

VERITAS Volume Manager™ 3.1.1

Installation Guide

Solaris

February 2001
30-000224-011


VERITAS

Disclaimer

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

Copyright

Copyright © 1998-2001 VERITAS Software Corporation. All rights reserved. VERITAS is a registered trademark of VERITAS Software Corporation in the US and other countries. The VERITAS logo and VERITAS Volume Manager are trademarks of VERITAS Software Corporation. All other trademarks or registered trademarks are the property of their respective owners.

Printed in the USA, February 2001.

VERITAS Software Corporation
1600 Plymouth St.
Mountain View, CA 94043
Phone 650-527-8000
Fax 650-527-8050
www.veritas.com



Contents

Preface	vii
Introduction	vii
Audience	vii
Organization	vii
Downloading and Running VRTSexplorer	ix
Chapter 1. Introduction	1
Where to go Next	1
VxVM Basics	1
Easier Storage Management	1
Physical Disks	2
The VxVM Connection	2
VM Disks	2
Disk Groups	3
Subdisks	3
Plexes	4
Volumes	4
Storage Layouts	4
Arrays	5
Pre-installation Instructions	5
Verify Disk Space Requirements	5
Set the Path Variable	6
Verify Directories and Permissions	6
Verify Disk Contents	8



Hot-Relocation	10
VMSA	11
A5x00 Devices	11
Claiming Devices in SENA Category	12
SDS to VxVM Conversion	13
Adding Disks After Installation	14
Chapter 2. Installing the VERITAS Volume Manager	15
First Time Installation	15
Pre-installation Instructions	15
Loading the Software from CD-ROM	16
The Packages	16
Mounting the CD-ROM with Solaris Volume Management	16
Disabling Hot-Relocation	19
Running the vxinstall Program	19
Quick Installation	23
Custom Installation	28
Prevent Multipathing/Suppress Devices from VxVM's View	33
Rebooting After vxinstall	38
VMSA Installation Overview	39
Installing the VMSA Server and/or Client (Solaris)	40
Installing the VMSA Client (Windows)	41
Customizing Security (Optional)	41
Starting the VMSA Client	44
The Read-Only Mode for VMSA	45
Enabling Cluster Support in VxVM (Optional)	46
Configuring Shared Disks	46
Disk Controller Reservation	48
Upgrading to VxVM Release 3.1.1	49
The vxdmp Driver	49



Upgrading VxVM and Solaris	52
Upgrading VxVM Only	54
Upgrading VxVM with Alternate Pathing Enabled	57
Upgrading Solaris Only	57
Upgrading VMSA	58
Checking VxVM Processes	59
Uninstalling VxVM	59
Shutting Down VxVM	61
Starting and Enabling the Configuration Daemon vxconfigd	62
Starting the Volume I/O Daemon vxiod	63
Striping Guidelines	66
Mirroring Guidelines	67
Mirroring and Striping Guidelines	68
Striping and Mirroring Guidelines	69
RAID-5 Guidelines	69
Index	73





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 the VxVM.

Audience

This guide is for system administrators who configure and maintain UNIX systems with the 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 the 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 <code>/etc/vx/tunefstab</code> file. See the <code>ls(1)</code> manual page for more information.
<i>italic</i>	New terms, book titles, emphasis, variables to be replaced by a name or value	See the <i>User's Guide</i> for details. The variable <code>ncsize</code> determines the value of...
monospace (bold)	User input; the “#” symbol indicates a command prompt	<code># mount -F vxfs /h/filesys</code>
<i>monospace</i> (bold and italic)	Variables to be replaced by a name or value in user input	<code># mount -F <i>fstype mount_point</i></code>

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

Getting Help

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

- ◆ U.S. and Canadian Customers: 1-800-342-0652
- ◆ International Customers: +1 (650) 527-8555
- ◆ E-mail: `support@veritas.com`

For license information:

- ◆ 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`

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.

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 3.1.1](#)” on page 49.

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




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	19 MB
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.

Package	/	/usr	/opt
VRTSvxvm	17 MB	30 MB	0
VRTSvmdev	0	0	2 MB
VRTSvmman	0	0	1 MB
VRTSvmdoc	0	0	30 MB
VRTSvmsa	0	0	33 MB

Set the Path Variable

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

Examples:

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

```
$ PATH=/sbin:/usr/sbin:/etc/vx/bin:$PATH; export PATH
```

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

```
% setenv PATH /sbin:/usr/sbin:/etc/vx/bin:$PATH
```

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.

Obtain a License Key

During installation, you will be prompted for a license key, which is issued by VERITAS Customer Support.

Before contacting VERITAS, you need the host ID and the machine type.

1. Obtain the host ID. Enter:

```
# hostid
```

2. Obtain the host machine type. Enter:

```
# uname -i
```

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

3. 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. (See “[Getting Help](#)” on page ix for contact information.)

Upgrades

If you are upgrading VxVM, you can view the currently installed licenses with the following command:

```
# vxlicense -p
```

and

```
# /usr/lib/vxvm/bin/vxliccheck -PV
```

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



Take a Disk Inventory

Disk Types

Decide which disks you want to place under VxVM control. The other disks in your SAN 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 disk 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.)

