

VERITAS Volume Manager™ 3.1

Installation Guide

Solaris

July 2000
30-000002-399


VERITAS

Disclaimer

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

Copyright

Copyright © 1998-2000 VERITAS Software Corporation. All rights reserved. VERITAS is a registered trademark of VERITAS Software Corporation in the US and other countries. The VERITAS logo and VERITAS Volume Manager are trademarks of VERITAS Software Corporation. All other trademarks or registered trademarks are the property of their respective owners.

Printed in the USA, July 2000.

VERITAS Software Corporation
1600 Plymouth St.
Mountain View, CA 94043
Phone 650-335-8000
Fax 650-335-8050
www.veritas.com



Contents

Preface	v
Introduction	v
Organization	v
Conventions	vi
Getting Help	vi
Chapter 1. Pre-installation Instructions	7
Introduction	7
Pre-installation Overview	7
Obtaining a License Key	11
SDS to VxVM Conversion	11
Chapter 2. Installing the Volume Manager	13
Introduction	13
Procedure Overview	14
Installation Overview	14
Loading the Software from CD-ROM	15
Mounting the CD-ROM With Solaris Volume Management	16
Mounting the CD-ROM at the Command Line	16
Disabling Hot-Relocation	16
Initializing the Volume Manager	17
Quick Installation	20
Custom Installation	24
Rebooting After vxinstall	29



Setting Up and Starting the Storage Administrator	30
Storage Administrator Installation Overview	31
Installing the Storage Administrator Server and/or Client (UNIX)	31
Installing the Storage Administrator Client (Windows)	32
Customizing Security (Optional)	34
Starting the Storage Administrator Server	35
Starting the Storage Administrator Client	37
Running the Storage Administrator in Read-Only Mode	37
Enabling Cluster Support in Volume Manager (Optional)	38
Configuring Shared Disks	39
Disk Controller Reservation	41
Upgrading to VxVM Release 3.1	41
Upgrading VxVM and Solaris	44
Upgrading VxVM Only	46
Upgrading Solaris Only	47
Upgrading the Storage Administrator	48
Checking Volume Manager Processes	48
Deinstalling the Volume Manager	49
Shutting Down the Volume Manager	51
Removing the Volume Manager Package	51



Preface

Introduction

This document provides information on installing and initializing the VERITAS® Volume Manager (VxVM®) 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.

Organization

This guide is organized with the following sections:

- ◆ [Pre-installation Overview](#)
- ◆ [Introduction](#)
- ◆ [Loading the Software from CD-ROM](#)
- ◆ [Initializing the Volume Manager](#)
- ◆ [Setting Up and Starting the Storage Administrator](#)
- ◆ [Enabling Cluster Support in Volume Manager \(Optional\)](#)
- ◆ [Upgrading to VxVM Release 3.1](#)
- ◆ [Checking Volume Manager Processes](#)
- ◆ [Deinstalling the Volume Manager](#)



Conventions

The following table describes the typographic conventions used in this guide.

Typeface	Usage	Examples
monospace	Computer output, files, directories, software elements such as command options, function names, and parameters	Read tunables from the <code>/etc/vx/tunefstab</code> file. See the <code>ls(1)</code> manual page for more information.
monospace (bold)	User input	<code># mount -F vxfs /h/filesys</code>
<i>italic</i>	New terms, book titles, emphasis, variables replaced with a name or value	See the <i>User's Guide</i> for details. The variable <code>ncsize</code> determines the value of...

Symbol	Usage	Examples
%	C shell prompt	
\$	Bourne/Korn shell prompt	
#	Superuser prompt (all shells)	
\	Continued input on the following line; you do not type this character	<code># mount -F vxfs \ /h/filesys</code>
[]	In a command synopsis, brackets indicates an optional argument	<code>ls [-a]</code>
	In a command synopsis, a vertical bar separates mutually exclusive arguments	<code>mount [suid nosuid]</code>

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

Email: support@veritas.com



Introduction

The following topics are covered in this chapter:

- ◆ [Pre-installation Overview](#)
 - ◆ [Obtaining a License Key](#)
 - ◆ [SDS to VxVM Conversion](#)

Pre-installation Overview

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` must exist and be writable.
- ◆ In order to initialize the Volume Manager, 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. See “[Obtaining a License Key](#)” on page 11 for details.
- ◆ Before initializing the Volume Manager, make sure you are aware of the contents of each disk. During the initialization process, you must 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 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 do not reboot, the initialization does 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 must be at the beginning or end of the disk and must 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 attempts to allocate space usually reserved for swap by shrinking the swap partition (which must be slice 1). This process is known as *swap relocation* and, if necessary, happens automatically during root disk encapsulation.

