

VERITAS Volume Manager™ 3.5

Installation Guide

Solaris

July 2002
N08835F


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 © 2002 VERITAS Software Corporation. All rights reserved. VERITAS, VERITAS SOFTWARE, the VERITAS logo and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation in the USA and/or other countries. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
Phone 650-527-8000
Fax 650-527-2908
www.veritas.com



Contents

Preface	vii
Introduction	vii
Audience	vii
Organization	vii
Conventions	viii
Getting Help	ix
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
New VERITAS Product Distribution Method	5
Verify Disk Space Requirements	5



Set the PATH and MANPATH Variables	7
Verify Directories and Permissions	7
Product Licensing	7
Using the VERITAS vLicense Website to Obtain a License	8
Add or Upgrade Licenses	9
Take a Disk Inventory	10
Verify Disk Contents	10
Hot-Relocation	12
Solaris Patches for VxVM	13
VERITAS Enterprise Administrator (VEA)	13
A5x00 Devices	13
DMP and Alternate Pathing	14
Claiming Devices in the SENA Category	14
Cluster Environment Requirements	14
Set up your SAN Configuration	15
Adding Disks After Installation	15

Chapter 2. Installing VERITAS Volume Manager 17

Installation and Upgrade Overview	17
Important Software Patch Information	18
Solaris Patches for VxVM	18
Installing or Upgrading VxVM with Alternate Pathing Enabled	20
Loading the Software from CD-ROM	21
The Packages	21
Mounting the CD-ROM with Solaris Volume Management	21
Mounting the CD-ROM from the Command Line	23
Installing VxVM Using the Installation Menu	23
Running the vxinstall Program	24
Quick Installation	29
Custom Installation	35

Prevent Multipathing/Suppress Devices from VxVM's View	40
Rebooting After vxinstall	45
Setting Up and Starting VEA	46
VMSA and VEA Co-Existence	46
VEA Installation Overview	47
Installing the VEA Server and/or Client (Solaris)	47
Installing the VEA Client (Windows)	48
Starting VEA	49
Starting the VEA Client	49
Enabling Cluster Support in VxVM (Optional)	50
Configuring Shared Disks	51
Disk Controller Reservation	53
Upgrading to VxVM Release 3.5	54
Before You Upgrade	54
VxVM and Solaris Release Support Matrix	56
VxVM and Solaris Upgrade Paths	57
Upgrading VxVM on an Encapsulated Root Disk	58
Upgrading VxVM and Solaris	63
Upgrading VxVM Only	66
Upgrading Solaris Only	71
Checking VxVM Processes	72
Starting and Enabling the Configuration Daemon vxconfigd	73
Starting the Volume I/O Daemon vxiod	74
System Setup Guidelines	74
Disabling Hot-Relocation	75
Hot-Relocation Guidelines	75
Striping Guidelines	77
Mirroring Guidelines	78
Dirty Region Logging (DRL) Guidelines	78
Mirroring and Striping Guidelines	79



RAID-5 Guidelines	79
Protecting Your System and Data	80
Uninstalling VxVM	82
Moving Volumes From an Encapsulated Root Disk	82
Moving Volumes To Disk Partitions	82
Shutting Down VxVM	84
Removing the VERITAS Volume Manager Packages	84
Restoring VERITAS Volume Manager 3.5 Configuration	86
Chapter 3. Moving Volumes onto Partitions - Example	87
Index	93



Preface

Introduction

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

Audience

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

- ◆ An understanding of system administration.
- ◆ A working knowledge of the UNIX operating system.
- ◆ A general understanding of file systems.

Organization

This guide is organized with the following chapters:

- ◆ [Introduction](#)
- ◆ [Installing VERITAS Volume Manager](#)
- ◆ [Moving Volumes onto Partitions - Example](#)

For information about technical support, product licensing, software updates, and other useful links, see the section called “[Getting Help](#)” on page ix.



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 <i>ncsize</i> determines the value of...
monospace (bold)	User input; the “#” symbol indicates a command prompt	# mount -F vxfs /h/filesys
<i>monospace</i> (bold and italic)	Variables to be replaced by a name or value in user input	# mount -F <i>fstype</i> <i>mount_point</i>

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

Note “Dynamic Disk Groups” were formerly known as “Disk Groups.” They are still sometimes referred to as “Disk Groups” in documentation, menu displays, and the CLI.

Getting Help

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

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

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

- ◆ Phone: 1-925-931-2464
- ◆ Email: license@veritas.com
- ◆ Fax: 1-925-931-2487

For software updates:

- ◆ Email: swupdate@veritas.com

For information on purchasing VERITAS products:

- ◆ Phone: 1-800-258-UNIX (1-800-258-8649) or 1-650-527-8000
- ◆ Email: vx-sales@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. Refer to the README file in the `support` directory on the VERITAS Installation CD, and the *VERITAS Volume Manager Release Notes* for information on using the VRTSexplorer program.





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 VERITAS Enterprise Administrator (VEA) 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 3.5](#)” on page 54.

VxVM Basics

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

Easier Storage Management

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

VxVM enhances data storage management by controlling these aspects:

- ◆ space—allocation and use
- ◆ performance—enhanced data delivery



- ◆ data availability—continuous operation and multisystem access
- ◆ device installation—centralized and optimized support
- ◆ system—multisystem support and monitoring of private/shared systems

Physical Disks

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

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

The VxVM Connection

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


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

VxVM objects include the following:

- ◆ VM Disks
- ◆ Disk Groups
- ◆ Subdisks
- ◆ Plexes
- ◆ Volumes

VM Disks

VM Disk Example



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

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

Disk Groups

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

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

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

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


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

Subdisks

Subdisk Example

VM Disk with One Subdisk

Subdisk




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

New VERITAS Product Distribution Method

The VERITAS products CD has an automated installation and licensing procedure that allows you to install packages using an Installation Menu rather than installing from the command line. VERITAS has also introduced a website from which you can quickly and easily obtain a license key for the products that you have purchased. The *VERITAS Storage Solutions 3.5 Product License and Installation Guide*, included as the liner notes with the VERITAS product CD, provides complete information. You should review the *VERITAS Storage Solutions 3.5 Product License and Installation Guide* before installing VxVM..

