VERITAS Volume Manager™ 3.2

Installation Guide

Solaris

VERITAS

Contents

Preface

Introduction

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

Audience

This guide is for system administrators who configure and maintain UNIX systems with VxVM, and assumes that you have:

◆ An understanding of system administration.

◆ A working knowledge of the UNIX operating system.

◆ A general understanding of file systems.

Organization

This guide is organized with the following chapters:

◆ Introduction

◆ Installing VERITAS Volume Manager

Conventions

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

Typeface	Usage	Examples
monospace	Computer output, file contents, files, directories, software elements such as command options, function names, and parameters	Read tunables from the `/etc/vx/tunefstab` file. See the `ls`(1) manual page for more information.
italic	New terms, book titles, emphasis, variables to be replaced by a name or value	See the *User's Guide* for details. The variable *ncsize* determines the value of...
monospace (**bold**)	User input; the "#" symbol indicates a command prompt	#**mount -F vxfs /h/filesys**
monospace (***bold and italic***)	Variables to be replaced by a name or value in user input	#**mount -F *fstype mount_point***

Symbol	Usage	Examples
%	C shell prompt	
$	Bourne/Korn/Bash shell prompt	
#	Superuser prompt (all shells)	
\	Continued input on the following line	#**mount -F vxfs ** **/h/filesys**
[]	In a command synopsis, brackets indicates an optional argument	`ls [-a]`
|	In a command synopsis, a vertical bar separates mutually exclusive arguments	`mount [suid | nosuid ]`

Getting Help

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

◆ U.S. and Canadian Customers: 1-800-342-0652

◆ International Customers: +1 (650) 527-8555

◆ Email: support@veritas.com

For license information (U.S. and Canadian Customers):

◆ Phone: 1-925-931-2464

◆ Email: license@veritas.com

◆ Fax: 1-925-931-2487

For software updates:

◆ Email: swupdate@veritas.com

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

◆ http://support.veritas.com (U.S. and Canada Customers)

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

◆ http://www.veritas.com

Downloading and Running VRTSexplorer

If you have access to the Internet, you can use the VRTSexplorer program to assist Technical Support in diagnosing the cause of your problem as follows:

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

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

Save the file to a temporary directory such as /tmp as shown in these instructions. If you download the file to a different directory, substitute its pathname for /tmp throughout.

2. Log in as root on the affected system, and use the following commands to extract the contents of the downloaded file to the directory /tmp/VRTSexplorer:

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

3. Run the `VRTSexplorer` program located in the VRTSexplorer directory by entering the following command:

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

4. When `VRTSexplorer` prompts you for a destination directory for the information that it collects, press Return to accept the default directory `/tmp`, or enter a pathname of your own choice. `VRTSexplorer` writes the results of its investigations to a compressed tar file named `VRTSexplorer.`*`case_number`*`.tar.Z` in the specified directory.

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

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

6. Call VERITAS Customer Support on 1-800-342-0652, inform them that you have run `VRTSexplorer` and tell them the name of the file that you transferred to the FTP site.

    Alternatively, if you have already been assigned a call ID number by Customer Support, e-mail `support@veritas.com` including your case ID number in the subject line.

For more information about the `VRTSexplorer` program, consult the `README` file located in the `VRTSexplorer` directory.

Introduction 1

This manual explains how to install (or initialize) VERITAS Volume Manager (VxVM). Although terms and basic concepts are introduced in this manual, refer to the *VERITAS Volume Manager Administrator's Guide* for more details.

After completing installation, there will be three utilities you can use: the Command Line Interface (CLI), the vxdiskadm menu, and the Volume Manager Storage Administrator (VMSA) interface. These utilities are explained in more detail in the *VERITAS Volume Manager Administrator's Guide.*

Where to Go Next

If you are a new VERITAS user, go to "VxVM Basics".

If you have used VERITAS Volume Manager products before, refer to "Pre-installation Instructions" on page 5.

If you are upgrading VxVM, go to "Upgrading to VxVM Release 3.2" on page 52.

VxVM Basics

This section introduces the basic features of VxVM. If you are a new user, read the following brief descriptions before proceeding. Details about the features are described in the *VERITAS Volume Manager Administrator's Guide.*

Easier Storage Management

VxVM is a storage management subsystem that allows you to manage physical disks as logical devices, which do not have the limitations that physical disks do.

VxVM enhances data storage management by controlling these aspects:

◆ space—allocation and use

◆ performance— enhanced data delivery

◆ data availability—continuous operation and multisystem access

◆ device installation—centralized and optimized support

◆ system—multisystem support and monitoring of private/shared systems

Physical Disks

A *physical disk* is the basic storage device (media) where the data is ultimately stored. A physical disk can be divided into one or more *partitions*, also known as *slices*.

When placed under VxVM control, a physical disk transforms to a virtual disk, or a volume that can be easier to monitor.

The VxVM Connection

The connection between physical objects and VxVM objects is made when you place a physical disk under VxVM control.

VxVM creates *virtual objects* and makes logical connections between the objects. The virtual objects are then used by VxVM to perform storage management tasks.

VxVM objects include the following:

◆ VM Disks

◆ Disk Groups

◆ Subdisks

◆ Plexes

◆ Volumes

VM Disks

VM Disk Example
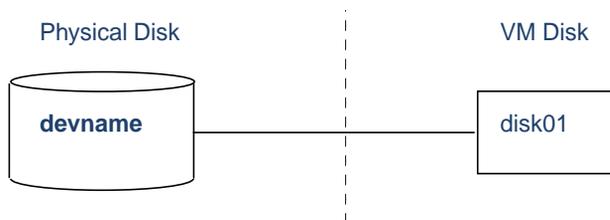
When active, the VM disk is associated with an underlying physical disk. VxVM logical objects map to physical objects and store data on stable storage.

A VM disk typically includes a *public region* and a *private region.* The public region on a physical disk is a region managed by VxVM and contains available space that is used for allocating subdisks. The private region contains VxVM internal configuration information.

Disk Groups

A VxVM volume is a logical *object.* VxVM creates other objects that you can operate, control, monitor, and query to optimize storage management.

To configure and maintain a volume for use, VxVM places physical disks under its control and collects the disk space into *disk groups.* A disk group is a collection of claimed disks organized into logical volumes. VxVM then allocates the space on those disks to logical volumes.

A *disk group* is a collection of VxVM disks that share a common configuration. A disk group configuration is a set of records with detailed information about related VxVM objects, their attributes, and their connections. The default disk group is rootdg (the root disk group).

You can create additional disk groups as necessary. Disk groups allow you to group disks into logical collections. A disk group and its components can be moved as a unit from one host machine to another.

Volumes are created within a disk group. A given volume must be configured from disks in the same disk group.

Subdisks

Subdisk Example

VM Disk with One Subdisk   Subdisk

```
disk01-01                      disk01-01
disk01
```
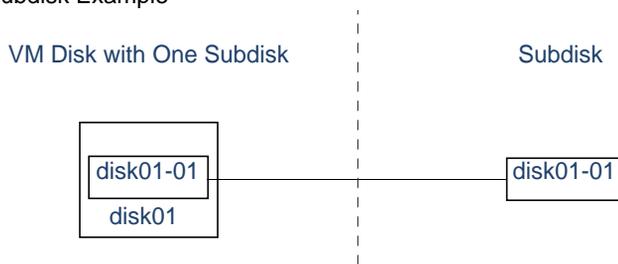
Subdisks are the smallest unit of storage in VxVM. They are like partitions on a physical disk and are used to build plexes.

Plexes

Plex with Two Subdisks
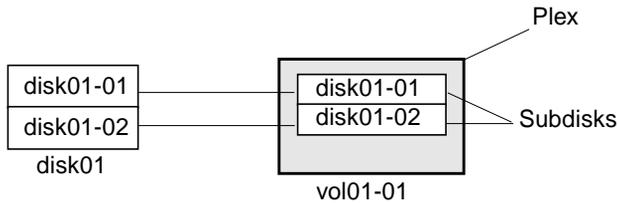
VxVM uses subdisks to build virtual objects called *plexes* (or mirrors). A plex consists of one or more subdisks located on one or more physical disks.

Volumes

Volume with One Plex
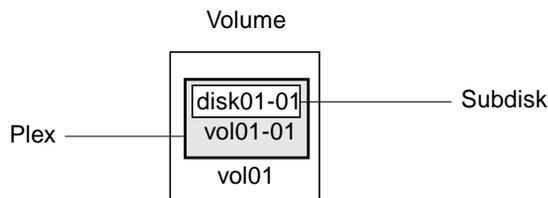
VxVM allows you to manage physical disks as logical devices called *volumes.* A volume is a logical device that appears to applications, databases, and file systems like a physical disk device, but does not have the physical limitations of a physical disk device. A volume is not restricted to a particular disk or a specific area of a disk.

Storage Layouts

Different volume layouts each provide different levels of storage service. A volume layout can be configured and reconfigured to match particular levels of desired storage service.

◆   Concatenation and spanning

◆   Striping (RAID-0)

◆   Mirroring (RAID-1)

◆   Mirroring plus striping

- ◆ Striping plus mirroring
- ◆ RAID-5 (striping with parity)

Arrays

A *disk array* is a collection of disks that appears to the system as one or more volumes.

Data is spread across several disks within an array, which allows the disks to share I/O operations. The use of multiple disks for I/O improves I/O performance by increasing the data transfer speed and the overall throughput for the array.

Pre-installation Instructions

Verify Disk Space Requirements

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

Package	Contents	Size
VRTSvxvm	Driver and Utilities	59 MB
VRTSlic	Licensing Utilities	2MB
VRTSvmdev	Developer Kit	2 MB
VRTSvmman	Manual Pages	1 MB
VRTSvmdoc	Documentation	30 MB
VRTSvmsa	Storage Administrator Server & Client (UNIX)	33 MB

The following table shows the approximate breakdown of the space requirements for each package against each file system (assuming the packages are installed in their default locations).

Set the Path Variable

Package	/	/usr	/opt
VRTSvxvm	27 MB	32 MB	0
VRTSlic	0	1MB	1MB
VRTSvmdev	0	0	2 MB
VRTSvmman	0	0	1 MB
VRTSvmdoc	0	0	30 MB
VRTSvmsa	0	0	33 MB

Most of the commands used in the installation are in the /usr/sbin directory or /sbin directory. You should add these directories to your *PATH* environment variable, making sure that /usr/sbin precedes /sbin in the path.

Similarly, set the path of your manpage environment variable to include /opt/VRTS/man.

Examples:

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

```
$ PATH=/usr/sbin:/sbin:/etc/vx/bin:/usr/lib/vxvm/bin:$PATH; export
PATH
$ MANPATH=/opt/VRTS/man:$MANPATH; export MANPATH
```

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

```
% setenv PATH /usr/sbin:/sbin:/etc/vx/bin:/usr/lib/vxvm/bin:$PATH
% setenv MANPATH /opt/VRTS/man:$MANPATH
```

Verify Directories and Permissions

If any of the documentation, developer kit, or manual page packages are to be installed, the directory /opt must exist and be writable. Also, /opt must not be a symbolic link. This is because the volumes not temporarily converted by the upgrade_start will be unavailable during the upgrade process. If you have a symbolic link from /opt to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in /opt will not be installed.