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

- ◆ Disks to be encapsulated must:
 - ◆ contain the required minimum free space (1024 sectors)
 - ◆ contain an `s2` slice that represents the full disk—this `s2` slice cannot contain a file system
 - ◆ contain two free partitions
- ◆ To encapsulate the boot disk, tag the swap partition as `swap` so that it is possible to dump to that partition later.



Before encapsulating your boot disk, set the EEPROM variable `use-nvramrc?` to **true** 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 must determine which disks are bootable yourself. 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 indicates 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)	33M

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	33M



- ◆ The hot-relocation feature is enabled by default. The associated daemon, `vxrelocd`, is automatically started during system startup. To disable hot-relocation, prevent `vxrelocd` from running after you load the VxVM software. See “[Disabling Hot-Relocation](#)” on page 16 and “[Upgrading VxVM and Solaris](#)” on page 44 for details.
- ◆ The following software must be installed and running on the machine to be administered by the Volume Manager Storage Administrator:
 - ◆ Solaris 2.6 or higher.
 - ◆ VERITAS Volume Manager Release 2.5.7, 3.0.4, or 3.1 (or higher). The `vxconfigd` daemon must be running on the machine to be administered.
 - ◆ The Volume Manager Storage Administrator server.
- ◆ 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](#)” on page 48 for details.
- ◆ To install the Volume Manager in a cluster environment, use the following procedure:
 - ◆ Obtain a license for the *optional* Volume Manager cluster feature from your Customer Support channel.
 - ◆ Before installing the Volume Manager, 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 the `rootdg` group. The `rootdg` group must exist and cannot be shared between systems. At least one disk must exist within the `rootdg` group while VxVM is running.
 - ◆ Before beginning the installation, 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, 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) is sufficient. For every additional two gigabytes of volume size, the log size increases by approximately one block per map (for example, a four-gigabyte volume with two nodes has 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 “Dirty Region Logging” in *Volume Manager Cluster Functionality* for more information about log sizes.



Obtaining a License Key

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

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

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

- ◆ the host ID
- ◆ the machine type

Obtain the host ID using the following command:

```
# hostid
```

Obtain the host machine type using the following command:

```
# uname -i
```

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

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

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

View the currently installed licenses with this command:

```
# vxlicense -p
```

SDS to VxVM Conversion

If Solstice™ DiskSuite™ (SDS) is configured on your system, VERITAS tools are available to assist you in converting your system to Volume Manager. VERITAS has a conversion tool which enables a *data conversion in-place*. This tool allows 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.





Introduction

The following topics are covered in this chapter:

- ◆ Procedure Overview
 - ◆ Installation Overview
- ◆ Loading the Software from CD-ROM
 - ◆ Mounting the CD-ROM With Solaris Volume Management
 - ◆ Mounting the CD-ROM at the Command Line
 - ◆ 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)
 - ◆ 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 3.1](#)
 - ◆ [Upgrading VxVM and Solaris](#)
 - ◆ [Upgrading Solaris 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](#)

Procedure Overview

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

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 VxVM](#)” on page 41. For information on how to upgrade the Storage Administrator, see “[Upgrading VxVM](#)” on page 48.

Installation Overview

To install the Volume Manager and the Storage Administrator for the first time, use the following procedure:



1. Install the software packages onto the system to be administered as described in the section “[Upgrading VxVM](#)” on page 15.
2. Initialize VxVM using `vxinstall` as described in the section “[Upgrading VxVM](#)” on page 17.
3. Set up and start the Storage Administrator as described in “[Upgrading VxVM](#)” on page 30.
4. If you are planning to use the Volume Manager in a cluster environment, follow the instructions in “[Upgrading VxVM](#)” on page 38.

Loading the Software from CD-ROM

The VERITAS CD-ROM contains the following packages:

- ◆ `VRTSvxvm`—the Volume Manager and Volume Replicator software.
- ◆ `VRTSvmdoc`—online copies of VERITAS Volume Manager guides.
If you do not want to be able to read or print the documentation, omit the `VRTSvmdoc` package.
- ◆ `VRTSvrdoc`—online copies of VERITAS Volume Replicator guides.
- ◆ `VRTSvmmman`—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 “[Upgrading VxVM](#)” on page 16. Otherwise, mount the CD-ROM as described in the section “[Upgrading VxVM](#)” on page 16.