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 100 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. Determine if you will encapsulate the boot (root) disk.
- a. Before encapsulating your boot disk, set the EEPROM variable `use-nvramrc?` to `true` to take advantage of VxVM boot disk aliases to identify the mirror of the boot disk if a replacement is needed. If this variable is set to `false`, you must determine which disks are bootable yourself. Set this variable to `true` as follows:


```
eprom "use-nvramrc?=true"
```
 - b. To encapsulate the boot disk, tag the swap partition as `swap` so that it is possible to dump to that partition later.



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 an `s2` slice that represents the full disk—this `s2` slice cannot contain a file system at the beginning or at the end of the disk
 - contain two free partitions
8. DMP considerations:
 - Will DMP be enabled?
 - If DMP will be enabled, will it co-exist with Sun's Alternate Pathing (AP)? See "[Alternate Pathing](#)" on page 12 and also the *VERITAS Volume Manager Administrator's Guide*.

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 "[Disabling Hot-Relocation](#)" on page 19 and "[Upgrading VxVM and Solaris](#)" on page 52 for details.



Solaris 7 Patch

If you are using Solaris 7 on your system, you must install Sun patch 106541 to ensure VxVM will operate properly.

Sun patches are available through the Sun web site, sunsolve.sun.com.

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, or 3.1.1). The `vxconfigd` and `VMSA_server` daemons must be running on the machine to be administered.
- ◆ The VMSA server. See [“Setting Up and Starting VMSA”](#) on page 39 on how to install the server.
- ◆ The VMSA 3.1.1 GUI client is not compatible with earlier versions of the VMSA server. When installing VMSA 3.1.1, 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 3.1.1”](#) on page 49 for details.

A5x00 Devices

If you have A5X00 devices connected to your system, then add the following Solaris patches before installing or upgrading to this release of VxVM:

Solaris 8—No patch is required.

Solaris 7—07473-03 or higher

Solaris 2.6—105375-20 or higher



Alternate Pathing

This release of VxVM will co-exist only with version 2.3.1 and patch 110722-01 and above of Sun's Alternate Pathing (AP) solution. Upgrade AP to 2.3.1 and patch 110722-01 before upgrading to this release of VxVM.

VxVM 3.1.1 will allow Dynamic Multipathing (DMP) to co-exist with Sun's AP software. For these features to co-exist, upgrade Sun's AP software to 2.3.1 and patch 110722-01 before installing 3.1.1.

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 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, the following patches must be installed before installing/upgrading VxVM.

Solaris 7—107473-03 or higher

Solaris 2.6—105375-20 or higher

If the required patches are not present on the system at the time of installation/upgrade to this version of VxVM, the A5x00 disk arrays will be claimed in the SEAGATE category. To claim these arrays in the SENA category, follow these steps:

1. Stop `vxconfigd`.

```
# vxconfigd -k
```

2. Add the required Solaris patch.

3. Copy `/usr/lib/liba5k.so.2` to `/etc/vx/slib`.

4. Copy `/usr/lib/libg_fc.so.2` to `/etc/vx/slib`.

5. Restart `vxconfigd`.

```
# vxconfigd
```

Cluster Environment Requirements

If your SAN 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 with VxVM in a cluster, consider leaving a small amount of space on the disk for these logs. The log size is proportional to the volume size and the number of nodes (each log has one recovery map plus one active map per node).

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

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

SDS to VxVM Conversion

If Solstice™ DiskSuite™ (SDS) is configured on your system, VERITAS tools are available to assist you in converting your system to VxVM. VERITAS has a conversion tool that enables a *data conversion in-place*. This tool allows user data to remain unchanged while SDS structural data is replaced with VxVM structural data.

To obtain and learn more about the VERITAS conversion tool, contact your sales representative.



Set up your 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 20 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 the VERITAS Volume Manager

2

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

- ◆ Package installation
- ◆ VxVM configuration and setup
- ◆ VMSA setup
- ◆ Shared disk setup (for cluster environments only)

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

First Time Installation

Pre-installation Instructions

To install VxVM and VMSA for the first time, use the following procedure:

1. This release of VxVM co-exists only with version 2.3.1 of Sun's Alternate Pathing (AP) driver. For VxVM to co-exist with AP successfully, follow these steps:
 - Upgrade to AP 2.3.1.
 - Install Solaris patch 110722-01.
 - Reboot the system.
 - Install to VxVM 3.1.1.

Note VxVM will not install if a previous version of AP is installed on the system.

2. Install the software packages onto the system to be administered as described in the section "[Loading the Software from CD-ROM](#)" on page 16.



3. Initialize VxVM using `vxinstall` as described in the section “[Running the vxinstall Program](#)” on page 19.
4. Set up and start VMSA as described in “[Setting Up and Starting VMSA](#)” on page 39.
5. If you are planning to use VxVM in a cluster environment, follow the instructions in “[Enabling Cluster Support in VxVM \(Optional\)](#)” on page 46.