Obtain a License Key

```
hostid uname -i
```

A License Key Request Form is included in your product package. Fill in the form and fax it to VERITAS Customer Support at 1-925-931-2487. If you are upgrading VxVM, you can view the currently installed licenses with the following command:

```
vxlicense -p
/usr/lib/vxvm/bin/vxliccheck -pv
```

Add new licenses with the following command:

```
vxlicense -c
```

Take a Disk Inventory

Disk Types

Decide which disks you want to place under VxVM control. The other disks in your configuration will not be affected. There are two types of disks:

◆ Encapsulated disks—Data in all existing file systems and partitions on the disk are preserved while running the vxinstall program.

◆ Initialized disks—Data on the disks are removed while running the vxinstall program.

Boot Disk

You cannot grow or shrink any volume associated with an encapsulated boot disk (rootvol, usr, var, opt, swapvol, etc.) because these map to a physical underlying partition on the disk and must be contiguous.

Verify Disk Contents

Verify the disk contents. Answer the following questions and list the data for your convenience.

1. Make sure you are aware of the contents of each disk. Determine which disks can be encapsulated (data is preserved) or initialized (data is removed).

2. Do you want to place the system root disk under VxVM control? (This is necessary in the case of a one-disk system or when you wish to mirror the root disk.)
   _____

3. Do you want to either encapsulate or initialize *all* disks on a controller together? Identify the controllers.
   _____
   _____
   _____

4. Identify the disks to encapsulate, initialize, or to be excluded in the following table.

Disk ID	Encapsulate, Initialize, Exclude

5. Verify the disks to be managed by VxVM have two free partitions and a small amount of free space (2048 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 VxVM 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. VxVM allocates approximately 2048 sectors 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 disks.

The boot disk is a special case. If no other space is available, VxVM 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 VxVM 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.

6. If you are encapsulating the boot (root) disk:

    a. eeprom "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.

7. If you have disks that will be encapsulated, they must meet the following criteria:

    - contain the required minimum free space (2048 sectors)

- contain two free partitions

8. DMP considerations:

- If your system uses Sun's AP software, this release of VxVM requires an AP upgrade and a Solaris patch. See "DMP and Alternate Pathing" on page 11.

- If you are upgrading from VxVM 3.1, DMP is automatically enabled in VxVM 3.2 You can use vxinstall option #3 to prevent DMP if you choose.

- If you are upgrading from VxVM 3.1.1, and you had prevented DMP, VxVM 3.2 retains your DMP setting. Your new installation will have DMP prevented. Use vxinstall option #3 to enable DMP if you chose.

Hot-Relocation

Hot-relocation automatically restores redundancy and access to mirrored and RAID-5 volumes when a disk fails. This is done by relocating the affected subdisks to disks designated as spares and/or free space in the same disk group.

The hot-relocation feature is enabled by default. The associated daemon, vxrelocd, is automatically started during system startup.

Follow these recommendations:

1. Leave the VxVM hot-relocation feature enabled to detect disk failures automatically. It will notify you of the nature of the failure, attempt to relocate any affected subdisks that are redundant, and initiate recovery procedures.

2. Configure at least one hot-relocation spare disk in each disk group. This will allow sufficient space for relocation in the event of a failure.

If you decide to disable hot-relocation, prevent vxrelocd from running after you load the VxVM software. See "Upgrading VxVM on an Encapsulated Root Disk" on page 56 for details.

Solaris Patches for VxVM

VxVM requires Solaris patches to work properly. See "Important Software Patch Information" on page 16 for details about required patches.

VMSA

For a machine to be administered by VMSA, the following software must be installed and running:

◆ Solaris 2.6 or higher.

◆ VERITAS Volume Manager Release (2.5.7, 3.0.4, 3.1, 3.1.1, or 3.2). The `vxconfigd` and `VMSA_server` daemons must be running on the machine to be administered.

◆ The VMSA 3.2 server. See "Setting Up and Starting VMSA" on page 42 on how to install the server.

◆ The VMSA 3.2 GUI client is not compatible with earlier versions of the VMSA server. When installing VMSA 3.2, both the VMSA server and VMSA GUI client must be updated.

◆ If an earlier VERITAS Volume Manager Storage Administrator package is installed on the machine, remove the old VMSA version before you install the new VMSA version. See "Upgrading to VxVM Release 3.2" on page 52 for details.

A5x00 Devices

Whether or not you have A5X00 devices connected to your system, you must add the Solaris patches shown in "Solaris Patches for VxVM" on page 16 to this release of VxVM.

DMP and Alternate Pathing

VxVM 3.2 allows Dynamic Multipathing (DMP) to co-exist with Sun's Alternate Pathing (AP) software. If your system uses Sun's AP software, this release of VxVM requires AP version 2.3.1 and Sun patch 110722 (latest version). Upgrade AP to 2.3.1 as described in "Installing or Upgrading VxVM with Alternate Pathing Enabled" on page 17. Then install the patch before installing or upgrading to this release of VxVM. See "Solaris Patches for VxVM" on page 16 for patch information.

DMP is described in more detail in the *VERITAS Volume Manager Administrator's Guide*.

Use the new interfaces provided through the utilities `vxinstall` and `vxdiskadm` to prevent DMP from multipathing certain or all devices connected to the system. These interfaces can also be used to suppress devices from VxVM.

Claiming Devices in the SENA Category

The A5x00 disk arrays will be claimed by DMP under the SENA category only if the required libraries are present on the system at the time of installation/upgrade of VxVM.

These libraries are present by default on Solaris 8. For Solaris 2.6 and Solaris 7, see "Solaris Patches for VxVM" to find the appropriate patch. These patches must be installed before installing or upgrading VxVM.

Cluster Environment Requirements

If your configuration has a cluster, which is a set of hosts that share a set of disks, follow these steps:

1. Obtain a license for the *optional* VxVM cluster feature from your Customer Support channel.

2. Decide where to place the rootdg for each node in the cluster. A system using VxVM has one or more disk groups, including the rootdg. The rootdg must exist and cannot be shared between systems. At least one disk must exist within the rootdg while VxVM is running.

3. Decide on the layout of shared disk groups. There may be one or more shared disk groups. Determine how many you wish to use.

4. If you plan to use Dirty Region Logging (DRL) with VxVM in a cluster, leave 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.

5. Refer to the *VERITAS Volume Manager Administrator's Guide* for more information on DRL.

Set up your SAN Configuration

If you have a SAN configuration, follow these general guidelines before you install VxVM.

◆ You must place at least one disk into the rootdg.

◆ If some disks in your SAN will not be under VxVM control, create a file called `/etc/vx/cntrls.exclude` or `/etc/vx/disks.exclude`. (See page 21 for more information.)

Adding Disks After Installation

Disks that are not initially placed under VxVM control by the `vxinstall` program can be added later using another VxVM interface (such as the VMSA or the `vxdiskadm` program).

Installing VERITAS Volume Manager 2

This section describes how to install and set up VERITAS Volume Manager (VxVM) and VERITAS Volume Manager Storage Administrator (VMSA). This also describes how to upgrade an existing version of VxVM and VMSA. Installation or upgrade and setup consists of the following procedures:

◆ VxVM installation - see "Installation and Upgrade Overview" on page 15

◆ VxVM upgrade - see "Upgrading to VxVM Release 3.2"

◆ VxVM configuration and setup

◆ VMSA setup

◆ Shared disk setup (for cluster environments only)

◆ Removing VxVM

Note Only users with superuser privileges can install and initialize VxVM and VMSA.

Installation and Upgrade Overview

1. For installations and upgrades, install the required patches. Refer to "Important Software Patch Information" on page 16.

2. For installations and upgrades, you may need to upgrade Sun's Alternate Pathing (AP) driver to version 2.3.1. Refer to "Installing or Upgrading VxVM with Alternate Pathing Enabled" on page 17.

3. For new installations, load the software packages onto the system to be administered as described in the section "Loading the Software from CD-ROM" on page 18.

4. For upgrades, you need to determine if your root disk is encapsulated, as this determines which upgrade process you should use. See "Upgrading to VxVM Release 3.2" for more information.

5. For new installations, install VxVM as described in "Running the vxinstall Program" on page 20.

6. Set up and start VMSA as described in "Setting Up and Starting VMSA" on page 42.

7. If you are planning to use VxVM in a cluster environment, follow the instructions in "Enabling Cluster Support in VxVM (Optional)" on page 49.

8. After installation, add any disk arrays that are unsupported by VERITAS to the JBOD category as described in the section "Hot-Relocation Guidelines" on page 70.

Important Software Patch Information

Caution	Please read this section completely *before you install* VERITAS Volume Manager (VxVM) 3.2. You must reboot your system after installing these patches. If you have already installed VxVM, install the patches and then reboot your system.

Solaris Patches for VxVM

Caution	Do not install VxVM until you have installed the following mandatory patches. The patches must be installed on your system for VxVM to work properly.

Some required system patches may already be present in your operating system. You should check to see if your system already contains the patches needed. Use the command showrev -p to display the patches included in your system. For more information, see showrev1M).

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

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

Required Solaris Patches

For Solaris Version	Sun Patch ID Number	Notes
Solaris 8	SUNWsan package (see your vendor to obtain) 109529 (-06 or later) 111413 (-02 or later)	Regardless of whether you have an A5k drive connected, you must install these A5k patches for VxVM to function. You must install the SUNWsan package and patch 109529 before installing patch 111413. You can then install VxVM. This order of these steps is important.
Solaris 7 Solaris 2.6	107473 (-07 or later) 105375 (-20 or later) You do not need to install patch #105357: Ignore the install script message.	Regardless of whether you have an A5k drive connected, you must install an A5k patch for VxVM to function. You should install this patch before installing VxVM.
Solaris 8, 7, 2.6	110722 (-01 or later)	If you have Sun's Alternate Pathing (AP), you must install Solaris patch 110722-01 for VxVM 3.2 to co-exist with the AP driver version 2.3.1. However, you must also upgrade AP. Please see "DMP Coexistence with Alternate Pathing" on page 10 before installing VxVM.
Solaris 7	106541 (-16 or later)	If you are running Solaris 7, you must install Sun patch ID 106541.

Installing or Upgrading VxVM with Alternate Pathing Enabled

Dynamic Multipathing (DMP) on VxVM 3.2 co-exists with Sun's Alternate Pathing (AP) driver 2.3.1. For VxVM 3.2 to function, the DMP driver (`vxdmp`) must always be present on the system. Upgrading to this version of VxVM automatically enables the DMP driver.

If you are do not know which version of AP is installed, run the following command:

```
# pkginfo -l SUNWapdv | grep VERSION
```

For VxVM to co-exist with AP successfully, you must upgrade AP before installing or upgrading to VxVM 3.2.

If AP is installed and the AP version is prior to 2.3.1, installation is aborted and you are instructed to upgrade your AP package to version 2.3.1.

Loading the Software from CD-ROM

The Packages

The VERITAS packages are found on the following CD-ROM.

CD-ROM Contents

◆ `VRTSvxvm`—VERITAS Volume Manager and Volume Replicator software.