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, install the packages using the following command:

```
# pkgadd -d /cdrom/CD_name/OS_version/pkgsrc VRTSvxvm \
VRTSvmsa VRTSvmdoc VRTSvrdoc VRTSvmmman 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 takes 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, mount the CD-ROM manually, assuming your CD driver is on controller 0. To mount the CD-ROM manually, use the following command:

```
# 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, install the packages using the following command:

```
# pkgadd -d /mnt/OS_version/pkgsrc VRTSvxvm VRTSvmsa \
VRTSvmdoc VRTSvrdoc VRTSvmmman 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 takes about five minutes.

Disabling Hot-Relocation

The hot-relocation feature is enabled by default and it is recommended that you leave it on. However, you can disable it 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 &
```



Initializing the Volume Manager

This section describes how to initialize the Volume Manager using the `vxinstall` program. The `vxinstall` program places all specified disks under Volume Manager control. These disks are automatically placed in the default disk group, `rootdg`. To place some disks in another disk group, use the `vxinstall` program to initialize at least one disk into `rootdg`. Then use the Storage Administrator or the `vxdiskadm` program 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 the `vxinstall` program can be added later using another Volume Manager interface (such as the Storage Administrator or the `vxdiskadm` program). However, you must place at least one disk under Volume Manager control with the `vxinstall` program 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 using the following procedure:

1. Log in as superuser.
2. To exclude one or more disks from Volume Manager control, create the `/etc/vx/disks.exclude` file and add the names of those disks to the file. The `vxinstall` program ignores any disks listed in this file and leaves them untouched. The following is an example of the contents of the `disks.exclude` file:

```
c0t1d0
```

To exclude all disks on an entire controller, add 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. To start the `vxinstall` program, use the following command:

```
# vxinstall
```

The `vxinstall` program first runs the `vxlicense` command to initialize the Volume Manager license key file. The `vxlicense` command displays 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. The `vxinstall` program does not prompt for another key.

The `vxinstall` program then examines all controllers attached to the system and lists them, as follows:

```
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, the `vxinstall` program displays a brief introduction to the installation process, as follows:

```
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.

The vxinstall program 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:

The options are defined as follows:

- ◆ 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 the vxinstall program.
 - ◆ 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, see “[Upgrading VxVM](#)” on page 20.
 - ◆ If you choose Custom Installation, see “[Upgrading VxVM](#)” on page 24.



Quick Installation brings 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` files, described in [step 2](#) on page 17. To prevent disks on your system from being placed under Volume Manager control, either list them in the `disks.exclude` file or choose the Custom Installation option, which allows you to exclude certain disks.

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. To encapsulate some disks on a given controller and initialize others, use Custom Installation instead.

Note You cannot grow or shrink any volume associated with an encapsulated bootdisk (`rootvol`, `usr`, `var`, `opt`, `swapvol`, etc.) because these map to a physical underlying partition on the disk and must be contiguous. Refer to *Volume Manager Rootability* in Chapter 3 of the *VERITAS Volume Manager Administrator's Guide*.

Quick Installation allows you to initialize disks on a controller, even if valid partitions are found on those disks. However, 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 are mounted on volumes instead.

To use the Quick Installation option, use the following procedure:

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

The `vxinstall` program prompts for each controller and disk connected to your system. When a default response is displayed in parentheses, press Return to accept that default. At any of the `vxinstall` prompts, select `q` to completely abandon the initialization and start over.

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

2. Indicate whether to encapsulate your boot disk (in this example, the boot disk name is `c0t3d0`) at the following prompt:

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 are left unaltered.
- ◆ If you enter **y**, your boot disk is encapsulated. At the prompt, enter a disk name for the boot disk and press Return. The `vxinstall` program encapsulates your root file system as a volume, along with your swap device and all other disk partitions found on your boot disk. The `/usr`, `/opt`, and `/var` file systems, and any other file systems on your boot disk, are also encapsulated.

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

```
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. Indicate whether you want to initialize all of the disks on the controller at the following prompt:

```
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**, the `vxinstall` program 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**, the `vxinstall` program prompts you for confirmation and warns you that this destroys any existing data on the disks, as follows:

```
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 is displayed:

```
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. Indicate whether you want default disk names to be assigned to all of the disks on the controller at the following prompt:

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

- ◆ If you enter **y**, the `vxinstall` program automatically assigns and lists default disk names for each disk, as follows:

```
The c1t0d0 disk will be given disk name disk01
The c1t1d0 disk will be given disk name disk02
The c1t2d0 disk will be given disk name disk03
The c1t3d0 disk will be given disk name disk04
```

Hit RETURN to continue.

