04
```

```
Hit RETURN to continue.
```

If encapsulation was selected earlier, this listing also includes the following message for each disk that can be encapsulated:

```
The clt0d0 disk has been configured for encapsulation.
```

or the following message for empty disks that will be initialized instead of encapsulated:

```
The c3t2d0 disk appears to be empty. Adding as a new disk.
```



- If you enter `n`, `vxinstall` prompts you for a disk name for each disk on the controller individually:

```
Enter disk name for c1t0d0 [<name>,q,?](default: disk01)
```

For each disk, respond with the desired disk name and press Return.

5. When all of the disks on the current controller have been named, press Return and repeat Steps 3 and 4 for the remaining controllers on your system.

When you have completed the `vxinstall` procedure for all controllers on your system, `vxinstall` displays a summary of the disks you have designated for initialization (New Disk) or encapsulation (Encapsulate) on each controller:

```
The following is a summary of your choices.
```

```
c0t5d0  New Disk
c1t0d0  New Disk
c1t1d0  New Disk
c1t2d0  New Disk
c2t2d3  New Disk
c3t0d0  Encapsulate
c3t1d0  Encapsulate
c3t2d0  New Disk
```

```
Is this correct [y,n,q,?] (default: y)
```

This is your last chance to review and alter your choices for how to handle any of the disks to be placed under Volume Manager control.

- If you enter `y`, `vxinstall` encapsulates all disks listed with Encapsulate and initializes (destroying any existing data) all disks listed with New Disk.
- If you enter `n`, `vxinstall` prompts you for the name of a disk to be removed from the list and excluded from Volume Manager control:

```
Enter disk to be removed from your choices. Hit return when done.
[<name>,q,?]
```



Enter the name of the disk to be removed from the list and press Return. Repeat for each disk to be removed and then press Return once more. `vxinstall` should now display an updated summary without the disks chosen for removal. If you are now satisfied with the summary, enter `y`.

Depending on how you chose to handle your disks during the `vxinstall`, you may be informed that you need to shut down your machine. If this happens, shut down and reboot your machine immediately or at the earliest convenient opportunity. Refer to “Rebooting After `vxinstall`” for further information.

CAUTION! Do not make any changes to your disk or file system configurations before shutting down and rebooting your system.

If a shutdown is not required, `vxinstall` indicates successful completion of the installation:

```
Installation is complete.
```



Custom Installation

Custom Installation allows you to control which disks are added to Volume Manager control and how they are added. This is important if you do not want all your disks under Volume Manager control. Custom Installation also gives you the option of initializing all disks on a controller, encapsulating all disks on a controller, or initializing some disks on a controller and encapsulating others.

To use the Custom Installation option, do the following:

1. Select menu item 2 (Custom Installation) from the `vxinstall` main menu.

`vxinstall` asks you a series of questions for each controller and disk it finds connected to your system. When a default response is displayed in parentheses, you can press Return to accept that default. At any of the `vxinstall` prompts, you can select `q` to completely abandon the initialization and then start over.

Note: All disks are encapsulated or initialized (according to your instructions) at the end of the `vxinstall` procedure. If you quit `vxinstall` before it enters its final phase of actually initializing or encapsulating the disks, all disks will be left as they were before `vxinstall` started.

2. At the following prompt, indicate whether to encapsulate your boot disk (this example assumes your boot disk is named `c0t3d0`):

The `c0t3d0` disk is your Boot Disk. You can not add it as a new disk. If you encapsulate it, you will make your root file system and other system areas on the Boot Disk into volumes. This is required if you wish to mirror your root file system or system swap area.

Encapsulate Boot Disk [y,n,q,?] (default: n)

- If you enter `n`, your root and other file systems will be left unaltered.
- If you enter `y`, `vxinstall` encapsulates your root file system as a volume, along with your swap device and all other disk partitions found on your boot disk. `/usr`, `/opt`, `/var`, and any other file systems on your boot disk are also encapsulated.



`vxinstall` now goes through each controller and asks you how to handle the disks contained on that controller. `vxinstall` begins this process for each controller by identifying the controller and generating a list of its disks:

```
Generating list of attached disks on c1....
```

```
The Volume Manager has detected the following disks on controller c1:
```

```
c1t0d0
```

```
c1t1d0
```

```
c1t2d0
```

```
c1t3d0
```

```
c1t0d4
```

```
Hit RETURN to continue.
```

If any disks are listed in the `/etc/vx/disks.exclude` file, they are listed here separately as excluded disks.

Press Return to continue.

3. At the following prompt, select the option that matches how you want the disks on this controller to be handled:

```
Installation options for controller c1
```

```
Menu: VolumeManager/Install/Custom/c1
```

```
1 Install all disks as pre-existing disks. (encapsulate)
```

```
2 Install all disks as new disks.(discards data on disks!)
```

```
3 Install one disk at a time.
```

```
4 Leave these disks alone.
```

```
?Display help about menu
```

```
??Display help about the menuing system
```

```
qExit from menus
```

```
Select an operation to perform:
```



- ▼ To encapsulate all disks on the controller, select 1.

Volumes will be created to encapsulate any partitions on all disks on this controller. The `/etc/vfstab` file will be updated to ensure that file systems previously mounted on disk partitions will be mounted as volumes instead.

- a. At the following prompt, indicate whether you want default disk names to be assigned to all of the disks on the controller.