Verify Disk Space Requirements

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

Package	Contents	Size
VRTSvxvm	Driver and Utilities	59 MB
VRTSvlic	Licensing Utilities	2MB
VRTSvmman	Manual Pages	1 MB
VRTSvmdoc	Documentation	30 MB
VRTSob, VRTSobgui	VEA Service VEA Administrator	33 MB



Package	Contents	Size
VRTSvmp _{pro}	VxVM Provider for VEA	5MB
VRTSfs _{pro}	VxVM Provider for VEA	8.5MB

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

Package	/	/usr	/opt
VRTSvx _{vm}	27 MB	32 MB	0
VRTSv _{lic}	0	1MB	1MB
VRTSv _{mman}	0	0	1 MB
VRTSv _{mdoc}	0	0	30 MB
VRTSv _{mpro}	0	0	5 MB
VRTSf _{spro}	0	0	8.5 MB
VRTSob, VRTSob _{gui}			33 MB

Note VERITAS also offers VERITAS Mapping Services Application, VxMS. This is a set of shared libraries that offers developers of data storage management applications a common set of tools to provide data mapping into their storage management products. VxMS is intended for both VERITAS and third party product development. There are no end user features of VxMS, and it is not to be sold as a stand-alone product. Instead, it will be licensed within other software applications. Please contact your VERITAS Sales Representative for further information.

Set the PATH and MANPATH Variables

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

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

Examples:

If you are using a shell such as `sh` or `ksh`, use the following commands:

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

If you are using a shell such as `csh` or `tcsh`, use the following commands:

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

Verify Directories and Permissions

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

Product Licensing

The VERITAS Volume Manager is a licensed product. Before you install VxVM, obtain a license key from VERITAS. If you have a temporary license key, you must obtain a permanent license key when you purchase the product. A License Key Request Form (LKRF) is included in the product package. The LKRF has all the information required to establish a user account on the VERITAS vLicense™ website and generate your license key. If you do not receive a License Key Request Form, contact your sales representative, or send an email with your sales order number to license@veritas.com.



To obtain a product license, use the vLicense website (see below) or complete the License Key Request Form, including your system's host ID and model type. Sign and date the completed form and fax it to VERITAS. You will receive a license key by email within a few business days. Retain the License Key Request Form for your records.

Using the VERITAS vLicense Website to Obtain a License

You can access the vLicense website at:

<http://www.veritas.com/vlicense>

To obtain a license from the website, provide the following information shown on the License Key Request Form:

- ◆ Your customer number
- ◆ Your order number
- ◆ Your serial number

To determine the host ID of your system, enter:

```
# hostid
```

To determine the machine type, enter:

```
# uname -i
```

After receiving a license key, record the number somewhere other than on the system where it is installed so that you can access it if the system becomes inoperable or the product requires reinstallation.

During the installation procedure, you must enter a license key. If you are replacing a temporary license with a permanent license, follow the instructions that accompany the license key to replace the old license with the new one.

The VERITAS licensing commands are provided in the software package VRTSvlic. The VRTSvlic package must be installed for the licensing process to work.

If you have any questions concerning licensing, contact VERITAS at the license information numbers listed under “[Getting Help](#)” on page ix.

Add or Upgrade Licenses

VxVM 3.5 introduces a new licensing package, `VRTSvlic`, that is common to all new VERITAS products. The `VRTSvlic` package replaces the functionality of the `VRTSlic` package that was used in VxVM 3.2 and earlier releases. The following table lists equivalency between licensing commands in `VRTSlic` and `VRTSvlic`:

VRTSlic command	VRTSvlic command	Description
<code>vxlicense -c</code>	<code>vxlicinst</code>	Installs a license.
<code>vxlicense -p</code>	<code>vxlicrep</code>	Reports installed licenses.
<code>vxlicense -t</code>	<code>vxlictest</code>	Tests a license.

Any existing valid VxVM license keys are accepted for backward compatibility on upgrading to VxVM 3.5. The new licensing package silently converts the old key to a new value. However, you may need to obtain additional licenses for features that are new in VxVM 3.5. You should retain the old `VRTSlic` package on your system if it is required by any installed VERITAS products other than VxVM. If you choose to remove the `VRTSlic` package to free up disk space, this does not delete the existing license key files.

If you are upgrading to VxVM 3.5 and have not yet added the `VRTSvlic` package to your system, you can use the following commands to view the currently installed licenses:

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

If you need to add new licenses to support features in VxVM 3.5, you must add the new licensing package, `VRTSvlic`, as described in the sections on upgrading VxVM in the chapter “[Installing VERITAS Volume Manager](#)” on page 17.

After the `VRTSvlic` package has been installed, you can view the currently installed licenses (for both `VRTSlic` and `VRTSvlic`) with the following command:

```
# vxlicrep
```

To add new licenses to VxVM and to other VERITAS products that use the `VRTSvlic` package for licensing, enter the following command:

```
# vxlicinst
```

For any installed VERITAS products that use the `VRTSlic` package for licensing, use the following command to add new licenses:

```
# vxlicense -c
```

For more information about the new licensing commands in the `VRTSvlic` package, refer to the `vxlicinst(1)`, `vxlicrep(1)` and `vxlictest(1)` manual pages.



Take a Disk Inventory

Disk Types

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

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

Boot Disk

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

Verify Disk Contents

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

1. Make sure you are aware of the contents of each disk. Determine which disks can be encapsulated (data is preserved) or initialized (data is removed).
2. Do you want to place the system root disk under VxVM control? (This is necessary in the case of a one-disk system or when you wish to mirror the root disk.)

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

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

Disk ID	Encapsulate, Initialize, Exclude

5. Verify that 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. This process is known as *swap relocation* and, if necessary, happens automatically during root disk encapsulation.

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

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

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



- b. Use the `prtvtoc(1M)` command to record the layout of the partitions on the unencapsulated boot disk (`/dev/rdisk/c0t0d0s2` in this example):

```
# prtvtoc /dev/rdisk/c0t0d0s2
```

Record the output from this command for future reference.

- c. 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 `/devices` directory, aliases may not function correctly. This might happen if the pathname is incorrectly entered at the command line, or if the device's pathname changes at a later date.

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:

- If your system uses Sun's AP software, this release of VxVM requires an AP upgrade and a Solaris patch. See "[DMP and Alternate Pathing](#)" on page 14.
- If you are upgrading from VxVM 3.1 or an older version, DMP is automatically enabled in VxVM 3.5. You can use `vxinstall` or `vxdiskadm` to prevent DMP, if you choose.
- If you are upgrading from VxVM 3.1.1, and you had prevented or suppressed DMP, VxVM 3.5 retains your DMP setting. Your new installation will have DMP prevented or suppressed. Use `vxinstall` option #3 to enable DMP if you choose.

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 75 for details.

Solaris Patches for VxVM

VxVM requires Solaris patches to work properly. See “[Important Software Patch Information](#)” on page 18 for details about required patches.

VERITAS Enterprise Administrator (VEA)

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

- ◆ Solaris 2.6 or higher.