If you selected encapsulation earlier, this listing also includes a message for each disk that can be encapsulated, as follows:

```
The c1t0d0 disk has been configured for encapsulation.
```

or a message for empty disks to be initialized instead of encapsulated, as follows:

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



- ◆ If you enter **n**, the `vxinstall` program prompts you for a disk name for each disk on the controller individually, as follows:

```
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, the `vxinstall` program displays a summary of the disks designated for initialization (`New Disk`) or encapsulation (`Encapsulate`) on each controller, as follows:

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**, the `vxinstall` program encapsulates all disks listed with `Encapsulate` and initializes (destroying any existing data) all disks listed with `New Disk`.

If you enter **n**, the `vxinstall` program 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 again. The `vxinstall` program displays an updated summary without the disks chosen for removal. If you are satisfied with the summary, enter **y**.



Depending on your input during the `vxinstall` program, you may be requested to shut down and reboot your machine. If this happens, shut down and reboot your machine immediately or at the earliest convenient opportunity. For more information, see “[Upgrading VxVM](#)” on page 29.

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, the `vxinstall` program indicates successful completion of the installation, as follows:

```
Installation is complete.
```

Custom Installation

Custom Installation allows you to control which disks are placed under Volume Manager control and how they are added. This is important if you do not want all your disks placed under Volume Manager control. Custom Installation also provides 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.

Note You cannot grow or shrink any volume associated with an encapsulated bootdisk (`rootvol`, `usr`, `var`, `opt`, `swapvol`, etc.) because these map to a physical underlying partition on the disk and must be contiguous. Refer to *Volume Manager Rootability* in Chapter 3 of the *VERITAS Volume Manager Administrator's Guide*.

To use the Custom Installation option, perform the following procedure:

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

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

Note All disks are encapsulated or initialized at the end of the `vxinstall` procedure. If you quit `vxinstall` before it initializes or encapsulates the disks, all disks are left as they were before the `vxinstall` program was started.

2. Indicate whether to encapsulate your boot disk (in this example, your boot disk is named `c0t3d0`) at the following prompt:



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

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

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. Indicate how you want the disks on this controller to be handled at the following prompt:

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
q Exit from menus
```

Select an operation to perform:

Encapsulating All Disks

1. To encapsulate all disks on the controller, select **1**.

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

2. 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**, the `vxinstall` program automatically assigns and lists default disk names for each disk, as follows:

```
The c1t0d0 disk will be given disk name disk01
```

```
The c1t0d0 disk has been configured for encapsulation.
```

```
The c1t1d0 disk will be given disk name disk02
```

```
The c1t1d0 disk has been configured for encapsulation.
```

```
Hit RETURN to continue.
```

If any disks are found to be empty, they are initialized instead of encapsulated, as described in the following message:

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

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

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

For each disk, enter the desired disk name and press Return.

3. When all of the disks on the current controller are named, press Return to move on to the next controller.



Initializing All Disks

1. 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 are 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 the `vxinstall` program 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 Volume Manager specifications.

2. Indicate whether you want default disk names to be assigned to all of the disks on the controller at the following prompt:

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

- ◆ If you enter **y**, the `vxinstall` program automatically assigns and lists default disk names for each disk, as follows:

```
The c1t0d0 disk will be given disk name disk01
The c1t1d0 disk will be given disk name disk02
```

```
Hit RETURN to continue.
```

- ◆ If you enter **n**, the `vxinstall` program prompts you for a disk name for each disk on the controller, as follows:

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

For each disk, enter the desired disk name and press Return.

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

Installing Individual Disks

1. To install one disk at a time, select 3.

Each disk is handled separately and you are 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.

2. At the following prompt, indicate how you want this disk to be handled:



```
Installation options for disk c1t0d0
Menu: VolumeManager/Install/Custom/c1/c1t0d0

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 determined how to handle this disk, see [step 3](#) on page 25.

3. Press Return to move on to the next disk. Repeat the previous step for the remaining disks on the current controller.
4. When all of the disks on the current controller have been named, press Return to move on to the next controller.

Leaving All Disks Unaltered

1. To leave all disks on the controller unaltered, select 4.

No changes are made to the disks and they are not 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, use this option to ensure that your applications continue to use the disks without modification.

2. Repeat [step 3](#) on page 25 for each of the remaining controllers on your system.

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

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

```
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. The `vxinstall` program displays an updated summary without the disks chosen for removal. If you are satisfied with the summary, enter **y**.

Depending on your input during the `vxinstall` procedure, you may be requested to shut down and reboot your machine. If this happens, shut down and reboot your machine immediately or at the earliest convenient opportunity. For more information, see [“Upgrading VxVM”](#) on page 29.

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, the `vxinstall` program indicates successful completion of the installation:

```
Installation is complete.
```

Rebooting After vxinstall

After you have determined how the `vxinstall` program processes 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 is definitely necessary.

The `vxinstall` program informs you when a shutdown and reboot is necessary. Conduct the shutdown and reboot as soon as possible.



Note The setup you choose can require several reboots.

When it is necessary to shut down and reboot your system, the `vxinstall` program displays 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 message is displayed, respond in one of the following ways:

- ◆ Enter **y** to begin an immediate shutdown.
- ◆ Enter **n** to cause the `vxinstall` program to exit without starting a shutdown. If you select this option, it is your responsibility 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.

For information on how to upgrade from a previous release of the Storage Administrator, see “[Upgrading VxVM](#)” on page 48.