◆ `VRTSlic`—VERITAS Volume Manager Licensing Utilities.

◆ `VRTSvmdoc`—online copies of VERITAS Volume Manager guides.

◆ `VRTSvrdoc`—online copies of VERITAS Volume Replicator guides.

◆ `VRTSvmman`—VxVM manual pages.

◆ `VRTSvmdev`—(optional) a VxVM developer's package (library and header files).

◆ `VRTSvmsa`—VERITAS Volume Manager Storage Administrator software.

◆ `VRTSvras`— VERITAS Volume Replicator Administrative Services software.

Note The VxVM and VVR Release Notes are also contained on the CD-ROM.

Mounting the CD-ROM with Solaris Volume Management

Follow these steps. When a default response is displayed in parentheses, press **Return** to accept that default.

1. Log in as superuser.

2. Insert the CD-ROM into a CD-ROM drive connected to your system.

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

3. After the CD-ROM is mounted, install the packages. Enter:

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

Note The order in which the packages are listed is significant. The VRTSlic package must be listed first, followed by the VRTSvxvm package, then the remaining packages.

Removing Previous Licensing Packages

If you added the VRTSlic and see the following message, you need to manually de-install the old package first.

```
VERITAS Licensing Utilities
(sparc) 1.0.1,REV=05.03.2001.19.22
Copyright (c) 1991 – 2000 VERITAS SOFTWARE CORP. ALL RIGHTS
RESERVED. THIS SOFTWARE IS THE PROPERTY OF AND IS LICENSED BY
VERITAS SOFTWARE, AND/OR ITS SUPPLIERS.

Current administration requires that a unique instance of the
<VRTSlic> package be created.  However, the maximum number of
instances of the package which may be supported at one time on the
same system has already been met.
No changes were made to the system.
```

1. To de-install the old package, enter pkgrm VRTSlic.

2. During de-install, if you see the following message:

```
WARNING: The VRTSvxfs package depends on the package currently
being removed. Dependency checking failed.
Do you want to continue with the removal of this package [y,n,?,q]
```

3. Enter y. You can now proceed to install the latest VRTSlic.

Note VRTSvxfs does not lose its package dependency when the latest VRTSlic is installed later in this process.

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.

If the installed version of AP is earlier than 2.3.1, a warning appears and installation is aborted.

If AP 2.3.1 is installed on the system but Solaris patch 110722-01 or later is not installed, a warning appears and installation is aborted.

Mounting the CD-ROM from the Command Line

If the Solaris volume management software is not available to mount the CD-ROM automatically, mount the CD-ROM manually. The following instructions assume that your CD driver is on controller 0.

When a default response is displayed in parentheses, press **Return** to accept that default.

Follow these steps:

1. Log in as superuser.

2. Insert the CD-ROM into a CD-ROM drive connected to your system.

3. To mount the CD-ROM manually, use the following command:

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

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

4. After the CD-ROM is mounted, use the following command to install the packages:

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

Note The order in which the packages are listed is significant. The VRTSlic package must be listed first, followed by the VRTSvxvm package, then the remaining packages.

If you added the VRTSlic and see a message that says no changes were made to the system, follow the instructions in "Removing Previous Licensing Packages" on page 19.

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

If the AP version is prior to 2.3.1, a warning appears and installation is aborted.

If AP 2.3.1 is installed on the system, but the Solaris patch 110722-01 is not installed, a warning appears and installation is aborted.

Running the vxinstall Program

This section describes how to initialize VxVM using the vxinstall program, which places all specified disks under VxVM control. These disks are automatically placed in the default disk group, rootdg.

Disks selected for configuration 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.

Required Reboots

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

Rootdg

You must initialize at least one disk into rootdg with the vxinstall program. At least one disk must remain in rootdg at all times while VxVM is running.

Placing Disks in another Disk Group

To place some disks in another disk group, use VMSA or the vxdiskadm program after completing the vxinstall program. Refer to the *VERITAS Volume Manager Administrator's Guide* for information on how to create other disk groups for your disks.

Run the vxinstall Program

Initialize VxVM using the following procedure:

1. Log in as superuser.

2. Identify controllers, disks and enclosures to exclude from being configured as VxVM devices by the vxinstall and vxdiskadm utilities. If there are no devices to exclude, continue to step 4.

Note These files do not exclude enclosures, controllers and disks from access by any other VxVM commands. See the vxdiskadm(1M) manual pages for more information.

If the vxinstall program does not find any exclude files, it prompts you about each disk in the system.

Caution Internal SEAGATE disks can be overwritten by the Volume Manager initialization routines when you run vxinstall. To avoid this, put the addresses of the disks in the /etc/vx/disks.exclude file, or put the addresses of their controllers in the /etc/vx/cntrls.exclude file.

3. To exclude devices from VxVM control, create or edit the `/etc/vx/cntrls.exclude, /etc/vx/disks.exclude` and `/etc/vx/enclr.exclude` files. You can:

- exclude one or more disks from VxVM control.

- exclude all disks on certain controllers from VxVM control.

- exclude all disks in specific enclosures from VxVM control.

The following section shows examples of each exclude file.

❖ To exclude one or more disks from VxVM control, create or edit the `/etc/vx/disks.exclude` file, and add the names of the disks to this file.

The following is an example of the contents of a `disks.exclude` file:

```
c0t1d0
```

❖ To exclude all disks on certain controllers from VxVM control, create or edit the `/etc/vx/cntrls.exclude` file, and add the names of the controllers to this file.

The following is an example of the contents of a `cntrls.exclude` file:

```
c0
c1
```

❖ To exclude all disks in specific enclosures from VxVM control, create or edit the `/etc/vx/enclr.exclude` file, and add the names of the enclosures to this file.

The following is an example of the contents of an `enclr.exclude` file:

```
sena0
emc1
```

If these files exist on a system prior to upgrading to this release of VxVM, their effect will still be visible. However, in release 3.1.1 of VxVM, a new set of interfaces has been added to the utilities `vxinstall` and `vxdiskadm` that can be used to hide devices from VxVM or prevent them from being multipathed by DMP.

4. Start the `vxinstall` program. Type:

```
# vxinstall
```

The `vxinstall` program first runs the `vxlicense` command to initialize the VxVM 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) automatically generates a key. The `vxinstall` program does not prompt for another key.

5. The `vxinstall` program then examines all controllers attached to the system and lists them. You are asked if you want to use enclosure-based naming. If you do not want to use enclosure-based names, enter **n**, and proceed to

```
Generating list of attached enclosures....

VxVM will use the following format to name disks on the host:

  <enclosurename>_<diskno>

In the above format, <enclosurename> is the logical name of the
enclosure to which the disk belongs. VxVM assigns default enclosure
names which can be changed according to the user requirements.

Some examples would be:

  hitachi0_2      - second disk detected in enclosure
                          'hitachi0'
  enggdept_2      - second disk detected in enclosure
                          'enggdept'
  dgc1_1          - first disk detected in enclosure 'dgc1'
  jbod_1          - first disk detected in the jbod category

For non-public loop disks (Ex. FC disks directly connected to the
host or through hubs), you could choose to use the c#t#d#s# naming
format.
Do you want to use enclosure based names for all disks ? [y,n,q,?]
(default: n) n
```

6. If you want to use enclosure-based names for the disks on the system, enter **y**. `vxinstall` detects and displays a list of disk arrays connected to your system.

```
Volume Manager Installation
Menu: VolumeManager/Install

The Volume Manager has detected the following categories of storage
connected to your system:

Enclosures:     enc01       sena0       sena1       sena3
                sena4       sena5
Others:         others0

Do you want to rename the enclosures [y,n,q,?] (default: y) Y
```

7. Enter **y** to assign a new name to the enclosures. This helps to identify their physical location. Enclosure names cannot exceed 31 characters in length.

```
Rename an enclosure
```

```
Menu: VolumeManager/Install/Rename enclosure

Enter an enclosure name: [list,q,?] test1
```

Enter the name of an enclosure you wish to rename.

```
Enter the new name for enclosure test1: enc0

   The enclosure name test1 has been changed to enc0

Hit RETURN to continue.
```

8. Enter **y** at the following prompt to rename more enclosures or **n** to continue with the installation.

```
Do you wish to rename more enclosures ? [y,n,q,?] (default: n) n
```

9. 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 want to exclude any devices from being seen by VxVM or not
be multipathed by vxdmp then use the Prevent
multipathing/Suppress devices from VxVM's view option, before you
choose Custom Installation or Quick Installation.

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.
```

10. Press Return to continue.

The vxinstall program displays a menu with the following options:

```
1  Quick Installation
2  Custom Installation
3  Prevent multipathing/Suppress devices from VxVM's view

?  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 VxVM control.

- `3 (Prevent multipathing/Suppress devices from VxVM's view)` excludes devices from VxVM or prevents them from being multipathed by DMP.

- `?` 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.

11. Select an option.

- If you choose `1`, see "Quick Installation" on page 25.

- If you choose `2`, see "Custom Installation" on page 31.

- If you choose `3`, see "Prevent Multipathing/Suppress Devices from VxVM's View" on page 36.

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. Encapsulation preserves existing data; initialization removes the data.

Initialization and Encapsulation Notes

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

◆ Quick Installation places every disk attached to your system under VxVM control, unless you have chosen to exclude those disks. Chose disks for exclusion using the `disks.exclude`, `cntrls.exclude` or `enclr.exclude` file or by using the

interface for suppressing devices as described earlier. If you have any disks on your system that you do not want placed under VxVM control, either exclude them using the `Prevent multipathing/Suppress devices from VxVM's view` in `vxinstall` option or `Custom Installation` option (which allows you to omit certain disks from installation).

◆ To encapsulate some disks in a given disk array and initialize others, use Custom Installation instead. (See "Custom Installation" on page 31.)

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

◆ 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.

Run Quick Installation

At any of the `vxinstall` prompts, you can select **q** to completely abandon the initialization and start over.

To use the Quick Installation option, follow these steps:

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

The `vxinstall` program prompts for each disk array and disk connected to your system. To accept a default response displayed in parentheses, press Return.

2. Indicate whether to encapsulate your boot disk at the following prompt. (In this example, the boot disk name is `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 are left unaltered.

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

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

```
Volume Manager Installation
Menu: VolumeManager/Install

Volume Manager Quick Installation
Menu: VolumeManager/Install/QuickInstall/enc0

Disk array serial number : 60020f20000005bd0000

Generating list of disks in disk array enc0....


Volume Manager has detected the following disks in disk array enc0:

  enc0_0 enc0_1 enc0_2 enc0_3

Hit RETURN to continue.
```

If any disks are listed in the `/etc/vx/disks.exclude`, `/etc/vx/cntrls.exclude` or `/etc/vx/enclr.exclude` files, the disks are listed here separately as excluded disks.