Note To install VEA on Solaris 2.6, 7, or 8, you must download and install the appropriate JAVA 1.3.1 patches for Solaris 2.6, 7, or 8 from <http://java.sun.com/j2se/1.3/install-solaris-patches.html>

- ◆ VERITAS Volume Manager Release 3.5. The `vxconfigd` and VEA service (`vxsvc`) must be running on the machine to be administered.
- ◆ The VEA server. See “[Setting Up and Starting VEA](#)” on page 46 on how to install the server.
- ◆ If a VERITAS Storage Administrator (VMSA) package is installed on the machine, refer to “[VMSA and VEA Co-Existence](#)” on page 46 before installing the new GUI.

A5x00 Devices

Whether or not you have A5x00 devices connected to your system, you must add the Solaris patches, shown in “[Solaris Patches for VxVM](#)” on page 18, to this release of VxVM. If you have A5x00 devices attached to your system, you need this patch to use them; if you don’t have A5x00 devices attached to your system, you still need the patches because they are used in the process of discovering whether such devices exist.



DMP and Alternate Pathing

VxVM 3.5 allows Dynamic Multipathing (DMP) to co-exist with Sun's Alternate Pathing (AP) software. If your system uses Sun's AP software, this release of VxVM requires AP version 2.3.1 and Sun patch 110722 (latest version). Upgrade AP to 2.3.1 as described in “[Installing or Upgrading VxVM with Alternate Pathing Enabled](#)” on page 20. Then install the patch before installing or upgrading to this release of VxVM. See “[Solaris Patches for VxVM](#)” on page 18 for patch information.

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

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

Claiming Devices in the SENA Category

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

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

Cluster Environment Requirements

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

1. Obtain a license for the *optional* VxVM cluster feature from your Customer Support channel.
2. Decide where to place the `rootdg` disk group for each node in the cluster. A system using VxVM has one or more disk groups, including the `rootdg`. The `rootdg` must exist, and it cannot be shared between systems. At least one disk must exist within the `rootdg` while VxVM is running.
3. Decide on the layout of shared disk groups. There may be one or more shared disk groups. Determine how many you wish to use.
4. If you plan to use Dirty Region Logging (DRL) with VxVM in a cluster, leave a small amount of space on the disk for these logs. The log size is proportional to the volume size and the number of nodes. Each log has one recovery map plus one active map per node.

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

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

Set up your SAN Configuration

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

- ◆ You must place at least one disk into the rootdg.
- ◆ If some disks in your SAN will not be under VxVM control, create a file called `/etc/vx/cntrls.exclude` or `/etc/vx/disks.exclude`. (See [“Running the vxinstall Program”](#) on page 24 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 VEA or the `vxdiskadm` program).





Installing VERITAS Volume Manager

2

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

- ◆ VxVM installation - see [“Installation and Upgrade Overview”](#) on page 17
- ◆ VxVM upgrade - see [“Upgrading to VxVM Release 3.5”](#) on page 54
- ◆ VxVM configuration and setup
- ◆ VEA setup
- ◆ Shared disk setup (for cluster environments only)
- ◆ Removing VxVM

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

Installation and Upgrade Overview

Note The VERITAS products CD has an automated installation and licensing procedure that allows you to install packages using an Installation Menu rather than installing from the command line. For more information, refer to [“Using the Installation Menu to Install”](#) on page 17, and the *VERITAS Storage Solutions 3.5 Product License and Installation Guide*, included as the liner notes with the VERITAS product CD.

1. For installations and upgrades, install the required patches. Refer to [“Important Software Patch Information”](#) on page 18.
2. For installations and upgrades, you may need to upgrade Sun’s Alternate Pathing (AP) driver to version 2.3.1. Refer to [“Installing or Upgrading VxVM with Alternate Pathing Enabled”](#) on page 20.



3. For new installations, load the software packages onto the system to be administered as described in the section “[Loading the Software from CD-ROM](#)” on page 21.
4. For upgrades, you need to determine if your root disk is encapsulated, as this determines which upgrade process you should use. See “[Upgrading to VxVM Release 3.5](#)” for more information.
5. For new installations, install VxVM as described in “[Running the vxinstall Program](#)” on page 24.
6. Set up and start VEA as described in “[Setting Up and Starting VEA](#)” on page 46.
7. If you are planning to use VxVM in a cluster environment, follow the instructions in “[Enabling Cluster Support in VxVM \(Optional\)](#)” on page 50.
8. After installation, add any disk arrays that are unsupported by VERITAS to the JBOD category as described in the section “[Hot-Relocation Guidelines](#)” on page 75.

Important Software Patch Information

Caution Please read this section completely *before you install* VERITAS Volume Manager (VxVM) 3.5. You must reboot your system after installing these patches.

Solaris Patches for VxVM

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

Check for Existing Patches

Some required system patches may already be present in your operating system. To display a list of patches that are already installed on your system, enter:

```
# showrev -p
```

For more information, see `showrev(1M)`.

Download Required Patches

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

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

Required Solaris Patches

For Solaris Version	Sun Patch ID Number	Notes
Solaris 9	None	
Solaris 8	<ol style="list-style-type: none"> 1. SUNWsan package (see your vendor to obtain) 2. 109529 (-06 or higher) 3. 111413 (-06 or higher) 	Regardless of whether you have an A5k drive connected, you must install these A5k patches for VxVM to function. You must install the SUNWsan package and patch 109529 before installing patch 111413. You can then install VxVM. This order of these steps is important.
	108827 (-19 or higher)	
	108528 (-14 or higher)	
	110722 (-01 or higher)	If you have Sun's Alternate Pathing (AP), you must install Solaris patch 110722-01 for VxVM 3.5 to co-exist with the AP driver version 2.3.1. However, you must also upgrade AP. Please see " DMP and Alternate Pathing " before installing VxVM.
Solaris 7	107473 (-07 or higher)	Regardless of whether you have an A5k drive connected, you must install an A5k patch for VxVM to function. You should install this patch before installing VxVM.
	110722 (-01 or higher)	If you have Sun's Alternate Pathing (AP), you must install Solaris patch 110722-01 for VxVM 3.5 to co-exist with the AP driver version 2.3.1. However, you must also upgrade AP. Please see " DMP and Alternate Pathing " before installing VxVM.
	106541 (-19 or higher)	If you are running Solaris 7, you must install Sun patch ID 106541.



Required Solaris Patches

For Solaris Version	Sun Patch ID Number	Notes
Solaris 2.6	110722 (-01 or higher)	If you have Sun's Alternate Pathing (AP), you must install Solaris patch 110722-01 for VxVM 3.5 to co-exist with the AP driver version 2.3.1. However, you must also upgrade AP. Please see “DMP and Alternate Pathing” before installing VxVM.
	105375 (-26 or higher)	Regardless of whether you have an A5k drive connected, it is recommended that you install A5k patch for VxVM to function. Install this patch before installing VxVM.
	105181 (-31 or higher)	

Note If 'patchadd -p' gives error as:

```
#patchadd -p | grep 111413
```

```
/usr/sbin/patchadd[193]: PatchArrElem: subscript out of range, then please install the patch 108987-08 (or later).
```

Installing or Upgrading VxVM with Alternate Pathing Enabled

Dynamic Multipathing (DMP) on VxVM 3.5 co-exists with Sun's Alternate Pathing (AP) driver 2.3.1. For VxVM 3.5 to function, the DMP driver (vxddmp) must always be present on the system. Upgrading to this version of VxVM automatically enables the DMP driver.