Note This release of the Storage Administrator requires Volume Manager Release 2.5.7, 3.0.4, 3.1 or higher. Before you install and run the Storage Administrator, you must install the Volume Manager and run the `vxinstall` program to initialize the Volume Manager.

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

Storage Administrator Installation Overview

Set up and start the Storage Administrator as follows:

1. Install the Storage Administrator server and client on the UNIX (Solaris) machine to be administered. See “[Upgrading VxVM](#)” on page 31.
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 “[Upgrading VxVM](#)” on page 34.
3. 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 “[Upgrading VxVM](#)” on page 31” or “[Upgrading VxVM](#)” on page 32.
4. Start the Storage Administrator server on the machine to be administered. See “[Upgrading VxVM](#)” on page 35.
5. Start the Storage Administrator client. See “[Upgrading VxVM](#)” on page 37.

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 “[Upgrading VxVM](#)” on page 15.

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



The Storage Administrator package for UNIX contains the following directories:

Directory	Contents
jre/	Java Runtime Environment (JRE)
bin/	vmsa_server—Storage Administrator server startup script vmsa—Storage Administrator startup script
vmsa/	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 using the following command:

```
$ VMSAHOME=Storage_Administrator_directory
$ export VMSAHOME
```

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 using the following command:

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

Installing the Storage Administrator Client (Windows)

The Volume Manager Storage Administrator client runs on Windows NT, Windows 2000, 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, ensure environment changes made during installation take effect by performing one of the following procedures:
 - ◆ For Windows NT or Windows 2000, 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 have access to the Storage Administrator. To set up a list of users who have 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 must include the user names of any users (including `root`) to 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 the `vrtsadm` group exists, `root` must be included in the `vrtsadm` group for `root` to access the Storage Administrator.

Changing the Group Name (Optional)

The `vrtsadm` group is the default VERITAS group name. However, you can change the `vrtsadm` group name to another name by setting the `vrts.server.adminGroup` property to another value. 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
```

Encryption

Data transferred between the VMSA client and server is encrypted by default. If necessary, disable encryption by setting the security property to `false` in the `/opt/VRTSvmsa/vmsa/properties` file, as follows:

```
vrts.security=false
```

Note Disabling encryption reduces the level of security and is not recommended.

Monitoring Access to the 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 1999: user xyz login succeeded
```



```
Fri Jan 23 10:59:52 PST 1999: 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.

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 must run continuously and must 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, as follows:

```
# cd $VM_SAHOME/bin
```

3. Start the server, as follows:

```
# ./vmsa_server &
```

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

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

To confirm that the Storage Administrator is running, use the following command:

```
# vmsa_server -q
```

To stop the Storage Administrator server, use the following command:

```
# vmsa_server -k
```

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

The VMSA server has an auto-start mode that starts the server on an as-needed basis. In auto-start mode, the VMSA server automatically starts whenever a VMSA client tries to connect to the server and the server is not already running. If no VMSA clients are connected for a certain amount of time (15 minutes, by default), the server automatically terminates and restarts at the next client connection.

To start the server in auto-start mode, type:

```
autostart on
```



To confirm that auto-start mode is enabled, run the following:

```
netstat -a | grep vmsa
```

The output should show that a VMSA service has a LISTEN state:

```
*.vmsa      *.*          0          0          0          0 LISTEN
```

Auto-start mode remains in effect across system reboots. To disable auto-start mode, type:

```
autostart off
```

For more information about auto-start mode, refer to the `autostart (1M)` manual page.

Starting the Storage Administrator Client

Note Only users with appropriate privileges can run the Storage Administrator. Refer to “[Upgrading VxVM](#)” on page 34 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 and the Storage Administrator server must be installed on the machine to be administered. The Volume Manager `vxconfigd` daemon 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, start the Storage Administrator in one of the following ways:

UNIX

To administer the *local* UNIX machine, use the following command:

```
# vmsa
```

To administer a *remote* UNIX machine, use the following command:

```
# vmsa remote_machine_name
```

Windows

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

Running the Storage Administrator in Read-Only Mode

The Storage Administrator can be run in a read-only mode that is useful for monitoring, training, or browsing purposes. Read-only mode allows you to view objects on the system through the Storage Administrator, but prevents administrative actions from taking effect. When the Storage Administrator is in read-only mode, the main window and dialog box title bars contain the following string:

```
[Read-only Mode]
```

The default mode for the Storage Administrator is full operational mode. Read-only mode can be enabled on a per-host or per-user basis.

To specify a list of users with permission to run the Storage Administrator in read-only mode, add a group named `vrtsro` to the `/etc/group` file or NIS (Network Information Name Service) group table on the machine to be administered. The `vrtsro` group must include the names of any users who are allowed to run the Storage Administrator in read-only mode, but are not allowed to run it in full operational mode.



Run the Storage Administrator in read-only mode using *one* of the following procedures:

- ◆ From the main window, choose Options > Read-Only Mode to turn read-only mode on (or off) for the current session.
- ◆ Start the server in read-only mode:

```
# vmsa_server -r
```

Read-only mode will remain in effect until you stop and restart the server. Use the `vmsa_server -k` command to stop the server.

- ◆ Start the client in read-only mode:

```
# vmsa -r
```

- ◆ Connect to a host machine as a user who only has permission to run the Storage Administrator in read-only mode on that host. This enables read-only mode on the specified host machine.

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.)

To enable the cluster functionality in the Volume Manager, use the following procedure

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 “[Upgrading VxVM](#)” on page 15.
3. Initialize VxVM using the `vxinstall` procedure as described in “[Upgrading VxVM](#)” on page 17.
4. Set up and start the Storage Administrator as described in “[Upgrading VxVM](#)” on page 30.
5. Configure shared disks as described in “[Upgrading VxVM](#)” on page 39.



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, verify that your shared disks still exist.

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 the shared disks are not being accessed from another node while you are performing the configuration. If you start the cluster on the node where you perform the configuration only, you can prevent disk accesses from other nodes because the quorum control reserves the disks for the single node (see “[Upgrading VxVM](#)” on page 41).

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 using the following procedure:

1. Start the cluster on at least one node.
2. Run the `vxdisksetup` program 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, the `vxdisksetup` program allows you to specify your choices.
3. Create disk groups on the shared disks.
Use the `vxdbg` program or the Storage Administrator to do this. Use the `-s` option of the `vxdbg` program to create shared disk groups.
4. Create volumes in the disk groups.
Use the `vxassist` program 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 “Dirty Region Logging” in *Volume Manager Cluster Functionality*.
5. If the cluster is only running with one node, bring up the other cluster nodes. Enter the `vxdbg list` command on each node to 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 using the following procedure:

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

```
# vxdg list
```

This displays the shared disk groups that existed earlier.

Converting Existing VxVM Disks to Shared Disks

If you are upgrading from VxVM 2.x to VxVM 3.x and you want to convert existing disk groups to shared disk groups, configure the shared disks using the following procedure:

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. Configure the disk groups using the following procedure.

To list all disk groups, use the following command:

```
# vxdg list
```

To deport disk groups to be shared, use the following command:

```
# vxdg deport groupname
```

To import disk groups to be shared, use the following command:

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

This procedure marks the disks in the shared disk groups as shared and stamps them with the ID of the cluster, enabling other nodes to recognize the shared disks.

If dirty region logs exist, ensure they are active. If not, replace them with larger ones.

To display the shared flag for all the shared disk groups, use the following command:

```
# 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. Enter the `vxdg list` command on each node to display the shared disk groups. This command displays the same list of shared disk groups displayed 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 `vxdisk list` command used 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 the `vxdisk` command 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 can be invisible to the booting system, and the `vxdisk` command may not display any of the shared disks. When the system joins the cluster, the shared disks become visible.

Upgrading to VxVM Release 3.1

If you are already running an earlier release of the Volume Manager, you can upgrade to VxVM Release 3.1 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. No new license is required for upgrade.

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

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 “Disk Tasks” in the *VERITAS Administrator’s Guide* for information on how to upgrade a disk group.



Table 1 indicates 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 3.1					Supported	Supported	Supported
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			
VxVM 2.2		Supported	Supported				
VxVM 2.1	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 2.0	Supported						