Loading the Software from CD-ROM

The Packages

The VERITAS Volume Manager CD-ROM contains the following packages:

- ◆ `VRTSvxvm`—the VERITAS Volume Manager and Volume Replicator software.
- ◆ `VRTSvmdoc`—online copies of VERITAS Volume Manager guides.
- ◆ `VRTSvrdoc`—online copies of VERITAS Volume Replicator guides.
- ◆ `VRTSvmman`—the VxVM manual pages.
- ◆ `VRTSvmdev`—(optional) a VxVM developer’s package (library and header files).
- ◆ `VRTSvmsa`—the VERITAS Volume Manager Storage Administrator software.
- ◆ `VRTSvras`—the online copy of VERITAS Volume Replicator Administrator Services guide.

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 VRTSvxvm \  
VRTSvmsa VRTSvmdoc VRTSvmman VRTSvmdev
```

Note List the `VRTSvxvm` package before the `VRTSvmsa` package because the VxVM must be installed before VMSA.

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 AP version is prior to 2.3.1 and patch 110722-01, a warning similar to the following displays:

```
Alternate Pathing (AP) version 2.3.0 has been detected.  
This version of VxVM co-exists only with AP 2.3.1 along with  
following patch(es):
```

```
  110722-01
```

```
Please upgrade your AP package and install the above patch(es)  
before you continue the VxVM installation. This installation will  
now be aborted.
```

If AP 2.3.1 is installed on the system but the Solaris patch 110722-01 is not installed, the following warning displays:

```
Alternate Pathing (AP) version 2.3.1 has been detected.  
This version of VxVM requires the following patch(es) to co-exist  
with AP 2.3.1:
```

```
  110722-01
```

```
Install the above patch(es) before you continue the VxVM  
installation. This installation will now be aborted.
```

If the DMP driver was disabled on the system prior to this upgrade, then the following message displays:

```
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.



Mounting the CD-ROM at the Command Line

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

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, type:

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

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

4. After the CD-ROM is mounted, install the packages. Type:

```
# pkgadd -d /mnt/pkg VRTSvxvm VRTSvmsa \  
VRTSvmdoc VRTSvman VRTSvmdev
```

Note List the `VRTSvxvm` package before the `VRTSvmsa` package because VxVM must be installed before VMSA.

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 similar to the following displays:

```
Alternate Pathing (AP) version 2.3.0 has been detected.  
This version of VxVM co-exists only with AP 2.3.1 along with  
following patch(es):
```

```
110722-01
```

```
Please upgrade your AP package and install the above patch(es)  
before you continue the VxVM installation. This installation will  
now be aborted.
```

If AP 2.3.1 is installed on the system but the Solaris patch 110722-01 is not installed, the following warning displays the following:

```
Alternate Pathing (AP) version 2.3.1 has been detected.  
This version of VxVM requires the following patch(es) to co-exist  
with AP 2.3.1:
```

```
110722-01
```

```
Install the above patch(es) before you continue the VxVM  
installation. This installation will now be aborted.
```

If the DMP driver was disabled on the system prior to this upgrade, then the following message displays:

```
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.

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 &
```

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 repeated reboots. You are usually asked for permission to reboot when the reboots are needed. If you do not reboot, the initialization does not complete. After the first reboot, the initialization process takes control of the system until the installation is complete.



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 the VxVM using the following procedure:

1. Log in as superuser.
2. The `vxinstall` and `vxdiskadm` utilities use the files `/etc/vx/cntrls.exclude` and `/etc/vx/disks.exclude` to automatically exclude controllers or disks so that they are not configured as VxVM devices. These files do not exclude controllers and disks from use by any other VxVM commands. See the `vxinstall(1M)` and `vxdiskadm(1M)` manual pages for more information.
 - a. To exclude one or more disks from VxVM control, create the `/etc/vx/disks.exclude` file and add the names of those disks to the file. The `vxinstall` program ignores any disks listed in this file and leaves them untouched.
 - If a `disks.exclude` file is not created, then the `vxinstall` program will prompt you about each disk.
 - In prior releases of VxVM, one or more disks could be added to `/etc/vx/disks.exclude` in order to make `vxinstall` and `vxdiskadm` ignore those disks. You could also make `vxinstall` and `vxdiskadm` ignore all disks on an entire controller by adding the device names of the controllers in `/etc/vx/cntrls.exclude`.

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

```
c0t1d0
```

- b. Create a `cntrls.exclude` file if you want to exclude all disks on a controller from VxVM control. Place this file in: `/etc/vx/cntrls.exclude`.

Add the name of that controller to the `/etc/vx/cntrls.exclude` file.

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

```
c0
c1
```

3. 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.

- If these files exist on a system prior to upgrading to this release of VxVM, their effect will still be visible. However, from this release 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.

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

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

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

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

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

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

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

Some examples would be:



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

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

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

4. Press Return to continue.

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.

5. 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.

6. Select an option.

- If you choose 1, see [“Quick Installation”](#) on page 23.
- If you choose 2, see [“Custom Installation”](#) on page 28.
- If you choose 3, see [“Prevent Multipathing/Suppress Devices from VxVM's View”](#) on page 33.

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 on a controller, even if valid partitions are found on those disks. However, avoid selecting the initialize option for any disks containing valid data.
- ◆ Quick Installation will bring every disk attached to your system under VxVM control, unless you have chosen to exclude those disks. Disks could be chosen for exclusion either through the `disks.exclude` or `cntrls.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 option `Prevent multipathing/Suppress devices` from VxVM's view in `vxinstall` or Custom Installation option (which allows you to leave certain disks alone).