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

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

For VxVM to co-exist with AP successfully, you must upgrade AP before installing or upgrading to VxVM 3.5.

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

Loading the Software from CD-ROM

The Packages

The VERITAS packages are found on the CD-ROM in the `volume_manager/pkgs` directory:.

CD-ROM Contents

- ◆ `VRTSvxvm`—VERITAS Volume Manager and Volume Replicator software.
- ◆ `VRTSvlic`—VERITAS Licensing Utilities.
- ◆ `VRTSvmdoc`—online copies of VERITAS Volume Manager guides.
- ◆ `VRTSvrdoc`—online copies of VERITAS Volume Replicator guides.
- ◆ `VRTSvmman`—VxVM manual pages.
- ◆ `VRTSob`—VERITAS Enterprise Administrator Service
- ◆ `VRTSobgui`—VERITAS Enterprise Administrator
- ◆ `VRTSvmpro`—VERITAS Virtual Disk Management Services Provider (required if you install `VRTSob` and `VRTSobgui`)
- ◆ `VRTSfspro`—VERITAS File System Provider (required even if you are not installing the VERITAS File System software)

The following related packages are also in the `volume_manager/pkgs` directory, but are installed by other VERITAS products:

- ◆ `VRTSvrw`—VERITAS Volume Replicator Web Console Package.
- ◆ `VRTSweb`—VERITAS Web GUI Common Web Server Package

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

Mounting the CD-ROM with Solaris Volume Management

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

1. Log in as superuser.
2. Insert the CD-ROM into a CD-ROM drive connected to your system.



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

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

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvlic VRTSvxvm VRTSvmdoc VRTSvmman
# pkgadd -a ../scripts/VRTSobadmin -d . VRTSob VRTSobgui
# pkgadd -d . VRTSfspro VRTSvmpro
```

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

Removing Previous Licensing Packages

If you add the `VRTSvlic` package and see the following message, you must first remove the existing `VRTSvlic` package:

```
Current administration requires that a unique instance of the
<VRTSvlic> package be created. However, the maximum number of
instances of the package which may be supported at one time on the
same system has already been met.
```

1. To de-install the old package, enter the following command:

```
# pkgrm VRTSvlic
```

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

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

3. Use `pkgadd` to install the latest `VRTSvlic` package.

Note `VRTSvxvm` does not lose its package dependency when the latest `VRTSvlic` is installed later in this process.

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

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

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

Mounting the CD-ROM from the Command Line

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

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

Follow these steps:

1. Log in as superuser.
2. Insert the CD-ROM into a CD-ROM drive connected to your system.
3. To mount the CD-ROM manually, use the following command:

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

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

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

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvlic VRTSvxvm VRTSvmdoc VRTSvmman
# pkgadd -a ../scripts/VRTSobadmin -d . VRTSob VRTSobgui
# pkgadd -d . VRTSfspro VRTSvmpro
```

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

If you added the `VRTSvlic` and see a message that says no changes were made to the system, follow the instructions in [“Removing Previous Licensing Packages”](#) on page 22.

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

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

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

Installing VxVM Using the Installation Menu

You can install the VxVM Solaris products using an installation menu provided on CD1:



1. Put the CD in the CD-ROM drive.
2. Enter `/cdrom/storage_solutions_solaris_3.5cd1/installer` and press Return.
Depending on whether the `VRTSlic` application is installed, one of the following is displayed:
 - A message asking if you want to install `VRTSlic`
If it is not installed, follow the instructions for installation.
 - The Product Status page, which shows:
 - Products available for installation
 - Whether a version of any of the products is currently installed
 - Whether any of the installed products are licensed
 - Options for operations you can initiate
3. At the prompt on the Product Status page, enter 1 to add a license key, and then press Return to begin.
4. At the next prompt, enter the license key and press Return.
5. At the Product Status page, enter 2 for the Installation Menu, and then press Return.
The Installation Menu is displayed.
6. At the Installation Menu, enter the number of the product you want to install, and then press Return.
The product installation begins automatically. When the operation finishes, the Installation Menu is displayed again.

Running the vxinstall Program

For new installations, use the `vxinstall` command. This installation has two options, Quick and Custom. See these sections for more information.

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

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

Required Reboots

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



Rootdg

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

Placing Disks in another Disk Group

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

Run the vxinstall Program

Initialize VxVM using the following procedure:

1. Log in as superuser.
2. Identify controllers, disks and enclosures to exclude from being configured as VxVM devices by the `vxinstall` and `vxdiskadm` utilities. If there are no devices to exclude, continue to [step 4](#).

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

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

3. To exclude devices from VxVM control, create or edit the `/etc/vx/cntrlr.exclude`, `/etc/vx/disks.exclude` and `/etc/vx/enclr.exclude` files. You can:
 - exclude one or more disks from VxVM control.
 - exclude all disks on certain controllers from VxVM control.
 - exclude all disks in specific enclosures from VxVM control.

Note The `vxinstall` and `vxdiskadm` utilities use the exclude files to avoid configuring the specified disks as VxVM devices. These files do not exclude disks from use by any other VxVM commands.

The following section shows examples of each exclude file.

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



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

```
c0t1d0
```

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

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

```
c0  
c1
```

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

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

```
sena0  
emc1
```

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

4. Start the `vxinstall` program. Type:

```
# vxinstall
```

The `vxinstall` program first runs the `vxlicinst` command to initialize the VxVM license key file. The `vxlicinst` command displays licensing information and then prompts you for a key. You must obtain and enter a valid key in order to proceed with the initialization.

Note The presence of certain hardware arrays (for example, A5x00) automatically generates a key. The `vxinstall` program does not prompt for another key.

5. The `vxinstall` program then examines all controllers attached to the system and lists them. You are asked if you want to use enclosure-based naming. If you do not want to use enclosure-based names, enter `n`, and proceed to [step 9](#) on page 28.

Generating list of attached enclosures....

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

```
<enclosurename>_<diskno>
```

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

Some examples would be:

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

For non-public loop disks (Ex. FC disks directly connected to the host or through hubs), you could choose to use the c#t#d#s# naming format.