Table 2 shows 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:	“Upgrading VxVM and Solaris” on page 44
VxVM 2.1 & Solaris 2.4 ► VxVM 3.1 & Solaris 2.6/7/8	
VxVM 2.2 & Solaris 2.4/2.5 ► VxVM 3.1 & Solaris 2.6/7/8	
VxVM 2.2.1 & Solaris 2.4/2.5/2.5.1 ► VxVM 3.1 & Solaris 2.6/7/8	
VxVM 2.3 & Solaris 2.4/2.5/2.5.1 ► VxVM 3.1 & Solaris 2.6/7/8	
VxVM 2.4/2.4.1 & Solaris 2.4/2.5/2.5.1 ► VxVM 3.1 & Solaris 2.6/7/8	
VxVM 2.5/2.5.x & Solaris 2.5/2.5.1/2.6 ► VxVM 3.1 & Solaris 2.6/7/8	
VxVM 3.0 & Solaris 2.5.1/2.6 ► VxVM 3.1 & Solaris 2.6/7/8	
VxVM 3.0.x & Solaris 2.5.1/2.6/7 ► VxVM 3.1 & Solaris 2.6/7/8	
VxVM Only:	“Upgrading VxVM” on page 44
VxVM 2.2.1 & Solaris 2.5.1 ► VxVM 3.1 & Solaris 2.5.1	
VxVM 2.3 & Solaris 2.5.1 ► VxVM 3.1 & Solaris 2.5.1	
VxVM 2.4 & Solaris 2.5.1 ► VxVM 3.1 & Solaris 2.5.1	
VxVM 2.4.1 & Solaris 2.5.1/2.6 ► VxVM 3.1 & Solaris 2.6	
VxVM 2.5/2.5.x & Solaris 2.5.1/2.6 ► VxVM 3.1 & Solaris 2.6	
VxVM 3.0 & Solaris 2.5.1/2.6 ► VxVM 3.1 & Solaris 2.6	
VxVM 3.0.x & Solaris 2.5.1/2.6/7/8 ► VxVM 3.1 & Solaris 2.6/7	
Solaris Only:	“Upgrading Solaris Only” on page 46
VxVM 3.1 & Solaris 2.5.1 ► VxVM 3.1 & Solaris 2.6/7/8	
VxVM 3.1 & Solaris 2.6 ► VxVM 3.1 & Solaris 7/8	
VxVM 3.1 & Solaris 7 ► VxVM 3.1 & Solaris 8	



Upgrading VxVM and Solaris

This section provides detailed instructions for upgrading both Solaris and VxVM. Use [Table 1](#) on page 42 as a reference to determine if an upgrade of the operating system is required in order to upgrade to VxVM 3.1.

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 3.1 license key. See “[Upgrading VxVM](#)” on page 11 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 “[Upgrading VxVM](#)” on page 15. The CD-ROM should then be visible as the file system mounted on `/CD_path`. If the CD-ROM is mounted automatically, its path is `/cdrom/CD_name`; if mounted manually, the path is `/mnt`.
4. Run the `upgrade_start` script to prepare the previous release of the Volume Manager for its removal using the following command:

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

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

5. Reboot to single-user mode (using a command such as `/etc/shutdown`).
6. If VxVM is installed, remove the old VxVM packages by entering:


```
# pkgrm VRTSvxvm VRTSvmsa VRTSvmdoc VRTSvmdev
```
7. If Sun Volume Manager is installed:



- ◆ Remove any `SUNWvxvm` related patches.
Refer to the SUN documentation for the procedure.
- ◆ Remove `SUNWvxvm`:

```
# pkgrm SUNWvxvm SUNWvmdoc SUNWvmdev
```


Run from the directory where the `VRTSvxvm` exists.

Note If you have made any additional modifications or have applied any patches, they must be removed before removing the `SUNWvxvm` package.

8. Shut down and halt the machine by using a command such as `/etc/shutdown`.
9. Reboot the machine to single-user mode using a command such as `/etc/shutdown`.
Remount the CD-ROM manually using the following command:

```
# mount -F hsfs -o ro /dev/c0t6d0s2 /mnt
```
10. Add the VxVM 3.0.x package using the following command:

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


If warnings are displayed that include the string `/etc/vx`, ignore them and continue.
11. Complete the upgrade using the following command:

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


If you are upgrading from VxVM 1.3, this command prompts you for licensing information.
12. To disable the hot-relocation feature, comment out the `vxrelocd` line in the startup file `/etc/rc2.d/S95vxvm-recover` using the following command:

```
#vxrelocd root &
```
13. Reboot to multi-user mode. At this point, your pre-upgrade configuration is in effect and any file systems previously defined on volumes are defined and mounted.
14. Using the `pkgadd` command, 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 “Disk Tasks” in the *VERITAS Administrator’s Guide* for information on how to upgrade a disk group.



16. See “[Upgrading VxVM](#)” on page 30 to set up and restart the Storage Administrator.

Upgrading VxVM Only

If you are already running a version of Solaris that is supported with VxVM 3.1, you can upgrade VxVM without upgrading Solaris. Upgrade VxVM only as follows:

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


```
# pkgrm VRTSvmsa
```
2. Add the VxVM 3.1 package using the following command:


```
# pkgadd -d CD_path/OS_version/pkg VRTSvxvm
```
3. Using the `pkgadd` command, install any or all of the additional packages `VRTSvmsa`, `VRTSvmdoc`, `VRTSvmman`, and `VRTSvmdev`.
4. See “*Setting Up and Starting the Storage Administrator*” on page 30, to set up and restart the Storage Administrator.

Note The Storage Administrator Release 3.x server is not backward compatible with Storage Administrator Release 1.x clients. Upgrade any existing client(s) to Release 3.1.

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`. You are informed that a reboot is required.