- ◆ To encapsulate some disks on a given controller and initialize others, use Custom Installation instead. (See page 28.)
- ◆ With encapsulation, the `/etc/vfstab` file is updated to ensure that file systems previously mounted on disk partitions are mounted on volumes instead.
- ◆ All disks are encapsulated or initialized (according to your instructions) at the end of the `vxinstall` procedure. If you quit the `vxinstall` program before it enters the final phase of initializing or encapsulating the disks, all disks are left as they were before the `vxinstall` program started.

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

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

```
Generating list of attached disks on c1....
The Volume Manager has detected the following disks on controller
c1:
  c1t0d0
  c1t1d0
  c1t2d0
  c1t3d0
  c1t4d0

Hit RETURN to continue.
```

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

3. Press Return to continue.
4. The following message displays. Indicate whether you want to initialize all of the disks on the controller.

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

```
Volume Manager Quick Installation For Controller c1 Menu:
VolumeManager/Install/QuickInstall/c1
Initialize all disks on this controller? (destroys data on these
disks) [y,n,q,?] (default: n)
```

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

```
Volume Manager will now try to encapsulate all the disk on this
controller.
Disks not having valid partitions will be initialized.
Hit RETURN to continue.
```

Press Return to continue.

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

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



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

```
Volume Manager will now initialize all the disks on this
controller 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 controller at the following prompt:

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

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

```
The c1t0d0 disk will be given disk name disk01
The c1t1d0 disk will be given disk name disk02
The c1t2d0 disk will be given disk name disk03
The c1t3d0 disk will be given disk name disk04
Hit RETURN to continue.
```

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

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

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

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

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

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

6. For each disk, respond with the desired disk name and press Return.
7. When all of the disks on the current controller 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 controllers on your system, the `vxinstall` program displays a summary of the disks designated for initialization (New Disk) or encapsulation (Encapsulate) on each controller, as follows:

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

```
c0t5d0  New Disk
c1t0d0  New Disk
```

```

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

```

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

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 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 38.

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 on a controller, encapsulating all disks on a controller, or initializing some disks on a controller 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 controller 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 and all other disk partitions found on your boot disk. The `/usr`, `/opt`, and `/var` file systems, and any other file systems on your boot disk, are also encapsulated.

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

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

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

```

c1t0d0
c1t1d0
c1t2d0
c1t3d0
c1t0d4

```

Hit RETURN to continue.

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

3. Press Return to continue.
4. Indicate how you want the disks on this controller to be handled at the following prompt:

```

Installation options for controller c1
Menu: VolumeManager/Install/Custom/c1

 1 Install all disks as pre-existing disks. (encapsulate)
 2 Install all disks as new disks.(discards data on disks!)
 3 Install one disk at a time.
 4 Leave these disks alone.

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

```

Select an operation to perform:

Encapsulating All Disks

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

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

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

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

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

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



```
The c1t0d0 disk has been configured for encapsulation.  
The c1t1d0 disk will be given disk name disk02  
The c1t1d0 disk has been configured for encapsulation.  
Hit RETURN to continue.
```

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

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

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

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

3. For each disk, enter the desired disk name and press Return.
4. When all of the disks on the current controller are named, press Return to move on to the next controller.

Initializing All Disks

1. To initialize all disks on the controller, select 2.

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

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

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

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

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

```
The c1t0d0 disk will be given disk name disk01  
The c1t1d0 disk will be given disk name disk02  
Hit RETURN to continue.
```


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

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

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

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

Installing Individual Disks

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

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

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

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

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

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

Select an operation to perform:

These options are similar to those in the Custom Installation main menu. Once you have determined how to handle this disk, see [step 4](#) on page 29.

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



Leaving All Disks Unaltered

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

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

2. Repeat [step 4](#) on page 29 for each of the remaining controllers on your system.

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

The following is a summary of your choices.

```
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**.

- 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 38.

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 disable multipathing.

- 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
```

- Enter `y`.
- 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:[ctrl_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 vxdumping 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/vxdmping.

- Option 4 defines a pathgroup in case of 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 `clt30d0` and `c2t30d0` are paths to the same disk and both are seen by VxVM as separate disks, `clt30d0` 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, `clt30d0 c2t30d0`.

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

The next three options allow you to exclude devices from `vxddmp`.

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

```
Exclude controllers from VxDMP
```

```
Menu: VolumeManager/Install/ExcludeDevices/CTRL-DMP
```

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

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 `vxddmp`. 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:[<ctrl-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 VxDMP
```

```
Menu: VolumeManager/Install/ExcludeDevices/PATH-DMP
```

Use this operation to exclude one or more disks from `vxddmp`.

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 option 1 through 7, you must reboot the system for device exclusion to take effect.

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 shutdown and reboot is required. If you chose to encapsulate any disks, a reboot will be necessary.

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

Note The setup you choose can require several reboots.