Do you want to use enclosure based names for all disks ? [y,n,q,?] (default: n) **n**

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

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

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

```
Enclosures:      enc01      sena0      sena1      sena3
                  sena4      sena5
Others:          others0
```

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

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

```
Rename an enclosure
Menu: VolumeManager/Install/Rename enclosure
```

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

Enter the name of an enclosure you wish to rename.

Enter the new name for enclosure test1: enc0

The enclosure name test1 has been changed to enc0

Hit RETURN to continue.



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

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

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

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

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

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

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

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

10. Press Return to continue.

The `vxinstall` program displays a menu with the following options:

```
1 Quick Installation
2 Custom Installation
3 Prevent multipathing/Suppress devices from VxVM's view
? Display help about menu
?? Display help about menuing system
q Exit from menus
```

Select an operation to perform:

The options are defined as follows:

- 1 (Quick Installation) uses default options for installation.

- 2 (Custom Installation) allows you to control whether and how each disk will be brought under VxVM control.
- 3 (Prevent multipathing/Suppress devices from VxVM's view) excludes devices from VxVM or prevents them from being multipathed by DMP.
- ? displays a help file describing the current operation or menu choices.
- ?? displays general information about using the vxinstall program.
- q exits from the current operation, or from the vxinstall program.

11. Select an option.

- If you choose 1, see “[Quick Installation](#)” on page 29.
- If you choose 2, see “[Custom Installation](#)” on page 35.
- If you choose 3, see “[Prevent Multipathing/Suppress Devices from VxVM's View](#)” on page 40.

Quick Installation

The Quick Installation option is the easiest to use. For each controller, Quick Installation gives you the option of either initializing or encapsulating *all* disks. Encapsulation preserves existing data; initialization removes the data.

Initialization and Encapsulation Notes

- ◆ Quick Installation allows you to initialize disks in a disk array, even if valid partitions are found on those disks. However, avoid selecting the initialize option for any disks containing valid data.
- ◆ Quick Installation places every disk attached to your system under VxVM control, unless you have chosen to exclude those disks. Chose disks for exclusion using the `disks.exclude`, `cntrls.exclude` or `enclr.exclude` file or by using the interface for suppressing devices as described earlier. If you have any disks on your system that you do not want placed under VxVM control, either exclude them using the Prevent multipathing/Suppress devices from VxVM's view in vxinstall option or Custom Installation option (which allows you to omit certain disks from installation).
- ◆ To encapsulate some disks in a given disk array and initialize others, use Custom Installation instead. (See “[Custom Installation](#)” on page 35.)
- ◆ With encapsulation, the `/etc/vfstab` file is updated to ensure that file systems previously mounted on disk partitions are mounted on volumes instead.



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

Run Quick Installation

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

To use the Quick Installation option, follow these steps:

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

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

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

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

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

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

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

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

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

```
Disk array serial number : 60020f20000005bd0000
```

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


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

```
enc0_0 enc0_1 enc0_2 enc0_3
```

Hit RETURN to continue.

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

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

Caution Initializing a disk destroys all data on that disk.

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

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

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

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

Disks not having valid partitions will be initialized.

Hit RETURN to continue.

Press Return to continue.

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

Are you sure (destroys data on these disks)

[y,n,q,?] (default: n)

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

Volume Manager will now initialize all the disks on this
disk array including disks having valid partitions.

Hit RETURN to continue.



Press Return to continue.

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

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

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

The enc0_0 disk will be given disk name disk02

The enc0_0 disk appears to be empty. Adding as a new disk.

The enc0_1 disk will be given disk name disk05

The enc0_1 disk appears to be empty. Adding as a new disk.

The enc0_2 disk will be given disk name disk06

The enc0_2 disk appears to be empty. Adding as a new disk.

The enc0_3 disk will be given disk name disk07

The enc0_3 disk appears to be empty. Adding as a new disk.

The clt0d0 disk will be given disk name disk01

The clt1d0 disk will be given disk name disk02

The clt2d0 disk will be given disk name disk03

The clt3d0 disk will be given disk name disk04

Hit RETURN to continue.

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

The clt0d0 disk has been configured for encapsulation.

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

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

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

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

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



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

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

The following is a summary of your choices.

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

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

8. This is your last chance to review and alter your choices for any of the disks to be placed under VxVM control.
 - If you enter **y**, the `vxinstall` program encapsulates all disks listed with Encapsulate and initializes (destroying any existing data) all disks listed with New Disk.
 - If you enter **n**, the `vxinstall` program prompts you for the name of a disk to be removed from the list and excluded from VxVM control:

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

9. Enter the name of the disk to be removed from the list and press Return.
10. Repeat these steps for each disk to be removed and then press Return again. The `vxinstall` program displays an updated summary
11. the disks chosen for removal. If you are satisfied with the summary, enter **y**.
12. 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 45.

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.

Custom Installation

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

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

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

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

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

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

2. Indicate whether to encapsulate your boot disk. (In this example, your boot disk is named `c0t3d0`.)

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

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

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

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

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

```
Disk array serial number : 60020f20000005bd0000
```



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

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

enc0_0 enc0_1 enc0_2 enc0_3

Hit RETURN to continue.

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

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

Installation options for enclosure enc0

Menu: VolumeManager/Install/Custom/enc0

- 1 Install all disks as pre-existing disks. (encapsulate)
- 2 Install all disks as new disks.(discards data on disks!)
- 3 Install one disk at a time.
- 4 Leave these disks alone.
- ? Display help about menu
- ?? Display help about the menuing system
- q Exit from menus

Select an operation to perform:

Encapsulating All Disks

1. To encapsulate all disks in the disk array, select 1.

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

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

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

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

```
The clt0d0 disk will be given disk name disk01
The clt0d0 disk has been configured for encapsulation.
The clt1d0 disk will be given disk name disk02
The clt1d0 disk has been configured for encapsulation.
Hit RETURN to continue.
```

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

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

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

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

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

Initializing All Disks

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

Caution Initializing a disk destroys all data on that disk.

All disks in the disk array are initialized. This destroys all data and partitions on each disk and makes the disk available as free space for allocating new volumes, or mirrors of existing volumes. If you are running the `vxinstall` program on new disks that have never been placed under VxVM control before, choose this option.

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

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

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

```
The clt0d0 disk will be given disk name disk01
The clt1d0 disk will be given disk name disk02
Hit RETURN to continue.
```



- If you enter **n**, 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 in the current disk array have been named, press Return to move on to the next controller.