```
Use default disk names for these disks? [y,n,q,?] (default: y)
```

- If you enter `y`, `vxinstall` automatically assigns and lists default disk names for each disk:

```
The clt0d0 disk will be given disk name disk01
```

```
The clt0d0 disk has been configured for encapsulation.
```

```
The clt1d0 disk will be given disk name disk02
```

```
The clt1d0 disk has been configured for encapsulation.
```

```
Hit RETURN to continue.
```

If any disks are found to be empty, the following message indicates that they will be initialized instead of encapsulated:

```
The c3t2d0 disk appears to be empty. Adding as a new disk.
```

- If you enter `n`, `vxinstall` prompts you for a disk name for each disk on the controller individually:

```
Enter disk name for clt0d0 [<name>,q,?](default: disk01)
```

For each disk, respond with the desired disk name and press Return.

- b. When all of the disks on the current controller have been named, press Return to move on to the next controller.



- ▼ To initialize all disks on the controller, select 2.

CAUTION! Be careful when using this option as it may result in loss of data.

All disks on this controller will be reinitialized. This destroys all data and partitions on each disk and makes the disk available as free space for allocating new volumes, or mirrors of existing volumes. If you are running `vxinstall` on new disks that have never been placed under Volume Manager control before, choose this option so that the disks will be repartitioned according to the Volume Manager's specifications.

- a. At the following prompt, indicate whether you want default disk names to be assigned to all of the disks on the controller:

```
Use default disk names for these disks? [y,n,q,?] (default: y)
```

- If you enter `y`, `vxinstall` automatically assigns and lists default disk names for each disk:

```
The clt0d0 disk will be given disk name disk01
```

```
The clt1d0 disk will be given disk name disk02
```

```
Hit RETURN to continue.
```

- If you enter `n`, `vxinstall` prompts you for a disk name for each disk on the controller:

```
Enter disk name for clt0d0 [<name>,q,?](default: disk01)
```

For each disk, respond with the desired disk name and press Return.

- b. When all of the disks on the current controller have been named, press Return to move on to the next controller.



- ▼ To install one disk at a time, select 3.

Each disk will be handled separately and you will be prompted for information on a per-disk basis. This allows you to install a disk as a pre-existing disk, install it as a new disk, or leave it alone.

- a. At the following prompt, select the option that matches how you want this disk to be handled:

```
Installation options for disk clt0d0
Menu: VolumeManager/Install/Custom/cl/clt0d0

1  Install as a pre-existing disk. (encapsulate)
2  Install as a new disk.(discards data on disks!)
3  Leave this disk alone.

? Display help about menu
??Display help about the menuing system
q Exit from menus
```

Select an operation to perform:

These options are similar to those in the Custom Installation main menu. Once you have selected how to handle this disk, follow the instructions provided for that option elsewhere in Step 3.

- b. Press Return to move on to the next disk. Repeat the previous step for the remaining disks on the current controller.
 - c. When all of the disks on the current controller have been named, press Return to move on to the next controller.
- ▼ To leave all disks on the controller unaltered, select 4.

No changes will be made to the disks and they will not be placed under Volume Manager control. If applications are currently using these disks and you do not want to upgrade these applications to use the Volume Manager, you can use this option to ensure that your applications can continue to use the disks without modification.



4. Repeat Step 3 for each of the remaining controllers on your system.

When you have completed the `vxinstall` procedure for all controllers on your system, `vxinstall` displays a summary of the disks you have designated for initialization (New Disk) or encapsulation (Encapsulate) on each controller:

The following is a summary of your choices.

c0t5d0	New Disk
c1t0d0	New Disk
c1t1d0	New Disk
c1t2d0	New Disk
c2t2d3	New Disk
c3t0d0	Encapsulate
c3t1d0	Encapsulate
c3t2d0	New Disk

Is this correct [y,n,q,?] (default: y)

This is your last chance to review and alter your choices for how to handle any of the disks to be placed under Volume Manager control.

- If you enter `y`, `vxinstall` proceeds to encapsulate all disks listed with Encapsulate and initialize (destroying any existing data) all disks listed with New Disk.
- If you enter `n`, `vxinstall` prompts you for the name of a disk to be removed from the list and excluded from Volume Manager control:

Enter disk to be removed from your choices. Hit return when done.

[<name>,q,?]

Enter the name of the disk to be removed from the list and press Return. Repeat for each disk to be removed and then press Return once more. `vxinstall` should now display an updated summary without the disks chosen for removal. If you are now satisfied with the summary, enter `y`.



Depending on how you chose to handle your disks during the `vxinstall`, you may be informed that you need to shut down and reboot your machine. If this happens, shut down and reboot your machine immediately or at the earliest convenient opportunity. Refer to the section “Rebooting After `vxinstall`” for further information.

CAUTION! Do not make any changes to your disk or file system configurations before shutting down and rebooting your system.

If a shutdown is not required, `vxinstall` indicates successful completion of the installation:

```
Installation is complete.
```



Rebooting After vxinstall

After you have selected how vxinstall should process all of the disks attached to your system, you may have to reboot the system to make changes to your disk partitioning that cannot be made while your disks are in use. The way in which you chose to handle your disks during the vxinstall session determines whether a shutdown and reboot is required. If you chose to encapsulate any disks, a reboot will definitely be necessary.

vxinstall should inform you when a shutdown and reboot is necessary. The shutdown and reboot should be conducted as soon as possible.

Note: Several reboots may be required, depending on the setup you chose.

When it is necessary to shut down and reboot your system, vxinstall should display a message similar to the following:

```
The system now must be shut down and rebooted in order to
continue the reconfiguration.
```