When it is necessary to shut down and reboot your system, the `vxinstall` program displays a message similar to the following:

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

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

If this message is displayed, respond in one of the following ways:

- Enter **y** to begin an immediate shutdown.
- Enter **n** to cause the `vxinstall` program to exit without starting a shutdown. If you select this option, shut down and 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.
- ◆ VERITAS Volume Manager Release 2.5.7, 3.0.4 3.1, or 3.1.1. The `vxconfigd` daemon must be running on the machine to be administered.
- ◆ The VMSA 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 58.

Note This release of VMSA requires VxVM Release 2.5.7, 3.0.4, 3.1 or 3.1.1. 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 VMSA server and client on the UNIX (Solaris) machine to be administered. See “[Installing the VMSA Server and/or Client \(Solaris\)](#)” on page 40.
2. If users other than `root` need to access VMSA, set up security and specify which users can run VMSA. See “[Customizing Security \(Optional\)](#)” on page 41.
3. If you plan to run the VMSA client from 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 40 or “[Installing the VMSA Client \(Windows\)](#)” on page 41.



4. Start the VMSA server on the machine to be administered. See “[Starting the VMSA Server](#)” on page 43.
5. Start the VMSA client. See “[Starting the VMSA Client](#)” on page 44.

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 16.

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 3.1.1 GUI client is not compatible with earlier versions of the VMSA server. When installing VMSA 3.1.1, 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. Go to the directory where you installed VMSA (/opt/VRTS/vmsa, by default).
4. Set the VMSAHOME environment variable using the following command:

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

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

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

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

Installing the VMSA Client (Windows)

The VMSA client runs on Windows NT, Windows 2000, Windows 98, or 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 and client on a UNIX (Solaris) machine.

Before you install VMSA on a Windows machine, you must uninstall any existing VERITAS Volume Manager Storage Administrator packages on that machine. Only one 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 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 any 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 `java` command at the end of the `/opt/VRTSvmsa/bin/vmsa_server` file:

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

Encryption

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

```
vrts.security=false
```

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.

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 &
```

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 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 and the server is not already running. If no VMSA clients are connected for a certain amount of time (15 minutes, by default), the server automatically terminates and restarts at the next client connection.

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 a VMSA service has 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 41 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)

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 may 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 16.
3. Initialize VxVM using the `vxinstall` procedure as described in “[Running the vxinstall Program](#)” on page 19.
4. Set up and start VMSA as described in “[Setting Up and Starting VMSA](#)” on page 39.
5. Configure shared disks as described in “[Configuring Shared Disks](#)” on page 46.

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 48).

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

Configuring New Disks

If you are installing and setting up VxVM for the first time, configure the shared disks using the following procedure:

1. Start the cluster on at least one node.

2. Run the `vxdisksetup` program to initialize each shared disk.

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

3. Create disk groups on the shared disks.

Use the `vxdbg` program or VMSA to do this. In the `vxdbg` program, use the `-s` option of to create shared disk groups.

4. Create volumes in the disk groups.

Use the `vxassist` program or the VMSA to do this.

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

5. If the cluster is only running with one node, bring up the other cluster nodes. Enter the `vxdbg list` command on each node to display the shared disk groups.

Verifying Existing Shared Disks

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

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

```
# vxdbg list
```

This displays the shared disk groups that existed earlier.



Converting Existing VxVM Disk Groups to Shared Disk Groups

If you are upgrading from VxVM 2.x to VxVM 3.1.1 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 list disk-group-name
```