To disable hot-relocation, comment out the `vxrelocd` line 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`, using the following command:

```
#vxrelocd root &
```

After editing this file, the system must be rebooted:

```
# reboot -r
```

Upgrading Solaris Only

If you are already running VxVM 3.1 with Solaris 2.6/7/8, you can upgrade to Solaris 2.6/7/8 using the following procedure:

If the boot disk is mirrored, detach the mirror.

1. If any of the file systems `/`, `/usr`, `/var`, or `/opt` are defined on volumes, ensure 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), an explanation of the problem is displayed and the upgrade does not proceed.
2. Run the `upgrade_start` script to prepare the previous release of the Volume Manager for its removal using the following command:

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

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

3. Shut down and halt the machine using a command such as the `/etc/shutdown` command.
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 the `/etc/shutdown` command.
6. Complete the upgrade using the following command:

```
# CD_path/scripts/upgrade_finish
```
7. Reboot to multi-user mode using a command such as the `/etc/shutdown` command.

Upgrading the Storage Administrator

If a previous release of the Storage Administrator is already installed, you must remove the previous `VRTSvmsa` package and then add the current `VRTSvmsa` package. To upgrade the Storage Administrator, use the following procedure.

1. Remove the previous `VRTSvmsa` package using the following command:

```
# pkgrm VRTSvmsa
```
2. See “[Upgrading VxVM](#)” on page 15 for instructions on adding the current `VRTSvmsa` package.

Note The Storage Administrator Release 3.x server is not backward compatible with Storage Administrator Release 1.x clients; therefore, upgrade any existing client(s) to Release 3.1.

Checking Volume Manager Processes

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

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

Output similar to the following is displayed:

```
root4581 80 10:23:34?0:00vxrelocd root
root4141 80 10:23:06?0:03vxconfigd
root47845839 10:23:37?0:00vxrelocd root
root479478 5 10:23:37?0:00vxnotify -f -w 15
```

Note If you have disabled hot-relocation, only the `vxconfigd` process is running.



Deinstalling the Volume Manager

To deinstall the Volume Manager when `root`, `swap`, `usr`, or `var` is a volume under Volume Manager control, use the following procedure:

1. Ensure that the `rootvol`, `swapvol`, `usr`, and `var` volumes have only one associated plex each. The plex must be contiguous, non-striped, non-spanned, and non-sparse. For information about the plex, use the following command:

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

If any of these volumes have more than one associated plex, remove the unnecessary plexes using the following command:

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

It is important that the plex remaining for the `rootvol` volume has been created using the `vxrootmir` program, which guarantees that underlying subdisks start on cylinder boundaries and that partitions are created for them.

2. Run the `vxunroot` program using the following command:

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

This command changes the volume entries in `/etc/vfstab` to the underlying disk partitions for the `rootvol`, `swapvol`, `usr`, and `var` volumes. The command also modifies `/etc/system` and prompts for a reboot so that disk partitions are mounted instead of volumes for the `root`, `swap`, `usr`, and `var` volumes.

Once the `root`, `swap`, `usr`, and `var` volumes have been completed, all remaining volumes must be moved to disk partitions. This can be done using one of three procedures, as follows:

- ◆ Back up the system fully onto tape and then recover from it.
- ◆ Back up each file system individually and then recover them all after creating new file systems on disk partitions.
- ◆ Move volumes incrementally onto disk partitions, as follows:
 - a. Evacuate as many disks as possible using the `vxdiskadm` program, the Storage Administrator, or the `vxevac` script.

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

- b. Remove the evacuated disks from Volume Manager control using the following commands:

```
# vxdg rmdisk diskname  
# vxdisk rm devname
```
- c. Decide which volume to move first. If the volume to be moved is mounted, unmount it.
- d. 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 synced.
- e. 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 for the first volume removed. Subsequent volumes can use the free space generated by the removal of this volume.
- f. Copy 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.
- g. Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly created partition.
- h. Mount the disk partition if the corresponding volume was previously mounted.
- i. Remove the volume from the Volume Manager using the following command:

```
# vxedit -rf rm volume_name
```
- j. 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 following command:

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

If the output is not 0, there are still some subdisks on this disk that must be subsequently removed. If the output is 0, remove the disk from Volume Manager control using the following commands:

```
# vxdg rmdisk diskname  
# vxdisk rm devname
```



- k. The free space now created can be used for adding the data in the next volume to be removed.
- l. After all volumes have been converted into disk partitions successfully, reboot the system. After the reboot, none of the volumes should be open. To verify that none of the volumes are open, use the following 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, use the following commands:

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

Removing the Volume Manager Package

To remove the installed Volume Manager packages, use the following command:

```
# 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. To determine whether both packages are present, use the following command:

```
# pkginfo | grep vxvm
```

Output is displayed that lists both the `VRTSvxvm` and `VRTSvxvm.2` packages. To remove both packages, use the following command:

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