```
Shutdown and reboot now [y,n,q,?] (default: n)
```

If this happens, respond in either of the following ways:

- Enter `y` to begin an immediate shutdown.
- Enter `n` to cause vxinstall to exit without starting a shutdown. If you select this option, it is your responsibility to remember to shut down and reboot as soon as possible.

Note: During the next one or more reboots, you may be asked several times if you wish to continue an operation. Press the Return key at all of these prompts to accept the default answer. If you select a different answer from the default for any of these prompts or press `q`, the initialization may fail.



Setting Up and Starting the Storage Administrator

This section describes how to install and set up the Volume Manager Storage Administrator. The Storage Administrator server must be installed and run on a UNIX (Solaris) machine running the Volume Manager. The Storage Administrator client can be installed and run on any machine that supports the Java 1.1 Runtime Environment.

You can run the Storage Administrator client as an application or from a Web browser. However, it is recommended that you run the Storage Administrator as an application to avoid performance issues associated with Web browsers.

For information on how to upgrade from a previous release of the Storage Administrator, see “Upgrading the Storage Administrator.”

Note: This release of the Storage Administrator requires Volume Manager Release 2.5 or higher. Before you install and run the Storage Administrator, you must install the Volume Manager and run `vxinstall` to initialize the Volume Manager.

Storage Administrator Installation Overview

Set up and start the Storage Administrator as follows:

1. Install the Storage Administrator server and client on the machine to be administered. This must be a UNIX (Solaris) machine.
(See “Installing the Storage Administrator Server and/or Client (UNIX).”)
2. If users other than `root` need to access the Storage Administrator, set up security and specify which users can run the Storage Administrator.
(See “Customizing Security (Optional).”)
3. If you plan to run the Storage Administrator client from a Web browser, set up a HyperText Transfer Protocol (HTTP) server on the machine to be administered. (See “Setting Up the HTTP Server (Optional).”)

If you plan to run the Storage Administrator client as an application, you do not need to set up an HTTP server on the machine to be administered.



4. If you plan to run the Storage Administrator client from a machine other than the machine to be administered, install the Storage Administrator client on the machine where the client will run.
(See “Installing the Storage Administrator Server and/or Client (UNIX)” or “Installing the Storage Administrator Client (Windows).”)
5. Start the Storage Administrator server on the machine to be administered.
(See “Starting the Storage Administrator Server.”)
6. Start the Storage Administrator client. (See “Starting the Storage Administrator Client.”)

Installing the Storage Administrator Server and/or Client (UNIX)

Install the Volume Manager Storage Administrator server and client on a Solaris machine as follows:

1. Log in as superuser.
2. Install the Storage Administrator package (as described in “Loading the Software from CD-ROM”).

The Storage Administrator client can be installed with or without the Storage Administrator server. To install the Storage Administrator server, answer *y*(es) at the appropriate prompt. The Storage Administrator package also contains an optional Web server package. At the appropriate prompt, indicate whether the Web package should be installed.

The Storage Administrator package (UNIX) contains the following directories:

Directory	Contents
httpd/	apache_1.2.0.tar.z – Apache HTTP server (optional)
jre/	Java Runtime Environment (JRE)
bin/	vmsa_server – Storage Administrator server startup script vmsa – Storage Administrator startup script



Directory	Contents
vmsa/	java/vmsa.jar – Storage Administrator class archive java/images/ – Storage Administrator image .gif files java/help/ – Storage Administrator help files server/ – Storage Administrator server libraries and properties index.html – Storage Administrator applet HTML page
man/	Volume Manager Storage Administrator manual pages

3. Go to the directory where you installed the Storage Administrator (/opt/VRTSvmsa, by default).
4. Set the VMSAHOME environment variable:

export VMSAHOME=*Storage_Administrator_directory*

where *Storage_Administrator_directory* is the directory where you installed the Volume Manager Storage Administrator.
5. Update your PATH environment variable to include the directory containing the Storage Administrator startup scripts:

```
export PATH=$PATH:$VMSAHOME/bin
```



Installing the Storage Administrator Client (Windows)

The Volume Manager Storage Administrator client runs on Windows NT, Windows 98, or Windows 95 machines. If you plan to run the Storage Administrator from a Windows machine, install the optional Windows package after you have installed the Storage Administrator server and client on a UNIX machine.

Before you install the Storage Administrator on a Windows machine, you must uninstall any existing Storage Administrator packages on that machine. Only one Storage Administrator package can be installed on a Windows machine at any given time.

Install the Storage Administrator client on a Windows machine as follows:

1. Log in as administrator.
2. Insert the CD-ROM containing the Volume Manager Storage Administrator.
3. Using Windows Explorer or a DOS Command window, go to the `win32` directory and execute the `setup.exe` program.
4. Follow the instructions presented by the `setup.exe` program.
5. After installation is complete, do one of the following for environment changes made during installation to take effect:
 - For Windows NT, log out and then log back in.
 - For Windows 98 or Windows 95, restart the computer.



Customizing Security (Optional)

The Volume Manager Storage Administrator can only be run by users with appropriate privileges. By default, only `root` can run the Storage Administrator.

After you install the Storage Administrator, you can specify which users will have access to the Storage Administrator. To set up a list of users with permission to use the Storage Administrator, add a group named `vrtsadm` to the group file (`/etc/group`) or NIS (Network Information Name Service) group table on the machine to be administered. The `vrtsadm` group should include the user names of any users (including `root`) who will have access to the Storage Administrator.

Example:

```
vrtsadm::999:root,joe,bill
```

If the `vrtsadm` group does not exist, only `root` has access to the Storage Administrator. If `vrtsadm` exists, `root` must be included in `vrtsadm` in order for `root` to access the Storage Administrator.

Changing the Group Name (Optional)

`vrtsadm` is the default VERITAS group name. However, you can change `vrtsadm` to another name by setting the `vrts.server.adminGroup` property to another value. You can reset the VERITAS group name by adding the following to the `java` command at the end of the `/opt/VRTSvmsa/bin/vmsa_server` file:

```
-Dvrts.server.adminGroup=new_groupname
```



Monitoring Access to the Volume Manager Storage Administrator

Once you have set up security for the Volume Manager Storage Administrator, you can monitor access to the Storage Administrator by reviewing the contents of the access log file. By default, the access log file is `/var/opt/vmsa/logs/access`. This file contains entries similar to the following:

```
Fri Jan 23 10:22:17 PST 1998: user xyz login succeeded
Fri Jan 23 10:59:52 PST 1998: user xyz login failed with error
"User password invalid"
```

Entries for failed access may be logged multiple times. This is due to a security requirement and is not an error.

Setting Up the HTTP Server (Optional)

If you plan to access the Volume Manager Storage Administrator via a Web browser, you must set up and run an HTTP server on the machine to be administered. If an HTTP server is not set up on the machine to be administered, you can install the Apache HTTP server provided in the Storage Administrator package. However, you can use any other HTTP server.

Install and set up the Apache HTTP server as follows:

1. Go to the HTTP server directory.
2. Uncompress and unpack the `apache_1.2.0.tar.Z` file:

```
uncompress apache_1.2.0.tar.Z
tar xvf apache_1.2.0.tar
```

3. Read the `README` file in the top-level Apache directory.

For more detailed instructions on how to set up the Apache HTTP server, refer to the `README` file in the top-level Apache directory and the documentation on the Apache Web site: <http://www.apache.org/docs>.

4. Go to the `htdocs` directory.
5. Link the following directories:

```
ln -s $VMSAHOME/vmsa vmsa
```



6. Go to the `conf` directory.
7. Copy all files named `file.conf-dist` to `file.conf`.
8. Edit the `httpd.conf` and `srn.conf` configuration files for your HTTP server.

Make the appropriate changes to the following entries in the `httpd.conf` file:

- `ServerAdmin`
- `ServerRoot`
- `ServerName`

In the `srn.conf` file, set `DocumentRoot` to the path to the `htdocs` directory:

```
DocumentRoot /opt/VRTSvmsa/httpd/apache_1.2.0/htdocs
```

9. Start the HTTP server.

If you do not want to use the Apache HTTP server provided with the Storage Administrator, set up another HTTP server as follows:

1. Follow the installation instructions provided with your HTTP server.
2. Go to the `htdocs` (or equivalent) directory.
3. Link the following directories:

```
ln -s $VMSAHOME/vmsa vmsa
```

4. Start the HTTP server.



Starting the Storage Administrator Server

In order to use the Volume Manager Storage Administrator, the Storage Administrator server must be running on the Solaris machine to be administered. Once started, the Storage Administrator server should run continuously and should not be stopped. Only one instance of the Storage Administrator server should be running at a time.

Start the Storage Administrator server manually as follows:

1. Log in as superuser.
2. Go to the directory where the Storage Administrator is installed:

```
cd $VMSAHOME/bin
```

3. Start the server:

```
./vmsa_server &
```

A file named `server.log` will be created in the `/var/opt/vmsa/logs` directory. If the server starts successfully, this file should contain the following (after a few seconds):

```
rebound //host:2410/vrts.remote.vrtsServer
```

You can use the following command to confirm that the Storage Administrator is running:

```
vmsa_server -q
```

If you need to stop the Storage Administrator server for some reason, use this command:

```
vmsa_server -k
```

Note: Once the Storage Administrator is installed, the Storage Administrator server starts up automatically at system startup.



Starting the Storage Administrator Client

Note: Only users with appropriate privileges can run the Storage Administrator. Refer to the section “Customizing Security (Optional)” for information on Storage Administrator security and access permissions.

The Volume Manager Storage Administrator can administer the local machine or a remote machine. However, Volume Manager 2.5 (or higher) and the Storage Administrator server must be installed on the machine to be administered. The Volume Manager daemon (`vxconfigd`) and the Storage Administrator server must be running on the machine to be administered.

After installing the Volume Manager and the Storage Administrator and starting the server, you can run the Storage Administrator in one of the following ways:

▼ Starting the Storage Administrator as an application

UNIX

- ▼ To administer the *local* UNIX machine, type:

```
vmrsa
```

Make sure that the Session Initiation dialog box contains the local machine name.

- ▼ To administer a *remote* UNIX machine, type:

```
vmrsa remote_machine_name
```

Make sure that the Session Initiation dialog box contains a remote machine name.

Windows

- ▼ To administer a *remote* UNIX machine, choose Start > Programs > VERITAS Volume Manager Storage Administrator > VERITAS Volume Manager Storage Administrator.

Specify the remote machine name in the Session Initiation dialog box.



▼ Starting the Storage Administrator from a Web browser

Note: Although you can run the Storage Administrator from a Web browser, it is recommended that you run it as an application. Due to the nature of Web browsers, running the Storage Administrator from a Web browser can significantly slow down response time.

If you want to access the Storage Administrator via a Web browser, you must set up and run an HTTP server on the machine to be administered. See “Setting Up the HTTP Server (Optional)” for details.

- ▼ To access the Storage Administrator from a Web browser, start the browser and then go to the following location (URL):

`http://machine_name/vmsa/index.html`

Click on the VERITAS Volume Manager Storage Administrator link in the browser window. Specify the appropriate machine name in the Session Initiation dialog box.

The default `index.html` file provides access to a single host and starts the Storage Administrator applet immediately. This release includes two additional HTML files, `vmsa2.html` and `vmsa3.html` (located in the same directory as `index.html`). These sample files provide access to multiple hosts and load the Storage Administrator applet on request. You can customize the Storage Administrator startup behavior by editing one of these files. Each file contains comments that describe the changes you need to make to customize the file for your environment.

The Storage Administrator uses the Java Development Kit (JDK) version 1.1. Currently, the only browsers that fully support JDK 1.1 are:

- HotJava 1.0 or higher

You can download HotJava from `http://www.javasoft.com`.

Before you run the Storage Administrator from HotJava, set Medium Security for Unsigned Applets (Edit > Preferences > Applet Security). If security is needed, make sure that the `CLASSPATH` environment variable is set to `/opt/VRTSvmsa/vxvm/java` (UNIX) or `%VMSAHOME%\vxvm\java` (Windows).



If security is not needed, you can disable Storage Administrator security by setting the `vrts.security` property to `false` in the `/opt/VRTSvmsa/vxvm/properties` file and changing the security parameter value to `false` in the Storage Administrator `index.html` file.

- Netscape 4.0x or higher with JDK 1.1 support

You can download Netscape from <http://www.netscape.com>.

Before you run the Storage Administrator from Netscape, add the following to your `preferences.js` (UNIX) or `prefs.js` (Windows) preferences file:

```
user_pref("signed.applets.codebase_principal_support", true);
```

Note: You must exit *all* instances of Netscape before you modify the user preferences file. Otherwise, your changes to the preferences file will be overwritten when Netscape exits.

Running the Storage Administrator in Read-Only Mode

The Volume Manager Storage Administrator server can be configured in a read-only mode that is useful for monitoring or browsing purposes. Read-only mode allows the administrator to view objects on the system through the Storage Administrator, but prevents administrative actions from taking effect. This mode is enabled via the properties file (`/opt/VRTSvmsa/vmsa/properties`).

To start the Storage Administrator server in read-only mode, add the following line to the properties file:

```
vrts.server.readonly=true
```

To restore the Storage Administrator server to full operational mode, use the following line:

```
vrts.server.readonly=false
```

Note: You must stop and restart the server for this change to take effect.

The default mode for the Storage Administrator is full operational mode.



Enabling Cluster Support in Volume Manager (Optional)

This Volume Manager release includes an *optional* cluster feature that enables VxVM to be used in a cluster environment. The cluster functionality in the Volume Manager allows multiple hosts to simultaneously access and manage a set of disks under Volume Manager control. A *cluster* is a set of hosts sharing a set of disks; each host is referred to as a *node* in the cluster.

Note: The Volume Manager cluster feature requires a license, which can be obtained from your Customer Support channel. (The presence of a SPARCstorage™ Array may serve as a license, but limits what may be done to private disk groups.)

Enable the cluster functionality in the Volume Manager as follows:

1. Obtain a license for the Volume Manager cluster feature.
2. Install the software packages onto each system (node) to be included in the cluster (as described in the section “Loading the Software from CD-ROM”).
3. Initialize VxVM using `vxinstall` (as described in the section “Initializing the Volume Manager”).
4. Set up and start the Storage Administrator (as described in “Setting Up and Starting the Storage Administrator”).
5. Configure shared disks (as described in “Configuring Shared Disks”).

Configuring Shared Disks

This section describes how to configure shared disks. If you are installing VxVM for the first time or adding disks to an existing cluster, you need to configure new shared disks. If you are upgrading VxVM, you should verify that your shared disks still exist.

Note: The shared disks should be configured from one node only. Since the Volume Manager software cannot tell whether a disk is shared or not, you must specify which are the shared disks.



Make sure that nobody else is accessing the shared disks from another node while you are performing the configuration. If you start the cluster on the node where you do the configuration only, you can prevent disk accesses from other nodes because the quorum control reserves the disks for the single node (see “Disk Controller Reservation”).

Note: This section assumes that the Sun Cluster cluster manager is installed on your system.

Configuring New Disks

If you are installing and setting up VxVM for the first time, configure the shared disks as follows:

1. Start the cluster on at least one node.
2. Run `vxdisksetup` to initialize each shared disk.

If you have decided not to put configuration information on every disk, or if you want larger areas for this information, `vxdisksetup` allows you to specify your choices.

3. Create disk groups on the shared disks.

You can use `vxdbg` or the Storage Administrator to do this. Use the `-s` option of `vxdbg` to create shared disk groups.

4. Create volumes in the disk groups.

You can use `vxassist` or the Storage Administrator to do this.

The volumes must be of type `gen`. Do not create RAID-5 volumes. Before creating any log subdisks, read the section on Dirty Region Logging in the *Volume Manager Cluster Functionality* document.

5. If the cluster is only running with one node, bring up the other cluster nodes.

When each node is ready, enter the command `vxdbg list` on it. This should display the shared disk groups.



Verifying Existing Shared Disks

If you are upgrading from a previous release of VxVM, verify that your shared disk groups still exist:

1. Start the cluster on all nodes.
2. Enter the following command on all nodes:

```
# vxdg list
```

This should display the shared disk groups that existed before.

Converting Existing VxVM Disks to Shared Disks

If you are upgrading from VxVM 2.x to VxVM 3.0.4 and you want to convert existing disk groups to shared disk groups, configure the shared disks as follows:

1. Start the cluster on at least one node.

For a two-node cluster, start the cluster on one node; for a four-node cluster, start the cluster on three nodes.

2. Issue the following commands:

List all disk groups by entering the command:

```
# vxdg list
```

Deport disk groups to be shared with the command:

```
# vxdg deport groupname
```

Import disk groups to be shared with the command:

```
# vxdg -s import groupname
```

This will mark the disks in the shared disk groups as shared and stamp them with the ID of the cluster, enabling other nodes to recognize the shared disks.

If there are dirty region logs, make sure they are active. If not, replace them with bigger ones.

The following command should now display the shared flag for all the shared disk groups:



```
# vxdg list
```

The disk groups are now ready to be shared.

3. If the cluster is only running with one node, bring up the other cluster nodes.

When each node is ready, enter the command `vxdg list` on it. This should display the same list of shared disk groups that appeared earlier.

Disk Controller Reservation

As part of its quorum control, the Sun Cluster cluster manager reserves the shared disk controllers when only one node is active. This prevents “rogue” hosts from accessing the shared disks. When this happens, the command `vxdisk list` on a node that has left the cluster may show all disks on such a controller as having an `error` status. The more detailed options of `vxdisk` will show the flag `unavailable`. When a new node joins the cluster, the Sun Cluster software releases the controllers. VxVM attempts to access these disks, and if that is successful, the disks return to an `online` status. (See the Sun Cluster documentation for further details.) If one system boots while the other system has the disks reserved, the disks may be invisible to the booting system, and `vxdisk` may display none of the shared disks. When the system joins the cluster, the shared disks become visible.



Upgrading to VxVM Release 3.0.4

If you are already running an earlier release of the Volume Manager, you can upgrade to VxVM Release 3.0.4 with one of the procedures that follow. The procedures in this section describe how to upgrade the Volume Manager, Solaris, or both. This section also contains information on how to upgrade the Storage Administrator.

Note: The upgrade procedures allow you to retain your existing VxVM configuration. After upgrading, you can resume using VxVM as before (without running `vxinstall` again).

Note: Importing a pre-3.x Volume Manager disk group does not automatically upgrade the disk group version to the VxVM 3.x level. You may need to manually upgrade each of your disk groups following a VxVM upgrade. See the chapter on disk tasks in the *VERITAS Volume Manager Administrator's Guide* for information on how to upgrade a disk group.



Table 1 shows which VxVM releases are supported on each Solaris release.

Table 1 Solaris and VxVM Upgrade Matrix

	Solaris 2.3	Solaris 2.4	Solaris 2.5	Solaris 2.5.1	Solaris 2.6	Solaris 7	Solaris 8
VxVM 3.0.4				Supported	Supported	Supported	Supported
VxVM 3.0.3				Supported	Supported	Supported	Supported
VxVM 3.0.2				Supported	Supported	Supported	
VxVM 3.0.1				Supported	Supported	Supported	
VxVM 3.0				Supported	Supported		
VxVM 2.5.4			Supported	Supported	Supported		
VxVM 2.5.3			Supported	Supported	Supported		
VxVM 2.5.2			Supported	Supported	Supported		
VxVM 2.5.1			Supported	Supported	Supported		
VxVM 2.5			Supported	Supported	Supported		
VxVM 2.4.1		Supported	Supported	Supported	Supported		
VxVM 2.4		Supported	Supported	Supported	Supported		
VxVM 2.3		Supported	Supported	Supported			
VxVM 2.2.1		Supported	Supported	Supported			



Table 1 Solaris and VxVM Upgrade Matrix *(Continued)*

	Solaris 2.3	Solaris 2.4	Solaris 2.5	Solaris 2.5.1	Solaris 2.6	Solaris 7	Solaris 8
VxVM 2.2		Supported	Supported				
VxVM 2.1	Supported	Supported					
VxVM 2.0	Supported						



Table 2 shows all possible VxVM/Solaris upgrade combinations. After reviewing this table, proceed to the appropriate section.

Table 2 Upgrade Reference

To perform these upgrades	Go to this section
VxVM and Solaris 2.5.1/2.6/7/8: VxVM 2.1 & Solaris 2.4 ► VxVM 3.0.4 & Solaris 2.5.1/2.6/7/8 VxVM 2.2 & Solaris 2.4/2.5 ► VxVM 3.0.4 & Solaris 2.5.1/2.6/7/8 VxVM 2.2.1 & Solaris 2.4/2.5/2.5.1 ► VxVM 3.0.4 & Solaris 2.6/7/8 VxVM 2.3 & Solaris 2.4/2.5/2.5.1 ► VxVM 3.0.4 & Solaris 2.5.1/2.6/7/8 VxVM 2.4/2.4.1 & Solaris 2.4/2.5/2.5.1 ► VxVM 3.0.4 & Solaris 2.6/7/8 VxVM 2.5/2.5.x & Solaris 2.5/2.5.1/2.6 ► VxVM 3.0.4 & Solaris 2.5.1/2.6/7/8	“Upgrading to VxVM 3.0.4 and Solaris 2.5.1/2.6/7/8”
VxVM 3.0 & Solaris 2.5.1/2.6 ► VxVM 3.0.4 & Solaris 2.5.1/2.6/7/8	
VxVM 3.0.1/3.0.2 & Solaris 2.5.1/2.6/7 ► VxVM 3.0.4 & Solaris 2.5.1/2.6/7/8	
VxVM Only: VxVM 2.2.1 & Solaris 2.5.1 ► VxVM 3.0.4 & Solaris 2.5.1 VxVM 2.3 & Solaris 2.5.1 ► VxVM 3.0.4 & Solaris 2.5.1 VxVM 2.4 & Solaris 2.5.1 ► VxVM 3.0.4 & Solaris 2.5.1 VxVM 2.4.1 & Solaris 2.5.1/2.6 ► VxVM 3.0.4 & Solaris 2.5.1/2.6 VxVM 2.5/2.5.x & Solaris 2.5.1/2.6 ► VxVM 3.0.4 & Solaris 2.5.1/2.6	“Upgrading VxVM Only”
VxVM 3.0 & Solaris 2.5.1/2.6 ► VxVM 3.0.4 & Solaris 2.5.1/2.6	
VxVM 3.0.1/3.0.2 & Solaris 2.5.1/2.6/7 ► VxVM 3.0.4 & Solaris 2.5.1/2.6/7	
Solaris Only:	“Upgrading Solaris Only”



Table 2 Upgrade Reference (*Continued*)

VxVM 3.0.3 & Solaris 2.5.1 ► VxVM 3.0.4 & Solaris 2.6/7/8	
VxVM 3.0.3 & Solaris 2.6 ► VxVM 3.0.4 & Solaris 7/8	
VxVM 3.0.3 & Solaris 7 ► VxVM 3.0.4 & Solaris 8	

Upgrading to VxVM 3.0.4 and Solaris 2.5.1/2.6/7/8

This section provides detailed instructions for upgrading both Solaris and VxVM. Table 1 should be used as a reference to determine if an upgrade of the operating system is required in order to upgrade to VxVM 3.0.4.

Update both VxVM and Solaris as follows:

Note: It is important that you follow these steps in the specified order.

1. Obtain and install a VxVM 3.0.4 license key.

Refer to the section entitled “Obtaining a License Key” for details.

2. If any of the file systems `/`, `/usr`, `/var`, or `/opt` are defined on volumes, make sure that at least one plex for each of those volumes is formed from a single subdisk that begins on a cylinder boundary.

This is necessary because part of the upgrade process involves temporarily converting file systems on volumes back to using direct disk partitions, and Solaris requires that disk partitions start on cylinder boundaries. The upgrade scripts automatically convert file systems on volumes back to using regular disk partitions, as necessary. If the upgrade scripts detect any problems (such as lack of cylinder alignment), they display an explanation of the problem and the upgrade does not proceed.

3. Load and mount the CD-ROM as described in the section entitled “Loading the Software from CD-ROM.”

It should then be visible as the file system mounted on `/CD_path`.

If the CD-ROM is mounted automatically, its path should be `/cdrom/CD_name`; if mounted manually, the path should be `/mnt`.

4. Run the `upgrade_start` script to prepare the previous release of the Volume Manager for its removal:



```
# CD_path/scripts/upgrade_start
```

The `upgrade_start` script looks for volumes containing file systems. If certain key file systems must be converted back to using partitions, the script will indicate that a reboot is necessary. If so, reboot to single-user mode before proceeding to the next step. If any volumes are in use, you must either unmount those volumes or reboot to single-user mode.

5. Reboot to single-user mode (using a command such as `/etc/shutdown`).
6. Remove the VxVM 1.3/3.0 packages by entering:

```
# pkgrm VRTSvxvm VRTSvmsa VRTSvmdoc VRTSvmdev
```

Note: If you have made any additional modifications or have applied any patches, those must be removed before removing the `VRTSvxvm` package. Also, if you are upgrading from Volume Manager version 1.x or 2.x.y, then the `VRTSvxva` package must be removed instead of the `VRTSvmsa` package.

7. Shut down and halt the machine by using a command such as `/etc/shutdown`.
8. Upgrade the operating system to Solaris 8.
9. Reboot the machine to single-user mode (using a command such as `/etc/shutdown`).

It is now necessary to remount the CD-ROM manually:

```
# mount -F hsfs -o ro /dev/c0t6d0s2 /mnt
```

10. Add the VxVM 3.0.x package by typing:

```
# pkgadd -d CD_path/OS_version/pkgs VRTSvxvm
```

If this returns warnings that include the string `/etc/vx`, you can ignore them and continue.

11. Complete the upgrade by entering:

```
# CD_path/scripts/upgrade_finish
```

If you are upgrading from VxVM 1.3, this will prompt you for licensing information.



12. If you want to disable the hot-relocation feature, comment out the `vxrelocd` line in the startup file `/etc/rc2.d/S95vxvm-recover`, as follows:

```
#vxrelocd root &
```

13. Reboot to multi-user mode.

At this point, your pre-upgrade configuration should be in effect and any file systems previously defined on volumes should be defined and mounted.

14. Using `pkgadd`, install any or all of the additional packages `VRTSvmsa`, `VRTSvmdoc`, `VRTSvmman`, and `VRTSvmdev`.
15. Importing a pre-3.x Volume Manager disk group does not automatically upgrade the disk group version to the VxVM 3.x level. You may need to manually upgrade each of your disk groups following a VxVM upgrade. See the chapter on disk tasks in the *VERITAS Command Line Interface Administrator's Guide* for information on how to upgrade a disk group.
16. Follow the instructions in "Setting Up and Starting the Storage Administrator" to set up and restart the Storage Administrator.

Note: If you were using the VxVM 1.3 workaround for Prestoserve™, you now need to implement the VxVM 3.0.4 workaround for Prestoserve. (See the VERITAS Volume Manager Release Notes for details.)

Upgrading VxVM Only

If you are already running a version of Solaris that is supported with VxVM 3.0.4, you can upgrade VxVM without upgrading Solaris. Upgrade VxVM only as follows:

1. Remove the `VRTSvmsa` package (if necessary):

```
# pkgrm VRTSvmsa
```

Note: If you are upgrading from Volume Manager version 1.x or 2.x.y, then the `VRTSvxva` package must be removed instead of the `VRTSvmsa` package.

2. Add the VxVM 3.0.4 package by typing:



```
# pkgadd -d CD_path/OS_version/pkgs VRTSvxvm
```

3. Using `pkgadd`, install any or all of the additional packages `VRTSvmsa`, `VRTSvmdev`, `VRTSvmman`, and `VRTSvmdev`.
4. Follow the instructions in “Setting Up and Starting the Storage Administrator” to set up and restart the Storage Administrator.

The Storage Administrator Release 3.0.2 server is not backward compatible with Storage Administrator Release 1.x clients, so you must upgrade any existing client(s) to Release 3.0.2.

Depending on your existing configuration, various messages and prompts may result from the `pkgadd`. Answer the prompts appropriately.

Messages similar to the following may list conflicting files and mention scripts that will be executed with superuser permission during installation:

```
## Checking for conflicts with packages already installed.
The following files are already installed on the system and are
being used by another package:
files...
Do you want to install these conflicting files [y,n,?,q]

Do you want to install these as setuid/setgid files [y,n,?,q]

This package contains scripts which will be executed with super-
user permission during the process of installing this package.

Do you want to continue with the installation of <VRTSvxvm.2>
[y,n,?]
```

To ensure that the VxVM package is installed correctly, you should enter `y` in response to this set of messages.

It may be necessary to reboot after running `pkgadd`. If this is the case, you should be informed that a reboot is required.

If you want to disable the hot-relocation feature for some reason, you can do so by preventing the `vxrelocd` daemon from starting up during system startup. This should be done after the VxVM packages have been installed.



To disable hot-relocation, comment out the `vxrelocd` line in the startup file `/etc/rc2.d/S95vxvm-recover`, as follows:

```
#vxrelocd root &
```

After editing this file, the system must be rebooted:

```
# reboot -- -r
```



Upgrading Solaris Only

If you are already running VxVM 3.0.4 with Solaris 2.5.1/2.6/7/8, you can upgrade to Solaris 2.6/7/8, as follows:

1. If any of the file systems `/`, `/usr`, `/var`, or `/opt` are defined on volumes, make sure that at least one plex for each of those volumes is formed from a single subdisk that begins on a cylinder boundary.

This is necessary because part of the upgrade process involves temporarily converting file systems on volumes back to using direct disk partitions, and Solaris requires that disk partitions start on cylinder boundaries. The upgrade scripts automatically convert file systems on volumes back to using regular disk partitions, as necessary. If the upgrade scripts detect any problems (such as lack of cylinder alignment), they display an explanation of the problem and the upgrade does not proceed.

2. Run the `upgrade_start` script to prepare the previous release of the Volume Manager for its removal:

```
# CD_path/scripts/upgrade_start
```

The `upgrade_start` script looks for volumes containing file systems. If certain key file systems must be converted back to using partitions, the script will indicate that a reboot is necessary. If so, reboot to single-user mode before proceeding to the next step. If any volumes are in use, you must either unmount those volumes or reboot to single-user mode.

3. Shut down and halt the machine (using a command such as `/etc/shutdown`).
4. Upgrade the operating system to Solaris 2.6, 7, or 8.
Refer to the Solaris installation documentation for instructions on how to upgrade Solaris.
5. Reboot to single-user mode (using a command such as `/etc/shutdown`).
6. Complete the upgrade by entering:

```
# CD_path/scripts/upgrade_finish
```

Reboot to multi-user mode (using a command such as `/etc/shutdown`).



Upgrading the Storage Administrator

If a previous release of the Storage Administrator is already installed, you need to remove the previous VRTSvmsa package and then add the current VRTSvmsa package. Upgrade the Storage Administrator as follows:

1. Remove the previous VRTSvmsa package with the command:

```
pkgrm VRTSvmsa
```

2. Follow the instructions in “Loading the Software from CD-ROM” to add the current VRTSvmsa package.

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.

Checking Volume Manager Processes

After the Volume Manager has been successfully installed and you have rebooted your system, you can use the following command to confirm that key Volume Manager processes (vxconfigd, vxnotify, and vxrelocd) are running:

```
ps -ef | grep vx
```

This should display output similar to the following:

```
root 458   1      80      10:23:34 ?    0:00  vxrelocd root
root 414   1      80      10:23:06 ?    0:03  vxconfigd
root 478  458     39      10:23:37 ?    0:00  vxrelocd root
root 479  478      5      10:23:37 ?    0:00  vxnotify -f -w 15
```

Note: If you have disabled Hot Relocation, only the vxconfigd process will be running.



Deinstalling the Volume Manager

The following procedures can be used to deinstall the Volume Manager.

If `root`, `swap`, `usr`, or `var` is a volume under Volume Manager control, perform the following steps:

1. Ensure that `rootvol`, `swapvol`, `usr`, and `var` volumes have only one associated plex each.

The plex must be contiguous, non-striped, non-spanned, and non-sparse. This information can be determined by issuing the following command:

```
# vxprint -ht rootvol swapvol usr var
```

If any of these volumes have more than one plex associated, remove the unnecessary ones. Unnecessary plexes can be removed using the command:

```
# vxplex -o rm dis plex_name
```

The plex remaining for `rootvol` should have been created using `vxrootmir` since this guarantees that the underlying subdisks will start on cylinder boundaries and that partitions will be created for them.

2. Run `vxunroot`:

```
# /etc/vx/bin/vxunroot
```

This will change the volume entries in `/etc/vfstab` to the underlying disk partitions for `rootvol`, `swapvol`, `usr`, and `var`. It will also modify `/etc/system` and ask for a reboot so that disk partitions are mounted instead of volumes for `root`, `swap`, `usr`, and `var`.

Once the `root`, `swap`, `usr`, and `var` volumes have been taken care of, all remaining volumes need to be moved to disk partitions. This can be done in one of three ways:

- ▼ One approach is to back up the system fully onto tape and then recover from it.
- ▼ Another approach is to back up each file system individually and then recover them all after creating new file systems on disk partitions.
- ▼ Yet another approach is to incrementally move volumes onto disk partitions using the following steps:



1. Evacuate as many disks as possible using `vxdiskadm`, the Storage Administrator, or the `vxevac` script.

Evacuation moves subdisks from the specified disks to target disks. The evacuated disks will provide the initial free disk space for volumes to be moved to disk partitions.

2. Remove the evacuated disks from Volume Manager control using:

```
# vxdg rmdisk diskname
```

```
# vxdisk rm devname
```

3. Decide which volume to move first. If the volume to be moved is mounted, unmount it.
4. If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that data on the volume is `syncd`.
5. Create a partition on free disk space of the same size as the volume.

If there is not enough free space for the partition, a new disk must be added to the system at least for the first volume which is removed. Subsequent volumes can use the free space generated by the removal of this volume.

6. `dd` the data on the volume onto the newly-created disk partition using a command similar to the following:

```
# dd if=/dev/vx/dsk/lhome of=/dev/dsk/c2t2d2s7
```

where `c2t2d2` is the disk outside of Volume Manager and `s7` is the newly-created partition.

7. Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly-created partition.
8. Mount the disk partition if the corresponding volume was previously mounted.
9. Remove the volume from the Volume Manager using the command:

```
# vxedit -rf rm volume_name
```
10. Remove any disks that have become free (have no subdisks defined on them) by removing volumes from Volume Manager control.



To check if there are still some subdisks remaining on a particular disk, use the command:

```
# vxprint -F "%sd num" diskname
```

If the output is not 0, there are still some subdisks on this disk which will get removed subsequently. If the output is 0, remove the disk from Volume Manager control as follows:

```
# vxdg rmdisk diskname
```

```
# vxdisk rm devname
```

11. The free space now created can be used for adding the data in the next volume to be removed.
12. After all volumes have been converted into disk partitions successfully, reboot the system.

After reboot, none of the volumes should be found open. This can be verified by using the command:

```
# vxprint -Aht -e v_open
```

If any volumes remain open, repeat the steps listed above.

Shutting Down the Volume Manager

To shut down the Volume Manager, enter:

```
# vxdctl stop
```

```
# vxiod -f set 0
```

Removing the Volume Manager Package

To remove the installed Volume Manager packages, enter:

```
# pkgrm VRTSvmsa VRTSvxvm VRTSvmdoc VRTSvmman VRTSvmdev
```

If VxVM 2.x is still installed and you want to remove VxVM 3.x, it is necessary to remove the VxVM 2.x package first. Use the following command to determine whether both packages are present:

```
# pkginfo | grep vxvm
```



This should provide output that lists both the `VRTSvxvm` and `VRTSvxvm.2` packages. Remove both packages with the command:

```
# pkgrm VRTSvmsa VRTSvxvm VRTSvxvm.2 VRTSvmdoc VRTSvmmman \  
VRTSvmdev
```