3. Press Return to continue.

4. The following message displays. Indicate whether you want to initialize all of the disks on the disk array.

Caution   Initializing a disk destroys all data on that disk.

```
Volume Manager Quick Installation For Disk Array enc0
Menu:VolumeManager/Install/QuickInstall/enc0

Initialize all disks on this disk array? (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 disk array instead and displays the following:

```
Volume Manager will now try to encapsulate all the disks on this
disk array.
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)
```

Enter **y** to continue with the disk initialization. The following message displays:

```
Volume Manager will now initialize all the disks on this
disk array including disks having valid partitions.
Hit RETURN to continue.
```

Press Return to continue.

5. Indicate whether you want default disk names to be assigned to all of the disks on the disk array 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 enc0_0 disk will be given disk name disk02

The enc0_0 disk appears to be empty. Adding as a new disk.
The enc0_1 disk will be given disk name disk05
The enc0_1 disk appears to be empty. Adding as a new disk.
The enc0_2 disk will be given disk name disk06
The enc0_2 disk appears to be empty. Adding as a new disk.
The enc0_3 disk will be given disk name disk07
The enc0_3 disk appears to be empty. Adding as a new disk.
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.
```

- If you did not select encapsulation, a message for empty disks to be initialized instead of encapsulated appears, 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 disk array individually, as follows:

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

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

7. When all of the disks on the current disk array have been named, press Return and repeat step 5 and step 6 for the remaining controllers on your system.

   When you have completed the vxinstall procedure for all disk arrays 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)
```

8. This is your last chance to review and alter your choices for any of the disks to be placed under VxVM 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 VxVM control:

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

9. Enter the name of the disk to be removed from the list and press Return.

10. Repeat these steps 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**.

11. 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 "Rebooting After vxinstall" on page 41.

A reboot is necessary if any disks are to be encapsulated.

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 VxVM control and how they are added. This is important if you do not want all your disks placed under VxVM control. Custom Installation also provides the option of initializing all disks in a disk array, encapsulating all disks in a disk array, or initializing some disks in a disk array and encapsulating others.

At any of the vxinstall prompts, you can select **q** to completely abandon the initialization and then start over.

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.

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 disk array and disk connected to your system. To accept a default response displayed in parentheses, press Return.

2. Indicate whether to encapsulate your boot disk. (In this example, 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 are left unaltered.

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

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

```
Volume Manager Custom Installation
Menu: VolumeManager/Install/Custom/enc0

Disk array serial number : 60020f20000005bd0000
```

```
Generating list of disks in disk array enc0....

The Volume Manager has detected the following disks in enclosure
enc0:

  enc0_0 enc0_1 enc0_2 enc0_3

Hit RETURN to continue.
```

If any disks are listed in the `/etc/vx/disks.exclude`,
`/etc/vx/cntrls.exclude`, or `/etc/vx/enclr.exclude` files, the disks are
listed here separately as excluded disks.

3. Press Return to continue.

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

```
Installation options for enclosure enc0
Menu: VolumeManager/Install/Custom/enc0

 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 in the disk array, 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 in the disk array individually:

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

3. For each disk, accept the default name or enter the desired disk name and press Return.

4. When all of the disks on the current disk array are named, press Return to move on to the next disk array.

Initializing All Disks

1. To initialize all disks in the disk array, select **2**.

Caution   Initializing a disk destroys all data on that disk.

All disks in the disk array are initialized. 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 VxVM control before, choose this option.

2. Indicate whether you want default disk names to be assigned to all of the disks in the disk array 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 in the current disk array 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 enc0
Menu: VolumeManager/Install/Custom/enc/enc0_0

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 4 on page 32.

3. Press Return to move on to the next disk. Repeat the previous step for the remaining disks on the current disk array.

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

Leaving All Disks Unaltered

1. To leave all disks in the disk array unaltered, select **4**.

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

2. Repeat step 4 on page 32 for each of the remaining disk arrays on your system.

    When you have completed the `vxinstall` procedure for all disk arrays 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.

      c0t5d0New Disk
      c1t0d0New Disk
      c1t1d0New Disk
      c1t2d0New Disk
      c2t2d3New Disk
      c3t0d0Encapsulate
      c3t1d0Encapsulate
      c3t2d0New Disk

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

3. This is your last chance to review and alter your choices for any of the disks to be placed under VxVM 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 VxVM control, as follows:

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

4. 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**.

5. 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 "Rebooting After vxinstall" on page 41.

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.
```

Prevent Multipathing/Suppress Devices from VxVM's View

This section describes how to exclude a device which is under VxVM or Dynamic Multipathing control. Follow these steps to prevent multipathing or suppress devices.

1. Select menu item 3 (`Prevent Multipathing/Suppress devices from VxVM's view`) from the `vxinstall` main menu.

   The following message displays:

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

   This operation might lead to some devices being suppressed from
   VxVM's view or prevent them from being multipathed by vxdmp. (This
   operation can be reversed using the vxdiskadm command).

   Do you want to continue? [y,n,q,?] (default: n) y
   ```

2. Enter **y**.

3. The following message displays. Select an operation.

   ```
   Volume Manager Device Operations
   Menu: VolumeManager/Install/Exclude Devices

   1 Suppress all paths through a controller from VxVM's view
   2 Suppress a path from VxVM's view
   3 Suppress disks from VxVM's view by specifying a VID:PID
   combination
   4 Suppress all but one path to a disk
   5 Prevent multipathing of all disks on a controller by VxVM
   6 Prevent multipathing of a disk by VxVM
   7 Prevent multipathing of disks by specifying a VID:PID combination
   8 List currently suppressed/non-multipathed devices
   ```

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

Select an operation to perform:
```

- If you select option 1, the following message displays:

```
Exclude controllers from VxVM
Menu: VolumeManager/Install/ExcludeDevices/CTLR-VXVM

Use this operation to exclude all paths through a controller from
VxVM.

This operation can be reversed using the vxdiskadm command.

You can specify a controller name at the prompt. A controller name
is of the form c#, example c3, c11 etc. Enter 'all' to exclude all
paths on all the controllers on the host. To see the list of
controllers on the system, type 'list'.

Enter a controller name:[ctlr_name,all,list,list-exclude,q,?]
```

Enter the controller name. The operation completes and the controller is disabled.

- If you select option 2, the following prompt displays:

```
Exclude paths from VxVM
Menu: VolumeManager/Install/ExcludeDevices/PATH-VXVM

Use this operation to exclude one or more paths from VxVM.

As a result of this operation, the specified paths will be excluded
from the view of VxVM. This operation can be reversed using the
vxdiskadm command.

You can specify a pathname or a pattern at the prompt. Here are
some path selection examples:

  all:all paths
  c4t2:all paths on controller 4, target 2
  c3t4d2:a single path
  list:list all paths on the system

Enter a pathname or pattern:[<Pattern>,all,list,list-exclude,q?]
```

The path specified here will be excluded from VxVM control.

- If you select option 3, you are required to enter the VID:PID combination of the
  device to be excluded from VxVM.

```
Exclude VID:PID from VxVM
Menu: VolumeManager/Install/ExcludeDevices/VIDPID-VXVM
```

```
Use this operation to exclude disks returning a specified
VendorID:ProductID combination from VxVM.
```

```
As a result of this operation, all disks that return
VendorID:ProductID matching the specified combination will be
excluded from the view of VxVM. This operation can be reversed
using the vxdiskadm command.
```

```
You can specify a VendorID:ProductID pattern at the prompt. The
specification can be as follows :
```

```
VID:PID          where VID stands for Vendor  ID
                       PID stands for Product ID
                 (The command vxdmpinq in /etc/vx/diag.d can be
                 used to obtain the Vendor ID and Product ID.)
```

```
Both VID and PID can have an optional '*' (asterisk) following
them.
If a '*' follows VID, it will result in the exclusion of all disks
returning Vendor ID starting with VID. The same is true for Product
ID as well. Both VID and PID should be non NULL. The maximum
allowed length for VendorID and ProductID are 8 and 16 characters
respectively.
```

```
Some examples of VID:PID specification are:
```

```
all- Exclude all disks
aaa:123- Exclude all disks having VID 'aaa' and PID '123'
aaa*:123- Exclude all disks having VID starting with 'aaa' and
   PID '123'
aaa:123*- Exclude all disks having VID 'aaa' and PID starting
   with '123'
aaa:*- Exclude all disks having VID 'aaa' and any PID
```

```
Enter a VID:PID combination:[<Pattern>,all,list,exclude,q,?]
```

The disks that match the VID:PID combination will be excluded from VxVM. The
Vendor ID and Product ID can be obtained from the Standard SCSI inquiry data
returned by the disk.

For example, the VID:PID combination for Sun's T3 disk array can be specified as
SUN:T3. The Vendor ID and Product ID of the disk can be obtained by the command
`/usr/lib/vxvm/diag.d/vxdmpinq`.

- Option 4 defines a pathgroup for disks which are not multipathed by VxVM.
  Only one pathgroup is made visible from the group. Specify the paths to be
  included in the pathgroup.

```
Exclude all but one paths to a disk
Menu: VolumeManager/Install/ExcludeDevices/PATHGROUP-VXVM
```

Use this operation to exclude all but one paths to a disk. In case
of disks which are not multipathed by vxdmp, VxVM will see each
path as a disk. In such cases, creating a pathgroup of all paths to
the disk will ensure only one of the paths from the group is made
visible to VxVM. The pathgroup can be removed using the vxdiskadm
command.

Example: If c1t30d0 and c2t30d0 are paths to the same disk and both
are seen by VxVM as separate disks, c1t30d0 and c2t30d0 can be put
in a pathgroup that only one of these paths is visible to VxVM.

The pathgroup can be specified as a list of blank separated paths,
for example, c1t30d0 c2t30d0.

Enter pathgroup: [<pattern>,list,list-exclude,q,?]

The next three options allow you to exclude devices from vxdmp.

- If option 5 is selected, the following screen displays:

Exclude controllers from DMP
Menu: VolumeManager/Install/ExcludeDevices/CTLR-DMP

Use this operation to exclude all disks on a controller from being
multipathed by vxdmp.

As a result of this operation, all disks having a path through the
specified controller will be claimed in the OTHER_DISKS category
and hence, not multipathed by vxdmp. This operation can be reversed
using the vxdiskadm command.

You can specify a controller name at the prompt. A controller name
is of the form c#, example c3, c11 etc. Enter 'all' to exclude all
paths on all the controllers on the host. To see the list of
controllers on the system, type 'list'.

Enter a controller name:[<ctlr-name>,all,list,list-exclude,q,?]

The controller entered here is excluded from DMP control.

- Option 6 excludes the specified path from multipathing. Enter the path name at
the prompt after this screen:

Exclude paths from DMP
Menu: VolumeManager/Install/ExcludeDevices/PATH-DMP

Use this operation to exclude one or more disks from vxdmp.

As a result of this operation, the disks corresponding to the
specified paths will not be multipathed by VxDMP. This operation
can be reversed using the vxdiskadm command.

You can specify a pathname or a pattern at the prompt. Here are
some path selection examples:

```
all:all paths
c4t2:all paths on controller 4, target 2
c3t4d2:a single path
list:list all paths on the system
```

Enter a pathname or pattern:[<pattern>,all,list,list-exclude,q,?]

If a path is specified, the corresponding disks will be claimed in the
OTHER_DISKS category and will not be multipathed.

-   Selecting option 7 excludes the disk corresponding to the VIP:PID specified.
    Enter the VID:PID combination at the prompt that follows this screen:

```
Exclude VID:PID from DMP
Menu: VolumeManager/Install/ExcludeDevices/VIDPID-DMP
```

Use this operation to prevent vxdmp from multipathing devices
returning a specific VID:PID combination.

As a result of this operation, all disks that return VID:PID
matching the specified combination will be claimed in the
OTHER_DISKS category(i.e. they will not be multipathed by vxdmp).
This operation can be reversed using the vxdiskadm command.

You can specify a VendorID:ProductID combination at the prompt. The
specification can be as follows:

```
VID:PID        where VID stands for Vendor ID
                     PID stands for Product ID
```

Both VID and PID can have an optional '*' (asterisk) following
them.
If a '*' follows VID, it will result in the exclusion of all disks
returning Vendor ID starting with VID. The same is true for Product
ID as well. Both VID and PID should be non NULL.

Some examples of VID:PID specification are:

```
all- Exclude all disks
aaa:123- Exclude all disks having VID 'aaa' and PID '123'
aaa*:123- Exclude all disks having VID starting with 'aaa' and
   PID '123'
aaa:123*- Exclude all disks having VID 'aaa' and PID starting
   with '123'
aaa:* - Exclude all disks having VID 'aaa' and any PID
```

Enter a VID:PID combination:[<pattern>,all,list,list-exclude,q,?]

All disks returning a VID:PID combination will be claimed in the OTHER_DISKS category and will not be multipathed.

Note If you selected any one of the options, reboot the system for device exclusion to take effect and re-run `vxinstall`. After the reboot for menu item #3 (`Prevent Multipathing/Suppress devices`) run `vxinstall` again to complete the installation.

Rebooting After vxinstall

After the `vxinstall` program processes the disks attached to your system, you may have to reboot the system to make changes to your disks. The way in which you chose to handle your disks during the `vxinstall` session determines whether a reboot is required. If you chose to encapsulate any disks, a reboot is necessary.

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

Note The setup you choose can require several reboots.

When it is necessary to 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, reboot as soon as possible.

Note During the next one or more reboots, you may be asked several times to continue an operation. Press the Return key at each of these prompts to accept the default answers. 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 VMSA

The following software must be installed and running on the machine to be administered by VMSA:

◆ Solaris 2.6 or higher.

◆ VxVM Release 2.5.7, 3.0.4 3.1, 3.1.1 or 3.2. The `vxconfigd` daemon must be running on the machine to be administered.

◆ The VMSA 3.2 server.

This section describes how to install and set up VMSA. The VMSA server must be installed and run on a Solaris machine running VxVM. The VMSA client can be installed and run on a Solaris or Windows machine that supports the Java 1.1 Runtime Environment.

For information on how to upgrade from a previous release of VMSA, see "Upgrading VMSA" on page 66.

Note This release of VMSA requires VxVM Release 2.5.7, 3.0.4, 3.1, 3.1.1 or 3.2. Before you install and run VMSA, you must install VxVM and run the `vxinstall` program to initialize VxVM.

The following are minimum system recommendations for the VMSA client:

Solaris:	SPARCstation 5 with 64M memory
Windows:	100MHz Pentium with 32M memory

VMSA Installation Overview

Set up and start VMSA as follows:

1. Install the VMSA server and client on the UNIX (Solaris) machine to be administered. See "Installing the VMSA Server and/or Client (Solaris)" on page 43.

2. If users other than `root` need access to VMSA, set up security and specify which users can run VMSA. See "Customizing Security (Optional)" on page 45.

3. If you plan to run the VMSA client on a machine other than the machine to be administered, install the VMSA client on the machine where the client will run. See "Installing the VMSA Server and/or Client (Solaris)" on page 43 or "Installing the VMSA Client (Windows)" on page 44.

4. Start the VMSA server on the machine to be administered. See "Starting the VMSA Server" on page 46.

5. Start the VMSA client. See "Starting the VMSA Client" on page 47.

Installing the VMSA Server and/or Client (Solaris)

Install the VMSA server and client on a Solaris machine as follows:

1. Log in as superuser.

2. Install the VERITAS Volume Manager Storage Administrator package as described in "Loading the Software from CD-ROM" on page 18.

The VMSA client can be installed with or without the VMSA server. To install the VMSA server, answer **y** at the appropriate prompt.

> **Note** The VMSA 3.2 GUI client is not compatible with earlier versions of the VMSA server. When installing VMSA 3.2, both the VMSA server and the GUI client must be updated.

The VMSA package for UNIX contains the following directories:

Directory	Contents
`jre/`	Java Runtime Environment (JRE)
`bin/`	`vmsa_server`—VMSA server startup script `vmsa`—VMSA startup script
`vmsa/`	`java/images/`—VMSA image `.gif` files `java/help/`—VMSA help files `server/`—VMSA server libraries and properties
`/opt/VRTS/man/`	VMSA manual pages

3. In your `.profile` file, update your PATH environment variable to include the directory containing the VMSA startup scripts using the following command:

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

where **_Storage_Administrator_directory_** is the directory where you installed VMSA.

Installing the VMSA Client (Windows)

The VMSA client runs on Windows NT, Windows 2000, Windows ME, Windows 98, and Windows 95 machines. If you plan to run VMSA from a Windows machine, install the optional Windows package after you have installed the VMSA server on a UNIX (Solaris) machine.

Before you install VMSA on a Windows machine, you must uninstall any existing VERITAS Volume Manager Storage Administrator packages and remove the old `setup.exe` from that machine. Only one VMSA package can be installed on a Windows machine at any given time.

Install the VMSA client on a Windows machine as follows:

1. Log in as administrator.

2. Insert the CD-ROM containing VMSA.

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 ME, Windows 98 or Windows 95, restart the computer.

Customizing Security (Optional)

By default, only `root` can run VMSA. However, you can allow other users to access VMSA by following these steps:

1. On the machine to be administered, add a group named `vrtsadm` to the group file `/etc/group`. Or, add a group named `vrtsadm` to the Network Information Name Service (NIS) group table.

2. The `vrtsadm` group line must include the user names of all users (including `root`) who will have access to VMSA.

   Example:

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

   In the example, `joe`, `bill`, and root can access VMSA.

   If you do not create a `vrtsadm` group, then only `root` has access to VMSA. If the `vrtsadm` group exists, `root` must be included in the `vrtsadm` group for `root` to access VMSA.

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 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`

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

Monitoring Access to VMSA

Once you have set up security for VMSA, you can monitor access to it 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 VMSA Server

To use VMSA, the VMSA server must be running on the Solaris machine to be administered. Once started, the VMSA server must run continuously. Only one instance of the VMSA server should be running at a time. By default, the VMSA server is started automatically when the system reboots. See " for other methods of starting the VMSA server.

Start the VMSA server manually as follows:

1. Log in as superuser.

2. Go to the directory where VMSA is installed, as follows:

```
# cd $VMSAHOME/bin
```

3. Start the server, as follows:

```
# ./vmsa_server &
```

Note If you are using the sh shell, use the `nohup` command as follows to allow the VMSA server to continue running after you logout:
`# nohup ./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
```

4. To confirm that VMSA is running, use the following command:

```
# ./vmsa_server -q
```

5. To stop the VMSA server, use the following command:

```
# ./vmsa_server -k
```

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

Auto-start Mode

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. 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. The default mode is auto-start disabled and the VMSA server is started automatically when the system reboots.

1. To start the server in auto-start mode, use the following command:

   ```
   # autostart on
   ```

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

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

   The output should show that the VMSA service is in a LISTEN state:

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

3. Auto-start mode remains in effect across system reboots. To disable auto-start mode, use the following command:

   ```
   # autostart off
   ```

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

Starting the VMSA Client

Note Only users with appropriate privileges can run VMSA. Refer to "Customizing Security (Optional)" on page 45 for information on VMSA security and access permissions.

VMSA can administer the local machine or a remote machine. However, VxVM and the VMSA server must be installed on the machine to be administered. The VxVM vxconfigd daemon and the VMSA server must be running on the machine to be administered.

After installing VxVM and VMSA and starting the server, start VMSA in one of the following ways.

UNIX (Solaris)

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**.

The Read-Only Mode for VMSA

VMSA 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, but it prevents administrative actions from taking effect. When VMSA is in read-only mode, the main window and dialog box title bars contain the following string:

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

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

Creating a Read-only Mode User List

To specify a list of users with permission to run VMSA 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 VMSA in read-only mode, but are not allowed to run it in full operational mode.

Running Read-only Mode

Run VMSA 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 VMSA in read-only mode on that host. This enables read-only mode on the specified host machine.

Enabling Cluster Support in VxVM (Optional)

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

This release includes an *optional* cluster feature that enables VxVM to be used in a cluster environment. The cluster functionality in the VxVM allows multiple hosts to simultaneously access and manage a set of disks under VxVM 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 VxVM 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 it limits what can be done to private disk groups.)

To enable the cluster functionality in VxVM, use the following procedure

1. Obtain a license for the VxVM cluster feature.

2. Install the software packages onto each system (node) to be included in the cluster as described in "Loading the Software from CD-ROM" on page 18.

3. Initialize VxVM using the `vxinstall` procedure as described in "Running the vxinstall Program" on page 20.

4. Set up and start VMSA as described in "Setting Up and Starting VMSA" on page 42.

5. Configure shared disks as described in "Configuring Shared Disks" on page 50.

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 VxVM 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 "Disk Controller Reservation" on page 52).

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. On one node, run the `vxdiskadm` program and choose option 1 to initialize new disks. When asked to add these disks to a disk group, choose `none` to leave the disks for future use.

3. On other nodes in the cluster, run `vxdisk scandisks` to see the newly initialized disks.

4. From the master node, create disk groups on the shared disks. To determine if a node is a master or slave, run `vxdctl -c mode`.

   Use the `vxdg` program or VMSA to create disk groups. In the `vxdg` program, use the `-s` option to create shared disk groups.

5. From the master node only, use `vxassist` or VMSA to create volumes in the disk groups.

> **Note** The volumes must be of type `gen`. Do not create RAID-5 volumes. Before creating any log subdisks, read the section on DRL in the *VERITAS Volume Manager Adminstrator's Guide.*

6. 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.

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 existing shared disk groups.

Converting Existing VxVM Disk Groups to Shared Disk Groups

If you are upgrading from VxVM 2.x to VxVM 3.2 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. 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 disk-group-name
    ```

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

    ```
    # vxdg -s import disk-group-name
    ```

    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 3.2

If you are running a release earlier than 3.2, you can upgrade your system and retain your existing VxVM configuration. After upgrading, you can resume using VxVM as before (without running the `vxinstall` program again). No new license is required for upgrading.

To upgrade, you can use one of the following processes:

To upgrade		
VxVM	On an encapsulated root disk	See "Upgrading VxVM on an Encapsulated Root Disk" on page 56.
VxVM	On an unencapsulated root disk	See "Upgrading VxVM Only" on page 62
VxVM and Solaris	On an encapsulated or unencapsulated root disk	See "Upgrading VxVM and Solaris" on page 59.
Solaris	On an encapsulated or unencapsulated root disk	See "Upgrading Solaris Only" on page 65

The following procedures explain how to upgrade VxVM, Solaris, or both. This section also contains information on how to upgrade VMSA.