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

```
# 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 3.1.1

If you are running a release earlier than 3.1.1, you can upgrade with one of the following procedures, which explain how to upgrade VxVM, Solaris, or both. This section also contains information on how to upgrade VMSA.

No new license is required for upgrading.

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

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

Note Importing a pre-3.x VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.x level. You may need to manually upgrade each of your disk groups following a VxVM upgrade. See “Upgrading a Disk Group” in the *VERITAS Administrator’s Guide* for information on how to upgrade a disk group.

The vxdump Driver

Starting with this release of VxVM, the `vxdump` driver must always be present on the system for VxVM to function. Upgrading to this release of VxVM enables `vxdump` even if it was disabled prior to this upgrade.

Adding Patches

If you have A5x00 devices connected to your system, then add the following Solaris patches before upgrading/installing to this release of VxVM:

- ◆ Solaris 7—107473-03 or higher
- ◆ Solaris 2.6—105375-20 or higher

This release of VxVM will co-exist only with versions 2.3.1 and patch 110722-01 of Sun’s AP solution. Upgrade AP to 2.3.1 and install the Solaris patch 110722-01 *before* upgrading to this release of VxVM.



The table below indicates which VxVM releases are supported on each Solaris release.

Solaris and VxVM Upgrade Matrix

	Solaris 2.3	Solaris 2.4	Solaris 2.5	Solaris 2.5.1	Solaris 2.6	Solaris 7	Solaris 8
VxVM 3.1.1					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 2.5.4			Supported	Supported	Supported		
VxVM 2.5.3			Supported	Supported	Supported		
VxVM 2.5.2			Supported	Supported	Supported		
VxVM 2.5.1			Supported	Supported	Supported		
VxVM 2.5			Supported	Supported	Supported		
VxVM 2.4.1		Supported	Supported	Supported	Supported		
VxVM 2.4		Supported	Supported	Supported	Supported		



Solaris and VxVM Upgrade Matrix (continued)

	Solaris 2.3	Solaris 2.4	Solaris 2.5	Solaris 2.5.1	Solaris 2.6	Solaris 7	Solaris 8
VxVM 2.3		Supported	Supported	Supported			
VxVM 2.2.1		Supported	Supported	Supported			
VxVM 2.2		Supported	Supported				
VxVM 2.1	Supported	Supported					
VxVM 2.0	Supported						

The following table shows possible VxVM/Solaris upgrade combinations. After reviewing this table, proceed to the appropriate section.

Upgrade Reference

To perform these upgrades	Go to this section
VxVM and Solaris:	“Upgrading VxVM and Solaris” on page 52.
VxVM 2.1 & Solaris 2.4 ► VxVM 3.1.1 & Solaris 2.6/7/8	
VxVM 2.2 & Solaris 2.4/2.5 ► VxVM 3.1.1 & Solaris 2.6/7/8	
VxVM 2.2.1 & Solaris 2.4/2.5/2.5.1 ► VxVM 3.1.1 & Solaris 2.6/7/8	
VxVM 2.3 & Solaris 2.4/2.5/2.5.1 ► VxVM 3.1.1 & Solaris 2.6/7/8	
VxVM 2.4/2.4.1 & Solaris 2.4/2.5/2.5.1 ► VxVM 3.1.1 & Solaris 2.6/7/8	
VxVM 2.5/2.5.x & Solaris 2.5/2.5.1/2.6 ► VxVM 3.1.1 & Solaris 2.6/7/8	
VxVM 3.0 & Solaris 2.5.1/2.6 ► VxVM 3.1.1 & Solaris 2.6/7/8	
VxVM 3.0.x & Solaris 2.5.1/2.6/7 ► VxVM 3.1.1 & Solaris 2.6/7/8	



Upgrade Reference (continued)

VxVM Only:	“Upgrading VxVM Only” on page 54.
VxVM 2.4/2.4.x & Solaris 2.6 ► VxVM 3.1.1 & Solaris 2.6	
VxVM 2.5/2.5.x & Solaris 2.6 ► VxVM 3.1.1 & Solaris 2.6	
VxVM 3.0 & Solaris 2.6 ► VxVM 3.1.1 & Solaris 2.6	
VxVM 3.0.x & Solaris 2.6/7/8 ► VxVM 3.1.1 & Solaris 2.6/7/8	
Solaris Only:	“Upgrading Solaris Only” on page 57.
VxVM 3.1.1 & Solaris 2.6 ► VxVM 3.1.1 & Solaris 7/8	
VxVM 3.1.1 & Solaris 7 ► VxVM 3.1.1 & Solaris 8	

Upgrading VxVM and Solaris

This section explains how to upgrade both Solaris and VxVM. (Refer to the [Table](#) on page 50 to determine if upgrading the operating system is required.)

Note You must use this procedure if you are upgrading encapsulated root disks. If your root disk is not encapsulated, refer to [“Upgrading VxVM Only”](#) on page 54.

Follow these steps:

1. If you have not already obtained and installed a VxVM 3.1.1 license key, do so now. See [“Obtain a License Key”](#) on page 7 for details.
2. If any of the file systems `/`, `/usr`, `/var`, or `/opt` are defined on volumes, make sure that at least one plex for each of those volumes is formed from a single subdisk that begins on a cylinder boundary.

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

3. Unmount the file systems that VxVM controls (`/`, `/usr`, `/var`, and `/opt`).

4. Load and mount the CD-ROM as described in “[Loading the Software from CD-ROM](#)” on page 16. The CD-ROM should then be visible as the file system mounted on `/CD_path`. If the CD-ROM is mounted automatically, its path is `/cdrom/CD_name`; if mounted manually, the path is `/mnt`.

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

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

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

6. Reboot to single-user mode (using a command such as `/etc/shutdown`).
7. 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.
8. If VxVM is installed, remove the old VxVM packages by entering:

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

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 `SUNWvxvm` package.

10. Shut down and halt the machine by using a command such as `/etc/shutdown`.
11. Reboot the machine to single-user mode using a command such as `/etc/shutdown`. If you are *not* upgrading Solaris, this reboot will complete the removal of the old version of VxVM and leave the system in a state where the new version of VxVM can be installed.
12. Remount the CD-ROM manually using the following command:

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



13. Add the VxVM 3.1.1 package using the following command:

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

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

14. Run a reboot with the `-r` option:

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

15. Complete the upgrade using the following command:

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

If you are upgrading from VxVM 1.3, this command prompts you for licensing information.

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

```
# vxrelocd root &
```

17. Reboot to multi-user mode. At this point, your pre-upgrade configuration is in effect and any file systems previously defined on volumes are defined and mounted.

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

19. Importing a pre-3.x VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.1.1 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.

20. See “[Setting Up and Starting VMSA](#)” on page 39 to set up and restart VMSA.

Upgrading VxVM Only

If you are already running a version of Solaris that is supported with VxVM 3.1.1, you can upgrade VxVM without upgrading Solaris. You will completely remove VMSA, install VxVM, and reinstall VMSA.

Caution If your root disk is encapsulated, you cannot follow the procedure in this section. Instead, follow the steps in the section, “[Upgrading VxVM and Solaris](#)” on page 52.

Upgrade VxVM as follows:

1. Log in as superuser.
2. Remove the VMSA software using the following command:


```
# pkgrm VRTSvmsa
```
3. Add the VxVM 3.1.1 package using the following command:


```
# pkgadd -d CD_path/pkgs/VRTSvxvm
```
4. Using the `pkgadd` command, install any or all of the additional packages: `VRTSvmsa`, `VRTSvmdoc`, `VRTSvmman`, and `VRTSvmdev`.
5. If you are prompted to run a reboot with the `-r` option, type:


```
# reboot -- -r
```
6. See “[Setting Up and Starting VMSA](#)” on page 39.

Note The VMSA Release 3.1.1 server is not backward compatible with VMSA Release 1.x clients. Upgrade any existing client(s) to Release 3.1.1.

7. Depending on your existing configuration, various messages and prompts may result.
8. If AP is installed on the system and the AP version is prior to 2.3.1, a warning similar to the following displays:

```
Alternate Pathing (AP) version 2.3.0 has been detected.
This version of VxVM co-exists only with AP 2.3.1 along with
following patch(es):
110722-01
Please upgrade your AP package and install the above patch(es)
before you continue the VxVM installation. This installation will
now be aborted.
```

If AP 2.3.1 is installed on the system but the Solaris patch 110722-01 is not installed, the following warning displays:

```
Alternate Pathing (AP) version 2.3.1 has been detected.
This version of VxVM requires the following patch(es) to co-exist
with AP 2.3.1:
110722-01
Install the above patch(es) before you continue the VxVM
installation. This installation will now be aborted.
```



If the DMP driver was disabled on the system prior to this upgrade, then the following message displays:

```
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.

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

```
Checking for conflicts with packages already installed.
The following files are already installed on the system and are being
used by another package:
```

```
files...
```

```
Do you want to install these conflicting files [y,n,?,q]
```

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

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

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

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

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

11. 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 VxVM with Alternate Pathing Enabled

DMP co-exists with AP driver 2.3.1 and patch 110722-01. If a previous version of AP is installed on your system, upgrade to AP version 2.3.1 and patch 110722-01 before upgrading VxVM.

If AP is installed and the AP version is prior to 2.3.1, a warning similar to the following displays:

```
Alternate Pathing (AP) version 2.3.0 has been detected.
This version of VxVM co-exists only with AP 2.3.1 along with
following patch(es):
110722-01
Please upgrade your AP package and install the above patch(es)
before you continue the VxVM installation. This installation will
now be aborted.
```

If AP 2.3.1 is installed on the system but the Solaris patch 110722-01 is not installed, the following warning displays:

```
Alternate Pathing (AP) version 2.3.1 has been detected.
This version of VxVM requires the following patch(es) to co-exist
with AP 2.3.1:
110722-01
Install the above patch(es) before you continue the VxVM
installation. This installation will now be aborted.
```

Upgrading Solaris Only

If you are running VxVM 3.1.1 with Solaris 2.6/7/8, 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`, or `/opt` are defined on volumes, ensure that at least one plex for each of those volumes is formed from a single subdisk that begins on a cylinder boundary. This is necessary because part of the upgrade process involves temporarily converting file systems on volumes back to using direct disk partitions, and Solaris requires that disk partitions start on cylinder boundaries.

The upgrade scripts (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. In the `/etc/vfstab` file, comment out any VxVM file systems.



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

If the CD-ROM is mounted automatically, enter:

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

If the CD-ROM is mounted manually, enter:

```
# /mnt/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.

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

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

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 3.1.1 server is not backward compatible with VMSA Release 1.x clients; therefore, upgrade any existing client(s) to Release 3.1.1.

1. Log in as superuser.
2. Remove the previous `VRTSvmsa` package using the following command:

```
# pkgrm VRTSvmsa
```

- See “[Loading the Software from CD-ROM](#)” on page 16 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:

```
root45818010:23:34?0:00vxrelocd root
root41418010:23:06?0:03vxconfigd
root4784583910:23:37?0:00vxrelocd root
root479478 510:23:37?0:00vxnotify -f -w 15
```

Uninstalling VxVM

There are two methods to uninstall VxVM.

If `root`, `swap`, `usr`, or `var` is a volume under VxVM control, use the following procedure.

- 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.

- 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`, `swap`, `usr`, and `var` volumes.

3. Once the `root`, `swap`, `usr`, and `var` volumes have been completed, all remaining volumes must be moved to disk partitions.

This can be done using one of three procedures, as follows:

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

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

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

```
# vxdg rmdisk disk_name
# vxdisk rm devname
```

- c. Decide which volume to move first. If the volume to be moved is mounted, unmount it.
- d. If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that data on the volume is synced.
- e. Create a partition on free disk space of the same size as the volume. If there is not enough free space for the partition, a new disk must be added to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this volume.
- f. Copy the data on the volume onto the newly created disk partition using a command similar to the following:

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

where `c2t2d2` is the disk outside of VxVM and `s7` is the newly created partition.

- g. Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly created partition.
- h. Mount the disk partition if the corresponding volume was previously mounted.
- i. Remove the volume from the VxVM using the following command:

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

- j. Remove any disks that have become free (have no subdisks defined on them) by removing volumes from 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
```