To upgrade, you need to:

1. Install patches.

2. Determine if the root disk is encapsulated by running the following command:

   # **df -k**

   If the output from this command includes a pathname that contains vx and rootvol as in /dev/vx/dsk/rootvol, then the root disk is encapsulated and you must upgrade both VxVM and Solaris. See "Upgrading VxVM on an Encapsulated Root Disk" on page 56.

Caution You need to use the procedures in "Upgrading VxVM on an Encapsulated Root Disk" on page 56 if you have a system with an encapsulated root disk. Use the appropriate upgrade process to upgrade your system.

The files array.info, vxvm.exclude and vxdmp.exclude are saved before an upgrade and restored after the VxVM package has been upgraded.

Note Importing a pre-3.2 VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.2 level. You may need to manually upgrade each of your disk groups following a VxVM upgrade. See "Upgrading a Disk Group" in the *VERITAS Administrator's Guide* for information on how to upgrade a disk group.

Adding Patches

You need to install at least one required patch before proceeding with your upgrade. Please refer to the "Solaris Patches for VxVM" section *before you upgrade* to VxVM 3.2.

VxVM and Solaris Release Support Matrix

The table below indicates which VxVM releases support these Solaris releases.VMSA

	Solaris 2.5.1	Solaris 2.6	Solaris 7	Solaris 8
VxVM 3.2		Supported	Supported	Supported
VxVM 3.1.1		Supported	Supported	Supported
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 and Solaris Upgrade Paths

The following table shows possible VxVM/Solaris upgrade paths.

Upgrade Paths

You can upgrade VxVM and Solaris versions:	To VxVM and Solaris versions:	
VxVM 2.0 to 3.0 & Solaris 2.3 to 7	VxVM 3.2 & Solaris 2.6, 7, 8	See "Upgrading VxVM and Solaris" on page 59
VxVM 3.1 & 3.1.1 & Solaris 2.6, 7	VxVM 3.2 & Solaris 7, 8	
You can upgrade VxVM on the following Solaris versions:	**To VxVM and Solaris versions:**	
VxVM 2.4 to 3.0 & Solaris 2.6	VxVM 3.2 & Solaris 2.6	For encapsulated root disks, see "Upgrading VxVM on an Encapsulated Root Disk" on page 56. For unencapsulated root disks, see "Upgrading VxVM Only" on page 62
VxVM 3.0.x & Solaris 2.6, 7, 8	VxVM 3.2 & Solaris 2.6, 7, 8	
VxVM 3.1 & Solaris 2.6, 7, 8	VxVM 3.2 & Solaris 2.6, 7, 8	
VxVM 3.1.1 & Solaris 2.6, 7, 8	VxVM 3.2 & Solaris 2.6, 7, 8	
You can upgrade Solaris only:	**To Solaris versions:**	
VxVM 3.2 & Solaris 2.6	VxVM 3.2 & Solaris 7, 8	"Upgrading Solaris Only" on page 65
VxVM 3.2 & Solaris 7	VxVM 3.2 & Solaris 8	

Upgrading VxVM on an Encapsulated Root Disk

This section explains how to upgrade VxVM if the root disk is encapsulated. Refer to the "VxVM and Solaris Release Support Matrix" to determine if upgrading the operating system is required.

Caution	If your root disk is not encapsulated, refer to "Upgrading VxVM Only" on page 62.

Follow these steps:

1. If you have not already obtained and installed a VxVM 3.2 license key, do so now. See "Obtain a License Key" on page 7 for details.

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

   This makes these file systems accessible during the upgrade process by temporarily converting file systems on volumes back to using direct disk partitions. 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. If any of the documentation, developer kit, or manual page packages are to be installed, the directory `/opt` must exist, be writable, and must not be a symbolic link. This is because the volumes not temporarily converted by the `upgrade_start` will be unavailable during the upgrade process. If you have a symbolic link from `/opt` to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in `/opt` will not be installed.

4. Boot the system using the following command. Note that this starts the system in single-user mode.

   ```
   # reboot -- -s
   ```

5. Load and mount the CD-ROM by starting the volmgt daemon. After starting the daemon, the path is `/cdrom/CD_name`.

   ```
   # /etc/init.d/volmgt start
   ```

6. Run the `upgrade_start` script to prepare the previous release of VxVM for its removal using the following command:

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

The upgrade_start script searches for volumes containing file systems.

7. Reboot to single-user mode.

   # **reboot -- -s**

Note You need to mount /opt manually before you can run pkgrm to remove VxVM packages if /opt is on it's own partition.

---

8. Remove the old VxVM packages by entering:

   # **pkgrm VRTSvmsa VRTSvmdoc VRTSvmdev VRTSvmman VRTSvxvm VRTSlic**

Note If **VRTSlic** is installed, you must remove it. The key remains even if the package is removed.

---

9. 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**

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

---

10. Reboot the machine to single-user mode by running the following command:

    # **reboot -- -s**

11. Load and mount the CD-ROM by starting the volmgt daemon.

    # **/etc/init.d/volmgt start**

Note You need to mount /opt manually before you can run pkgadd to add VxVM packages if /opt is on it's own partition.

---

12. Add the VxVM 3.2 licensing package using the following command:

    # **pkgadd -d /CD_path/pkgs VRTSlic**

13. Add the VxVM 3.2 package using the following command:

    # **pkgadd -d /CD_path/pkgs VRTSvxvm**

If warnings are displayed that include the string `/etc/vx`, ignore them and continue.

14. Complete the upgrade using the following command:

/_CD_path_/scripts/upgrade_finish

15. Perform reconfiguration reboot:

 # reboot -- -r

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

vxrelocd root &

17. Using the `pkgadd` command, install any or all of the additional packages `VRTSvmsa`, `VRTSvmdoc`, `VRTSvmman`, and `VRTSvmdev`.

18. Importing a pre-3.2 VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.2 level. You may need to manually upgrade each of your disk groups following a VxVM upgrade. See "Administering Disks" in the *VERITAS Volume Manager Administrator's Guide* for information on how to upgrade a disk group.

19. See "Setting Up and Starting VMSA" on page 42 to set up and restart VMSA.

Upgrading VxVM and Solaris

This section explains how to upgrade both Solaris and VxVM. (Refer to the "VxVM and Solaris Release Support Matrix" to determine if upgrading the operating system is required.)

Follow these steps:

1. If you have not already obtained and installed a VxVM 3.2 license key, do so now. See "Obtain a License Key" on page 7 for details.

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

This makes these file systems accessible during the upgrade process by temporarily converting file systems on volumes back to using direct disk partitions. 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. If any of the documentation, developer kit, or manual page packages are to be installed, the directory /opt must exist, be writable, and must not be a symbolic link. This is because the volumes not temporarily converted by the upgrade_start will be unavailable during the upgrade process. If you have a symbolic link from /opt to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in /opt will not be installed.

4. Boot the system using the following command. Note that this starts the system in single-user mode.

```
# reboot -- -s
```

5. Load and mount the CD-ROM by starting the volmgt daemon. After starting the daemon, the path is /cdrom/*CD_name*.

```
# /etc/init.d/volmgt start
```

6. Run the upgrade_start script to prepare the previous release of VxVM for its removal using the following command:

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

The upgrade_start script searches for volumes containing file systems.

7. Reboot to single-user mode.

```
# reboot -- -s
```

8. If you are upgrading Solaris, upgrade the operating system to Solaris 2.6, 7, or 8. Refer to the Solaris installation documentation for instructions on how to upgrade Solaris.

9. If you upgraded Solaris, boot up in single-user mode.

```
# reboot -- -s
```

Note You need to mount /opt manually before you can run pkgrm to remove VxVM packages if /opt is on it's own partition.

10. Remove the old VxVM packages by entering:

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

Note If **VRTSlic** is installed, you must remove it. The key remains even if the package is removed.

11. 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
```

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

12. Reboot the machine to single-user mode by running the following command:

```
# reboot -- -s
```

13. Load and mount the CD-ROM by starting the volmgt daemon.

```
# /etc/init.d/volmgt start
```

Note You need to mount /opt manually before you can run pkgadd to add VxVM packages if /opt is on it's own partition.

14. Add the VxVM 3.2 licensing package using the following command:

```
# pkgadd -d /CD_path/pkgs VRTSlic
```

15. Add the VxVM 3.2 package using the following command:

pkgadd -d /*CD_path*/pkgs VRTSvxvm

If warnings are displayed that include the string /etc/vx, ignore them and continue.

16. Complete the upgrade using the following command:

/*CD_path*/scripts/upgrade_finish

17. Perform reconfiguration reboot:

 # reboot -- -r

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

vxrelocd root &

19. Using the pkgadd command, install any or all of the additional packages VRTSvmsa, VRTSvmdoc, VRTSvmman, and VRTSvmdev.

20. Importing a pre-3.2 VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.2 level. You may need to manually upgrade each of your disk groups following a VxVM upgrade. See "Administering Disks" in the *VERITAS Volume Manager Administrator's Guide* for information on how to upgrade a disk group.

21. See "Setting Up and Starting VMSA" on page 42 to set up and restart VMSA.

Upgrading VxVM Only

If you are already running a version of Solaris that is supported with VxVM 3.2, you can upgrade VxVM without upgrading Solaris. You will completely remove VMSA, install VxVM, and reinstall VMSA.

Caution	Step 2 of this procedure instructs you to find out if your root disk is encapsulated. If it is encapsulated, you must use the procedure found in "Upgrading VxVM on an Encapsulated Root Disk" on page 56.

Upgrade VxVM as follows:

1. Log in as superuser.

2. Find out if the root disk is encapsulated by running the following command:

```
# df -k
```

If the output from this command includes a pathname that contains vx and rootvol as in /dev/vx/dsk/rootvol, then the root disk is encapsulated and you should follow the steps in the section "Upgrading VxVM on an Encapsulated Root Disk" on page 56.

If the root disk is not encapsulated, continue to the next step in this procedure.

3. Close all volumes through the applications that use them. To stop all volumes:

```
vxvol -f stopall
```

4. Stop the vxconfigd daemon:

```
vxdctl stop
```

5. Stop the vxiod daemon:

```
vxiod -f set 0
```

6. Remove the VMSA software using the following command:

```
# pkgrm VRTSvmsa
```

7. If the VRTSlic license package has been installed, remove it by entering:

```
# pkgrm VRTSlic
```

Note	If VRTSlic is installed, you must remove it. The key remains even if the package is removed.

8. Add the VxVM 3.2 licensing package using the following command:

> # **pkgadd -d /*CD_path*/pkgs VRTSlic**

9. If you added the VRTSlic and see a message that says no changes were made to the system, follow the instructions in "Removing Previous Licensing Packages" on page 19.

10. Add the VxVM 3.2 package using the following command:

> # **pkgadd -d /*CD_path*/pkgs VRTSvxvm**

11. Using the pkgadd command, install any or all of the additional packages: VRTSvmsa, VRTSvmdoc, VRTSvmman, and VRTSvmdev.

12. If you are prompted to run a reboot with the -r option, type:

> # **reboot -- -r**

13. See "Setting Up and Starting VMSA" on page 42.

Note The VMSA Release 3.2 server is not backward compatible with VMSA Release 1.x clients. Upgrade any existing clients to Release 3.2.

14. Depending on your existing configuration, various messages and prompts may result.

15. If AP is installed on the system and the AP version is prior to 2.3.1, a warning appears and installation is aborted. If AP 2.3.1 is installed on the system but the Solaris patch 110722-01 or later is not installed, a warning appears and installation is aborted.

16. DMP is automatically enabled in VxVM 3.2. You will see the following message:

```
DMP driver was previously disabled, but as this version of VxVM
requires DMP to be enabled, it will be enabled. If you wish to
exclude certain or all devices from being multipathed by DMP, you
can use vxinstall or vxdiskadm utilities to do so after the upgrade
is complete. Please refer to documentation on these utilities for
more details.
Continue installation? [y,n,?] (default: n):
```

- If you enter **y**, DMP will be enabled and installation will proceed. If you enter **n**, installation will abort.

17. 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 the previous set of messages.

18. It may be necessary to reboot after running `pkgadd`. You are informed that a reboot is required.

19. 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.

a. In the startup file, `/etc/rc2.d/S95vxvm-recover`, comment out the `vxrelocd` line:

```
# vxrelocd root &
```

b. After editing this file, the system must be rebooted:

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

Upgrading Solaris Only

If you are running VxVM 3.2 with Solaris 2.6 or 7, you can upgrade Solaris using the following procedure.

1. If the boot disk is mirrored, detach the mirror. You can use VMSA or line commands.

2. If any of the file systems /, /usr, /var, /opt, /export/home, /var/adm or /usr/kvm 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 (located in the /scripts directory on the CD) 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. If any of the documentation, developer kit, or manual page packages are to be installed, the directory /opt must exist, be writable, and must not be a symbolic link. This is because the volumes not temporarily converted by the upgrade_start will be unavailable during the upgrade process. If you have a symbolic link from /opt to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in /opt will not be installed.

4. In the /etc/vfstab file, comment out any data volumes on VxVM with VxFS.

5. Reboot to single user mode:

   # reboot -- -s

6. Load and mount the CD-ROM by starting the volmgt daemon. After starting the daemon, the path is /cdrom/*CD_name*.

   # **/etc/init.d/volmgt start**

7. Run the upgrade_start script to prepare the previous release of VxVM for its removal.

   # **/*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 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.

8. Reboot to single user mode:

   # reboot -- -s

9. Upgrade the operating system to Solaris 7 or 8. Refer to the Solaris installation documentation for instructions on how to upgrade Solaris. After installing Solaris, install the required patches shown in "Solaris Patches for VxVM" on page 16.

10. Reboot to single-user mode:

    # reboot -- -s

11. Load and mount the CD-ROM by starting the volmgt daemon. After starting the daemon, the path is /cdrom/*CD_name*.

    # **/etc/init.d/volmgt start**

12. Complete the upgrade by entering:

    # /*CD_path*/**scripts/upgrade_finish**

13. Reboot to multi-user mode using a command such as the /etc/shutdown command.

14. Edit the /etc/vfstab file to remove comments and restore any changes made before running **upgrade_start**.

Upgrading VMSA

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

Note The VMSA Release 3.2 server is not backward compatible with VMSA Release 1.x clients; therefore, upgrade any existing client(s) to Release 3.2.

1. Log in as superuser.

2. Remove the previous VRTSvmsa package using the following command:

   # **pkgrm VRTSvmsa**

3. See "Loading the Software from CD-ROM" on page 18 for instructions on adding the current VRTSvmsa package.

Checking VxVM Processes

After VxVM has been successfully installed and you have rebooted your system, confirm that key VxVM processes (vxconfigd, vxnotify, and vxrelocd) are running.

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

Use the following command:

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

Output similar to the following displays:

```
root458 180 10:23:34? 0:00 vxrelocd root
root 414 180 10:23:06? 0:03 vxconfigd
root 478 414 10:23:37? 0:00 vxrelocd root
root 479 478 5 10:23:37? 0:00 vxnotify -f -w 15
```

Adding Unsupported Disk Arrays to the JBOD Category on Systems Running VxVM 3.2

Use this procedure when adding new array support after VxVM is installed. For more information, see vxddladm help addjbod.

1. Enter the following command to add the new JBOD category:

   ```
   # vxddladm addjbod vid=vendorid
   ```

   where *vendorid* is FUJITSU, IBM, or SEAGATE as appropriate.

2. Enter the following command to bring the arrays under VxVM control:

   ```
   # vxdctl enable
   ```

This procedure ensures that dynamic multipathing is set up correctly on the array. Otherwise, Volume Manager treats the independent paths to the disks as separate devices, which can result in data corruption.

Starting and Enabling the Configuration Daemon vxconfigd

The VxVM configuration daemon (`vxconfigd`) maintains VxVM disk and disk group configurations. The `vxconfigd` communicates configuration changes to the kernel and modifies configuration information stored on disk.

Startup scripts usually invoke `vxconfigd` at system boot time. The `vxconfigd` daemon must be running for VxVM to operate properly.

The following procedures describe how to check `vxconfigd` has been started, whether it is enabled or disabled, how to start it manually, or how to enable it as required.

To determine whether `vxconfigd` is enabled, use the following command:

    # **vxdctl mode**

The following message indicates that the `vxconfigd` daemon is running and enabled:

    mode: enabled

This message indicates that `vxconfigd` is not running:

    mode: not-running

To start the `vxconfigd` daemon, enter the following command:

    # **vxconfigd**

This message indicates that `vxconfigd` is running, but not enabled:

    mode: disabled

To enable the volume daemon, enter the following command:

    # **vxdctl enable**

Once started, `vxconfigd` automatically becomes a background process.

By default, `vxconfigd` writes error messages to the console. However, you can configure it to write errors to a log file. For more information, see the `vxconfigd`(1M) and `vxdctl`(1M) manual pages.

Starting the Volume I/O Daemon vxiod

The volume I/O daemon (`vxiod`) provides extended I/O operations without blocking calling processes. Several `vxiod` daemons are usually started at system boot time after your initial installation, and they should be running at all times. The procedure below describes how to verify that the `vxiod` daemons are running, and how to start them if necessary.

To verify that `vxiod` daemons are running, enter the following command:

    # **vxiod**

Note The vxiod daemon is a kernel thread and is not visible using the ps command.

If any vxiod daemons are running, the following message displays:

```
N volume I/O daemons running
```

where *N* is the number of vxiod daemons currently running.

If no vxiod daemons are currently running, start some by entering this command:

```
# vxiod set N
```

where *N* equals the desired number of vxiod daemons. It is recommended that at least one vxiod daemon should be run for each CPU in the system.

For more information, see the vxiod(1M) manual page.

System Setup Guidelines

This section contains information to help you set up your system for efficient storage management. For specific setup tasks, refer to the *VERITAS Volume Manager Storage Administrator Administrator's Guide*.

The following system setup sequence is typical and is used as an example. Your system requirements may differ. The system setup guidelines provide helpful information for specific setup configurations.

▼ **Initial Setup**

1. Place disks under VxVM control.

2. Create new disk groups (if you do not want to use rootdg or you require additional disk groups).

3. Create volumes.

4. Put file systems on volumes.

▼ **Options**

❖ Encapsulate the boot/root disk and mirror it to create an alternate boot disk.

❖ Designate hot-relocation spare disks.

❖ Add mirrors to volumes.

▼ **Maintenance**

❖ Resize volumes and file systems.

❖ Add more disks or disk groups.

❖ Create snapshots.

The following general guidelines are to help you understand and plan an efficient storage management system. See the cross-references in each section for more information about the featured guideline.

Disabling Hot-Relocation

The hot-relocation feature detects disk failures automatically and notifies you of the nature of the failure, attempts to relocate any affected subdisks that are redundant, and initiates recovery procedures. 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 &
```

Hot-Relocation Guidelines

Refer to the following guidelines when using hot-relocation.

◆ The hot-relocation feature is enabled by default. Although it is possible to disable hot-relocation, it is advisable to leave it enabled.

◆ Although hot-relocation does not require you to designate disks as spares, designate at least one disk as a spare within each disk group. This gives you some control over which disks are used for relocation. If no spares exist, VxVM uses any available free space within the disk group. When free space is used for relocation purposes, it is possible to have performance degradation after the relocation.

◆ After hot-relocation occurs, designate one or more additional disks as spares to augment the spare space. Some of the original spare space may be occupied by relocated subdisks.

◆ If a given disk group spans multiple controllers and has more than one spare disk, set up the spare disks on different controllers (in case one of the controllers fails).

◆ For a mirrored volume, configure the disk group so that there is at least one disk that does not already contain a mirror of the volume. This disk should either be a spare disk with some available space or a regular disk with some free space and the disk is not excluded from hot-relocation use.

◆ For a mirrored and striped volume, configure the disk group so that at least one disk does not already contain one of the mirrors of the volume or another subdisk in the striped plex. This disk should either be a spare disk with some available space or a regular disk with some free space and the disk is not excluded from hot-relocation use.

◆ For a RAID-5 volume, configure the disk group so that at least one disk does not already contain the RAID-5 plex (or one of its log plexes) of the volume. This disk should either be a spare disk with some available space or a regular disk with some free space and the disk is not excluded from hot-relocation use.

◆ If a mirrored volume has a DRL log subdisk as part of its data plex, you cannot relocate the data plex. Instead, place log subdisks in log plexes that contain no data.

◆ Hot-relocation does not guarantee to preserve the original performance characteristics or data layout. Examine the locations of newly-relocated subdisks to determine whether they should be relocated to more suitable disks to regain the original performance benefits.

◆ Hot-relocation is capable of creating a new mirror of the root disk if the root disk is mirrored and it fails. The rootdg should contain sufficient contiguous spare or free space to accommodate the volumes on the root disk. (`rootvol` and `swapvol` require contiguous disk space.)

◆ Although it is possible to build VxVM objects on spare disks (using `vxmake` or the VMSA interface), it is recommended that you use spare disks for hot-relocation only.

Striping Guidelines

Refer to the following guidelines when using striping. See *VERITAS Volume Manager Administrator's Guide* for more information.

◆ Do not place more than one column of a striped plex on the same physical disk.

◆ Calculate stripe unit sizes carefully. In general, a moderate stripe unit size (for example, 64 kilobytes, which is also the default used by `vxassist`) is recommended.

◆ If it is not feasible to set the stripe unit size to the track size, and you do not know the application I/O pattern, use 64 kilobytes for the stripe unit size.

> **Note** Many modern disk drives have *variable geometry*. This means that the track size differs between cylinders, so that outer disk tracks have more sectors than inner tracks. It is therefore not always appropriate to use the track size as the stripe unit size. For these drives, use a moderate stripe unit size (such as 64 kilobytes), unless you know the I/O pattern of the application.

◆ Volumes with small stripe unit sizes can exhibit poor sequential I/O latency if the disks do not have synchronized spindles. Generally, striping over disks without synchronized spindles yields better performance when used with larger stripe-unit sizes and multi-threaded, or largely asynchronous, random I/O streams.