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
```

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

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

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

Shutting Down VxVM

To shut down VxVM, use the following commands:

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



Removing the VERITAS Volume Manager Package

To remove the installed packages, use the following command:

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

If VxVM 2.x is still installed and you want to remove VxVM 3.1.1, it is necessary to remove the VxVM 2.x package first. To determine whether both packages are present, use the following command:

```
# pkginfo | grep vxvm
```

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

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

To verify the removal of the packages:

```
# pkginfo | grep VRTS
```

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:

```
# vxctl 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 vxctl(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.

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.





Index

- A**
 - A5x00 devices 11
 - AP (alternate pathing) 12
 - arrays 5
 - auto-start mode, VMSA 43
- B**
 - boot disk 20
 - encapsulated 23
 - encapsulation requirements 9
 - restrictions 8
- C**
 - cluster functionality
 - enabling 46
 - environment requirements 13
 - shared disks 46
 - cluster requirements 13
 - cntrls.exclude file 21
 - configuration daemon (vxconfigd), starting 62
 - configuring
 - new disks 47
 - shared disks 46
 - controllers
 - excluding disks 21
 - controllers, mirroring guidelines 67
 - custom installation method 28
- D**
 - daemons, checking 59
 - data
 - protecting 70
 - data conversion tool 13
 - deleting VxVM packages 62
 - devices
 - suppress devices 33
 - disk arrays 5
 - disk drives
 - variable geometry 66
 - disk groups 3
 - disks
 - initializing all 23
 - initializing selected 28
 - placing in other disk group 20
 - rootdg 20
 - specifying disks without VxVM control 20
 - stripe unit size 66
 - types of 8
 - disks.exclude file 20
 - DMP
 - prevent multipathing 33
 - DMP (dynamic multipathing)
 - AP co-existence 12
 - preventing 33
 - SENA category devices 12
 - driver, VxDMP 49
 - DRL guidelines 68
- E**
 - encapsulated disks 8
 - encapsulating
 - all disks 29
 - boot disks 23
 - excluding
 - disks from VxVM control 20
 - disks on a controller 21
- G**
 - guidelines
 - DRL 68
 - mirroring 67
 - mirroring and striping 68
 - RAID-5 69
- H**
 - hot-relocation 10, 19



-
- I**
 - I/O daemon (vxiod), starting 63
 - initialized disks 8
 - initializing
 - all disks 30
 - installing
 - individual disks 31
 - VMSA client (Windows) 41
 - VMSA server (UNIX) 40
 - L**
 - layouts, types of 4
 - license keys, obtaining 7
 - license, viewing 7
 - log subdisks 68
 - logical objects 3
 - M**
 - mirroring
 - and striping guidelines 68
 - guidelines 67
 - root disk 70
 - mirroring controllers 67
 - N**
 - new disks, configuring 47
 - O**
 - overlapped seeks 67
 - P**
 - patches
 - A5x00 devices 11
 - Solaris 7 operating system 11
 - path variable 6
 - pkgrm command 62
 - plexes 4
 - Prevent Multipathing/Suppress Devices from VxVM's view 33
 - preventive maintenance 70
 - private region 3
 - processes, checking 59
 - protecting data 70
 - public region 2
 - Q**
 - quick installation method 23
 - R**
 - RAID-5
 - guidelines 69
 - removing VxVM packages 62
 - replacing SDS structural data 13
 - root disk
 - alternate 70
 - rootdg 20
 - rootdg (root disk group) 3
 - S**
 - SANs, configuring 14
 - SDS 13
 - security, VMSA user access 41
 - shared disks, configuring 46
 - shutting down VxVM 61
 - starting vxconfigd configuration daemon 62
 - starting vxiod daemon 63
 - storage layouts 4
 - storage management 1
 - stripe unit size recommendations 66
 - striping
 - and mirroring guidelines 69
 - striping guidelines 66
 - subdisks 3
 - physical disk placement 68
 - RAID-5 plex, configuring 69
 - suppress devices 33
 - swap relocation 9
 - U**
 - upgrading
 - VMSA 58
 - VxVM and Solaris 52
 - VxVM only 54
 - VxVM with AP 57
 - upgrading to 3.1.1 49
 - V**
 - VM disks 2
 - VMSA
 - auto-start mode 43
 - client, starting 44
 - installation requirements 11
 - read-only mode 45
 - setting user access 41
 - upgrading 58
 - VMSA client (UNIX), installing 40
 - VMSA client (Windows), installing 41
 - VMSA server, starting 43
 - volume length, RAID-5 guidelines 69
 - volumes 4
 - DRL 68
 - vrtsadm group 41
 - VRTExplorer ix

vxassist mirror command 70
vxconfigd configuration daemon, starting 62
vxdctl mode command 62
VxDMP driver 49
vxinstall 21
vxinstall program, running 20
vxiod I/O daemon, starting 63
VxVM
 obtaining system information ix
VxVM packages, removing 62