◆ Typically, the greater the number of physical disks in the stripe, the greater the improvement in I/O performance; however, this reduces the effective mean time between failures of the volume. If this is an issue, combine striping with mirroring to combine high-performance with improved reliability.

◆ If only one plex of a mirrored volume is striped, set the policy of the volume to `prefer` for the striped plex. (The default read policy, `select`, does this automatically.)

◆ If more than one plex of a mirrored volume is striped, configure the same stripe-unit size for each striped plex.

◆ Where possible, distribute the subdisks of a striped volume across drives connected to different controllers and buses.

◆ Avoid the use of controllers that do not support overlapped seeks. (Such controllers are rare.)

The `vxassist` command automatically applies and enforces many of these rules when it allocates space for striped plexes in a volume.

Mirroring Guidelines

Refer to the following guidelines when using mirroring. See *VERITAS Volume Manager Administrator's Guide* for more information.

◆ Do not place subdisks from different plexes of a mirrored volume on the same physical disk. This action compromises the availability benefits of mirroring and degrades performance. Using the `vxassist` command precludes this from happening.

◆ To provide optimum performance improvements through the use of mirroring, at least 70 percent of physical I/O operations should be read operations. A higher percentage of read operations results in even better performance. Mirroring may not provide a performance increase or may even result in performance decrease in a write-intensive workload environment.

Note The UNIX operating system implements a file system cache. Read requests can frequently be satisfied from the cache. This can cause the read/write ratio for physical I/O operations through the file system to be biased toward writing (when compared to the read/write ratio at the application level).

◆ Where possible, use disks attached to different controllers when mirroring or striping. Most disk controllers support overlapped seeks. This allows seeks to begin on two disks at once. Do not configure two plexes of the same volume on disks that are attached to a controller that does not support overlapped seeks. This is important for older controllers or SCSI disks that do not cache on the drive. It is less important for modern SCSI disks and controllers. Mirroring across controllers allows the system to survive a failure of one of the controllers. Another controller can continue to provide data from a mirror.

◆ A plex exhibits superior performance when striped or concatenated across multiple disks, or when located on a much faster device. Set the read policy to prefer the faster plex. By default, a volume with one striped plex is configured to prefer reading from the striped plex.

Dirty Region Logging (DRL) Guidelines

Dirty Region Logging (DRL) can speed up recovery of mirrored volumes following a system crash. When DRL is enabled, VxVM keeps track of the regions within a volume that have changed as a result of writes to a plex. VxVM maintains a bitmap and stores this information in a log subdisk. Log subdisks are defined for and added to a volume to provide DRL. Log subdisks are independent of plexes, are ignored by plex policies, and are only used to hold the DRL information.

Note Using Dirty Region Logging can impact system performance in a write-intensive environment.

Refer to the following guidelines when using DRL:

◆ For DRL to be in effect, the volume must be mirrored.

◆ At least one log subdisk must exist on the volume for DRL to work. However, only one log subdisk can exist per plex.

◆ The subdisk that is used as the log subdisk should not contain necessary data.

◆ Mirror log subdisks by having more than one log subdisk (but only one per plex) in the volume. This ensures that logging can continue, even if a disk failure causes one log subdisk to become inaccessible.

◆ Log subdisks must be configured with an even number of sectors. (The last sector in a log subdisk with an odd number of sectors is not used.) The log subdisk size should usually be configured as 33 sectors per 2 GB of volume. A log subdisk of 33 sectors is sufficient for a volume larger than 2 GB. In general, the default log subdisk length configured by the `vxassist` command is suitable for use.

◆ Do not place the log subdisk on a heavily-used disk, if possible.

◆ Use persistent (non-volatile) storage disks for log subdisks.

Mirroring and Striping Guidelines

Refer to the following guidelines when using mirroring and striping together. For more information, see the *VERITAS Volume Manager Administrator's Guide* for details.

◆ Make sure enough disks are available for the striped and mirrored configuration. At least two disks are required for the striped plex and one or more *additional* disks are needed for the mirror.

◆ Never place subdisks from one plex on the same physical disk as subdisks from the other plex.

Striping and Mirroring Guidelines

Refer to the following guidelines when using striping and mirroring together. For more information, see *VERITAS Volume Manager Administrator's Guide*.

◆ Make sure enough disks are available for the striped and mirrored configuration. At least two disks are required for the striped plex, and one or more *additional* disks are needed for the mirror.

◆ Never place subdisks from one plex on the same physical disk as subdisks from the other plex.

RAID-5 Guidelines

Refer to the following guidelines when using RAID-5. For more information, see *VERITAS Volume Manager Administrator's Guide*.

In general, the guidelines for mirroring and striping together also apply to RAID-5. The following guidelines should also be observed with RAID-5:

◆ Only one RAID-5 plex can exist per RAID-5 volume (but there can be multiple log plexes).

◆ The RAID-5 plex must be derived from at least two subdisks on two or more physical disks. If any log plexes exist, they must belong to disks other than those used for the RAID-5 plex.

◆ RAID-5 logs can be mirrored and striped.

◆ If the volume length is not explicitly specified, it is set to the length of any RAID-5 plex associated with the volume; otherwise, it is set to zero. If you specify the volume length, it must be a multiple of the stripe unit size of the associated RAID-5 plex, if any.

◆ If the log length is not explicitly specified, it is set to the length of the smallest RAID-5 log plex that is associated, if any. If no RAID-5 log plexes are associated, it is set to zero.

◆ Sparse RAID-5 log plexes are not valid.

Protecting Your System and Data

A disk failure can cause loss of data on the failed disk and loss of access to your system. Loss of access is due to the failure of a key disk used for system operations. VxVM can protect your system from these problems.

To maintain system availability, data important to running and booting your system must be mirrored. The data must be preserved so it can be used in case of failure.

The following are suggestions for protecting your system and data:

◆ Place the disk containing the root file system (the root or boot disk) under VxVM control through encapsulation. Encapsulation converts the root and swap devices to volumes (rootvol and swapvol).

◆ Mirror the root disk so that an alternate root disk exists for booting purposes. By mirroring disks critical to booting, you ensure that no single disk failure leaves your system unbootable and unusable.

For maximum availability of the system, create mirrors for the rootvol, swapvol, usr, and var volumes. For more information, see the *VERITAS Volume Manager Troubleshooting Guide*.

◆ Use mirroring to protect data against loss from a disk failure. To preserve data, create and use mirrored volumes that have at least two data plexes. The plexes must be on different disks. If a disk failure causes a plex to fail, the data in the mirrored volume still exists on the other disk.

When you use the vxassist mirror command to create mirrors, it locates the mirrors so the loss of one disk does not result in a loss of data. By default, the vxassist command does not create mirrored volumes; edit the file

`/etc/default/vxassist` to set the default layout to mirrored. For information on the `vxassist` defaults file, see *VERITAS Volume Manager Administrator's Guide* and the `vxassist`(1M) manual page.

◆ Leave the VxVM hot-relocation feature enabled to detect disk failures automatically. It will notify you of the nature of the failure, attempt to relocate any affected subdisks that are redundant, and initiate recovery procedures. Configure at least one hot-relocation spare disk in each disk group. This will allow sufficient space for relocation in the event of a failure.

If the `root` disk is mirrored, hot-relocation can automatically create another mirror of the `root` disk if the original `root` disk fails. The rootdg must contain enough contiguous spare or free space for the volumes on the root disk (`rootvol` and `swapvol` volumes require contiguous disk space).

◆ Use the DRL feature to speed up recovery of mirrored volumes after a system crash. Make sure that each mirrored volume has at least one log subdisk.

Note `rootvol`, `swapvol`, and `usr` volumes cannot be DRL volumes.

◆ Use logging to prevent corruption of recovery data in RAID-5 volumes. Make sure that each RAID-5 volume has at least one log plex.

◆ Perform regular backups to protect your data. Backups are necessary if all copies of a volume are lost or corrupted. Power surges can damage several (or all) disks on your system. Also, typing a command in error can remove critical files or damage a file system directly. Performing regular backups ensures that lost or corrupted data is available to be retrieved.

Uninstalling VxVM

Before uninstalling, you should move volumes from VxVM control. If the root volume is encapsulated, use the procedures in "Moving Volumes From an Encapsulated Root Disk" on page 77. If your root disk is not encapsulated, you can start with the procedures in "Moving Volumes To Disk Partitions."

Moving Volumes From an Encapsulated Root Disk

If `root`, `swap`, `usr`, or `var` is a volume under VxVM 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 of 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`, `swapvol`, `usr`, opt, home and `var` volumes.

   Continue to "Moving Volumes To Disk Partitions" on page 77 to complete the removal.

Moving Volumes To Disk Partitions

1. All volumes must be moved to disk partitions. This can be done using one of these procedures:

   - 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.

- Use VxVM to move volumes incrementally onto disk partitions. Use the following instructions to do this.

Move Volumes onto Disk Partitions Using VxVM

1. Evacuate disks using the `vxdiskadm` program, VMSA, or the `vxevac` script. You should consider the amount of target disk space required for this before you begin.

   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.

2. Remove the evacuated disks from VxVM control using the following commands:

   ```
   # vxdg rmdisk disk_name
   # 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 synced.

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 for the first volume removed. Subsequent volumes can use the free space generated by the removal of this volume.

6. `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 VxVM 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 VxVM using the following command:

   ```
   # vxedit -rf rm volume_name
   ```

10. Remove any disks that have become free (have no subdisks defined on them) by removing volumes from VxVM control. To check if there are still some subdisks remaining on a particular disk, use the following command:

    # **vxprint -F "%sd num" _disk_name_**

11. 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 VxVM control using the following commands:

    # **vxdg rmdisk _disk_name_**
    # **vxdisk rm _devname_**

12. The free space now created can be used for adding the data in the next volume to be removed.

13. 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**

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

Shutting Down VxVM

To shut down VxVM, use the following commands:

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

Removing the VERITAS Volume Manager Package

Before removing the packages, you should determine if any other VERITAS products are installed on your system. These products may depend on the VRTSlic package license that you may be removing.

To check for other installed VERITAS products:

    pkginfo | grep -i vrts

You get a warning when you try to remove VRTSlic if other products are using that license.

To remove the installed packages, use the following command:

    # **pkgrm VRTSvmdoc VRTSvmman VRTSvmdev VRTSvmsa VRTSvxvm VRTSlic**

If VxVM 2.x is still installed and you want to remove VxVM 3.2, 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 VRTSvmdoc VRTSvmman \
VRTSvmdev VRTSvxvm.2 VRTSvxvm VRTSlic
```

To verify the removal of the packages:

```
# pkginfo | grep VRTS
```

You can now remove `/VxVM*UPGRADE`, `/etc/vfstab.prevm`, `/etc/vfstab.vmsa-sav` and `.vmsa` from `/` and from each VMSA user's home directory.

If you do not have any other VERITAS products installed on the system, you can remove the `/etc/vx` directory, the `/usr/lib/vxvm` directory and the `/opt/VRTS*` directories.

This should complete the removal of the Volume Manager product.

Index