Installing Individual Disks

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

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

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

```
Installation options for disk enc0
Menu: VolumeManager/Install/Custom/enc/enc0_0

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

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

Select an operation to perform:

These options are similar to those in the Custom Installation main menu. Once you have determined how to handle this disk, see [step 4](#) on page 36.

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

Leaving All Disks Unaltered

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

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

2. Repeat [step 4](#) on page 36 for each of the remaining disk arrays on your system.

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

The following is a summary of your choices.

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

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

3. This is your last chance to review and alter your choices for any of the disks to be placed under VxVM control.
 - If you enter **y**, the `vxinstall` program proceeds to encapsulate all disks listed with Encapsulate and initialize (destroying any existing data) all disks listed with New Disk.
 - If you enter **n**, the `vxinstall` program prompts you for the name of a disk to be removed from the list and excluded from VxVM control, as follows:


```
Enter disk to be removed from your choices. Hit return when
done.
[<name>,q,?]
```
4. Enter the name of the disk to be removed from the list and press Return. Repeat for each disk to be removed and then press Return once more. The `vxinstall` program displays an updated summary without the disks chosen for removal. If you are satisfied with the summary, enter **y**.



5. Depending on your input during the `vxinstall` procedure, you may be requested to shut down and reboot your machine. If this happens, shut down and reboot your machine immediately or at the earliest convenient opportunity. For more information, see [“Rebooting After vxinstall”](#) on page 45.

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

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

Installation is complete.

Prevent Multipathing/Suppress Devices from VxVM's View

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

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

The following message displays:

Volume Manager Installation

Menu: VolumeManager/Install/Exclude Devices

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

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

2. Enter **y**.

3. The following message displays. Select an operation.

Volume Manager Device Operations

Menu: VolumeManager/Install/Exclude Devices

1 Suppress all paths through a controller from VxVM's view

2 Suppress a path from VxVM's view

3 Suppress disks from VxVM's view by specifying a VID:PID combination

4 Suppress all but one path to a disk

5 Prevent multipathing of all disks on a controller by VxVM

6 Prevent multipathing of a disk by VxVM

7 Prevent multipathing of disks by specifying a VID:PID combination

8 List currently suppressed/non-multipathed devices

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

Select an operation to perform:

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

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

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

This operation can be reversed using the `vxdiskadm` command.

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

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

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

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

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

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

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

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

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

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

The path specified here will be excluded from VxVM control.

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

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



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

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

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

VID:PID where VID stands for Vendor ID
 PID stands for Product ID
(The command 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/vxdmpinq.

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

Exclude all but one paths to a disk

Menu: VolumeManager/Install/ExcludeDevices/PATHGROUP-VXVM

Use this operation to exclude all but one paths to a disk. In case of disks which are not multipathed by vxddmp, 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 vxddiskadm 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 DMP

Menu: VolumeManager/Install/ExcludeDevices/CTLR-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 vxddiskadm command.

You can specify a controller name at the prompt. A controller name is of the form c#, example c3, cl1 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:[<ctrlr-name>,all,list,list-exclude,q,?]

The controller entered here is excluded from DMP control.

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

Exclude paths from DMP

Menu: VolumeManager/Install/ExcludeDevices/PATH-DMP

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

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

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

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

Both VID and PID can have an optional '*' (asterisk) following them.

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

Some examples of VID:PID specification are:

all- Exclude all disks

aaa:123- Exclude all disks having VID 'aaa' and PID '123'

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

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

aaa:* - Exclude all disks having VID 'aaa' and any PID

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

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

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

Rebooting After `vxinstall`

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

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

Note The setup you choose can require several reboots.

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

```
The system now must be shut down and rebooted in order to continue
the reconfiguration.
Shutdown and reboot now [y,n,q,?] (default: n)
```

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

- Enter **y** to begin an immediate shutdown.
- Enter **n** to cause the `vxinstall` program to exit without starting a shutdown. If you select this option, reboot as soon as possible.

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



Setting Up and Starting VEA

To use VEA, the following packages need to be installed and running on the server:

- ◆ VERITAS Enterprise Administrator Service (VRTSob)
- ◆ VERITAS Volume Manager Service Provider (VRTSvmpro)
- ◆ VERITAS File System Service Provider (VRTSfspro)

Note All these packages are dependent on one another, and all must be present on the server for VEA to run. Furthermore, VRTSob must already be installed before you can install VRTSvmpro. Once installed, you cannot remove VRTSob before removing (deinstalling) VRTSvmpro.

In addition, one of the following packages needs to be installed and running on the client:

- ◆ VERITAS Enterprise Administrator (VRTSobgui)
This is the client package for Solaris.
- ◆ VERITAS Enterprise Administrator for Windows® (win32/VRTSobgui.msi)
This is the client package for Windows.

Note To install VEA on Solaris 2.6, 7, or 8, you must install the appropriate java 1.3.1 patches for Solaris 2.6, 7, or 8.

This section describes how to install and set up VEA. The VEA server must be installed and run on a Solaris machine running VxVM 3.5. The VEA client can be installed and run on a Solaris or Windows machine that supports the Java 1.1 Runtime Environment.

Note This release of VEA requires VxVM Release 3.5. Before you install and run VEA, you must install VxVM and run the `vxinstall` program to initialize VxVM.

The following are minimum system recommendations for the VEA client:

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

VMsa and VEA Co-Existence

If you do not plan to use VMsa to administer other (pre-VxVM 3.5) machines, then you should uninstall this package before installing VEA.

Note You can later do a client-only install, if you want to run the VMSA client-only on your machine.

If you do not remove VMSA then, during boot, you will get the following warning:

```
VERITAS VM Storage Administrator Server terminated.
Stopping VERITAS VM Storage Administrator Server
### Terminated
```

This message informs you that the `vmsa_server` is unable to start. You can safely ignore this warning message.

You can avoid this message by removing the boot-time start up of the `vmsa_server`, using the command:

```
# /opt/VRTSvmsa/bin/autostart on
```

Alternatively, you can remove the files `/etc/init.d/vmsa-server` and `/etc/rc2.d/S96vmsa-server`.

VEA Installation Overview

Set up and start VEA as follows:

1. Install the VEA server and client on the UNIX (Solaris) machine to be administered. See [“Installing the VEA Server and/or Client \(Solaris\)”](#) on page 47.
2. If you plan to run the VEA client on a machine other than the machine to be administered, install the VEA client on the machine where the client will run. See [“Installing the VEA Server and/or Client \(Solaris\)”](#) on page 47 or [“Installing the VEA Client \(Windows\)”](#) on page 48.
3. Start the VEA server on the machine to be administered. See [“Starting VEA”](#) on page 49.
4. Start the VEA client. See [“Starting the VEA Client”](#) on page 49.

Installing the VEA Server and/or Client (Solaris)

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

1. Log in as superuser.
2. Install the VEA package, using `pkgadd`, as described in [“Loading the Software from CD-ROM”](#) on page 21.



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

Note VEA 3.5 is not compatible with earlier versions of VMSA, the previous VERITAS Volume Manager GUI. You cannot run VMSA with VxVM version 3.5.

The VEA package for UNIX contains the following directories:

Directory	Contents
/opt/VRTS/jre/	Java Runtime Environment (JRE)
/opt/VRTS/bin/	vxsvc—VEA server startup and stop script vea—VEA startup script
/opt/VRTS/man/	VEA manual pages

Installing the VEA Client (Windows)

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

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

Note If you plan to install the VEA client on Windows NT 4.0, Windows Installer must be upgraded to version 2.0. For more information about upgrading Windows Installer, visit <http://www.microsoft.com>. If you are using Windows NT 4.0, it is also recommended that you use Windows NT 4.0 Service Pack 6.

Install the VEA client on a Windows machine as follows:

1. Log in as administrator.
2. Insert the CD-ROM containing VEA.
3. Using Windows Explorer or a DOS Command window, go to the **win32** directory and execute the `vrtsobgui.msi` program with Windows installer.

4. Follow the instructions presented by the Windows installer 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 XP, ME, Windows 98 or Windows 95, restart the computer.

Starting VEA

After installing the VEA packages, the VEA service needs to be stopped and restarted.

To stop the VEA backend service, enter:

```
# /opt/VRTSob/bin/vxsvc -k
```

You can also stop the VEA backend service manually by killing this process.

Note The VEA backend service is automatically started on a reboot.

To start the VEA backend service, enter:

```
# /opt/VRTSob/bin/vxsvc
```

To check the state of the VEA backend service, enter:

```
# /opt/VRTSob/bin/vxsvc -m
```

Starting the VEA Client

Only users with appropriate privileges can run VEA. VEA can administer the local machine or a remote machine. However, VxVM and the VEA server must be installed on the machine to be administered. The VxVM `vxconfigd` daemon and the VEA server must be running on the machine to be administered.

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

UNIX (Solaris)

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

```
# /opt/VRTSob/bin/vea &
```

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

```
# /opt/VRTSob/bin/vea -host remote_machine_name
```



Note You might want to set your DISPLAY variable when administering a remote machine.

Windows

To administer a *remote* UNIX machine, select **Start > Programs > VERITAS Enterprise Administrator > VERITAS Enterprise Administrator**.

Enabling Cluster Support in VxVM (Optional)

Note This section assumes that you are using SunCluster as the cluster monitor on your system.

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

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

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

1. Obtain a license for the VxVM cluster feature.
2. Install the software packages onto each system (node) to be included in the cluster as described in “[Loading the Software from CD-ROM](#)” on page 21.
3. Initialize VxVM using the `vxinstall` procedure as described in “[Running the vxinstall Program](#)” on page 24.
4. Set up and start VEA as described in “[Setting Up and Starting VEA](#)” on page 46.
5. Configure shared disks as described in “[Configuring Shared Disks](#)” on page 51.

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 53).

Configuring New Disks

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

1. Start the cluster on at least one node.
2. On one node, run the `vxdiskadm` program and choose option 1 to initialize new disks. When asked to add these disks to a disk group, choose `none` to leave the disks for future use.
3. On other nodes in the cluster, run `vxctl enable` to see the newly initialized disks.
4. From the master node, create disk groups on the shared disks. To determine if a node is a master or slave, run `vxctl -c mode`.

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

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

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

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



Verifying Existing Shared Disks

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

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

```
# vxdg list
```

This displays the existing shared disk groups.

Converting Existing VxVM Disk Groups to Shared Disk Groups

If you are upgrading from VxVM 2.x to VxVM 3.5 and you want to convert existing disk groups to shared disk groups, configure the shared disks as follows:

1. Start the cluster on at least one node.

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

2. Configure the disk groups using the following procedure.

To list all disk groups, use the following command:

```
# vxdg list
```

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

```
# vxdg deport disk-group-name
```

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

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

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

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

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

```
# vxdg list
```

The disk groups are now ready to be shared.

3. If the cluster is only running with one node, bring up the other cluster nodes. Enter the `vxdg list` command on each node to display the shared disk groups. This command displays the same list of shared disk groups displayed earlier.

Disk Controller Reservation

As part of its quorum control, the Sun Cluster cluster manager reserves the shared disk controllers when only one node is active. This prevents “rogue” hosts from accessing the shared disks. When this happens, the `vxdisk list` command used on a node that has left the cluster may show all disks on such a controller as having an `error` status. The more detailed options of the `vxdisk` command show the flag `unavailable`. When a new node joins the cluster, the Sun Cluster software releases the controllers. VxVM attempts to access these disks, and if that is successful, the disks return to an `online` status. (See the Sun Cluster documentation for further details.) If one system boots while the other system has the disks reserved, the disks can be invisible to the booting system, and the `vxdisk` command may not display any of the shared disks. When the system joins the cluster, the shared disks become visible.



Upgrading to VxVM Release 3.5

If you are running a release earlier than 3.5, you can upgrade your system and retain your existing VxVM configuration. After upgrading, you can resume using VxVM as before (without running the `vxinstall` program again). Unless your license has expired, no new license is required for upgrading. However, you must install the new licensing package, which uses the existing license information.

Before You Upgrade

Before you upgrade, do the following:

- 1. Log in as superuser.
- 2. Install required patches.
- 3. Determine if the root disk is encapsulated by running the following command:

```
# mount | grep "/" on"
```

If the output from this command includes a pathname containing `vx` and `rootvol`, as in `/dev/vx/dsk/rootvol`, then your root disk is encapsulated. See [“Upgrading VxVM on an Encapsulated Root Disk”](#) on page 58.

Caution Use only the appropriate upgrade process to upgrade your system. Doing a `pkgrm` and `pkgadd` of VxVM does not support the process of upgrading from a previous version of VxVM to VxVM 3.5. You can only use `pkgrm` to remove VxVM packages when you are using the `upgrade_start/upgrade_finish` procedures.

Caution Do not run `vxinstall` during or after the upgrade.

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

To upgrade	Where	See
VxVM only	On an encapsulated root disk	“Upgrading VxVM on an Encapsulated Root Disk” on page 58.
VxVM only	On an unencapsulated root disk	“Upgrading VxVM Only” on page 66.
VxVM and Solaris	On an encapsulated or unencapsulated root disk	“Upgrading VxVM and Solaris” on page 63.

To upgrade	Where	See
Solaris only	On an encapsulated or unencapsulated root disk	“Upgrading Solaris Only” on page 71.

Note Importing a pre-3.5 VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.5 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.



VxVM and Solaris Release Support Matrix

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

	Solaris 2.5.1	Solaris 2.6	Solaris 7	Solaris 8	Solaris 9
VxVM 3.5		Supported	Supported	Supported	Supported
VxVM 3.2		Supported	Supported	Supported	Only VxVM 3.2S9 supports Solaris 9
VxVM 3.1.1		Supported	Supported	Supported	
VxVM 3.1		Supported	Supported	Supported	
VxVM 3.0.4	Supported	Supported	Supported	Supported	
VxVM 3.0.3	Supported	Supported	Supported	Supported	
VxVM 3.0.2	Supported	Supported	Supported		
VxVM 3.0.1	Supported	Supported	Supported		
VxVM 3.0	Supported	Supported			



VxVM and Solaris Upgrade Paths

Possible Upgrade Paths for VxVM and/or Solaris Upgrades

The following table shows possible VxVM/Solaris upgrade paths.

Upgrade Paths

You can upgrade VxVM and Solaris versions:	To VxVM and Solaris versions:	Notes (see also Note, below)
VxVM 2.0 to 3.0 & Solaris 2.3 to 7	VxVM 3.5 & Solaris 2.6, 7, 8, 9	See “Upgrading VxVM and Solaris” on page 63
VxVM 3.1 or 3.1.1 & Solaris 2.6, 7, 8	VxVM 3.5 & Solaris 7, 8, 9	
VxVM 3.2 & Solaris 2.6, 7, 8	VxVM 3.5 & Solaris 7, 8, 9	
You can upgrade VxVM on the following Solaris versions:	To VxVM and Solaris versions:	
VxVM 2.4 to 3.0 & Solaris 2.6	VxVM 3.5 & Solaris 2.6	For encapsulated root disks, see “Upgrading VxVM on an Encapsulated Root Disk” on page 58. For unencapsulated root disks, see “Upgrading VxVM Only” on page 66
VxVM 3.0.x & Solaris 2.6, 7, 8	VxVM 3.5 & Solaris 2.6, 7, 8, 9	
VxVM 3.1 & Solaris 2.6, 7, 8	VxVM 3.5 & Solaris 2.6, 7, 8, 9	
VxVM 3.1.1 & Solaris 2.6, 7, 8	VxVM 3.5 & Solaris 2.6, 7, 8, 9	
VxVM 3.2 & Solaris 2.6, 7, 8	VxVM 3.5 & Solaris 2.6, 7, 8, 9	
You can upgrade Solaris only:	To Solaris versions:	
VxVM 3.5 & Solaris 2.6	VxVM 3.5 & Solaris 7, 8, 9	“Upgrading Solaris Only” on page 71
VxVM 3.5 & Solaris 7	VxVM 3.5 & Solaris 8, 9	
VxVM 3.5 & Solaris 8	VxVM 3.5 & Solaris 9	“Upgrading Solaris Only” on page 71

Note The `upgrade_start` function will currently fail if it detects enclosure-based rootdisk names. You should run `vxdiskadm, option 20` (Change the disk naming scheme), to change enclosure-based rootdisk names to regular names before using `upgrade_start`.



Upgrading VxVM on an Encapsulated Root Disk

This section explains how to upgrade VxVM if the root disk is encapsulated. Refer to the “[VxVM and Solaris Release Support Matrix](#)” to determine if upgrading the operating system is required.

The process for upgrading VxVM on an encapsulated root disk has changed since VxVM 3.2. You now have two choices of tools to use for upgrading. You can use either:

- ◆ The `pkgadd` command on top of your existing software.
- ◆ The `upgrade_start/upgrade_finish` method to install VxVM 3.5.

Both methods have advantages and disadvantages, as shown in the following:

	Using <code>pkgadd</code>	Using <code>upgrade_start/upgrade_finish</code>
Advantages	Usually requires only 1 reboot	<ol style="list-style-type: none"> 1. There will be only a single VRTSvxvm package displayed by <code>pkginfo</code>. 2. VxVM configuration data is backed up. 3. The root disk is unencapsulated. Therefore, there is no risk of not being able to boot up the system. 4. VRTSlic can be deinstalled, if there are no other packages that require it.
Disadvantages	<ol style="list-style-type: none"> 1. In the highly unlikely event that a system failure occurs during the <code>pkgadd</code> process, it is possible that the encapsulated rootdisk may become unbootable. 2. The new package is displayed by <code>pkginfo</code> as VRTSvxvm.2. <p>Note You can avoid having multiple instances by specifying the <code>-a admin</code> option with <code>instance=overwrite</code>. See <code>pkgadd(1M)</code> for more details.</p> <ol style="list-style-type: none"> 3. You may not be able to remove VRTSlic if another installed VRTSvxvm package is still using the license. 	This process usually requires three reboots.

Caution If your root disk is not encapsulated, refer to “[Upgrading VxVM Only](#)” on page 66.

These procedures follow.

Upgrade with `pkgadd` command

1. If you have not already obtained and installed a VxVM 3.5 license key, do so now. See “[Product Licensing](#)” on page 7 for details.

2. Bring the system down to single-user mode using the following command:

```
# init s
```

3. Stop the `vxconfigd` daemon:

```
# vxdctl stop
```

4. Stop the `vxiod` daemon:

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

5. (Optional) Remove the VMSA software with the following command:

```
# pkgrm VRTSvmsa
```

Note If you still have clients running previous versions of VxVM, you can skip this step. change Step 7 to:

6. Add the VxVM 3.5 licensing package with the following:

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvlic
```

Note VxVM 3.5 requires the 3.5 licensing package. You do not need to remove the previous versions of VRTSlic.

7. Add the VxVM 3.5 package with the following:

```
# pkgadd -d . VRTSvxvm
```

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

8. Using the `pkgadd` command, install the additional packages:

```
# pkgadd -d . VRTSvmdoc VRTSvmman
# pkgadd -a ../scripts/VRTSobadmin -d . VRTSob VRTSobgui
```



```
# pkgadd -d . VRTSfspro VRTSvmpo
```

Note It is recommended that you install both VRTSfspro and VRTSvmpo if you intend to use the VEA GUI.

9. Perform a reconfiguration reboot:

```
# cd /; shutdown -g0 -y -i6
```

or

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

10. Importing a pre-3.5 VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.5 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.

11. See “[Setting Up and Starting VEA](#)” on page 46 to set up and restart VEA.

Upgrade with upgrade_start/upgrade_finish commands

Follow these steps:

1. If you have not already obtained and installed a VxVM 3.5 license key, do so now. See “[Product Licensing](#)” on page 7 for details.
2. If any of the file systems `/`, `/usr`, `/var`, `/opt`, `/export/home`, `/var/adm` or `/usr/kvm` are defined on volumes, make sure that at least one plex for each of those volumes is created from a single subdisk that begins on a cylinder boundary.

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

Note If any of the documentation, developer kit, or manual page packages are to be installed, the directory `/opt` must exist, be writable, and must not be a symbolic link. This is because the volumes not temporarily converted by the

`upgrade_start` will be unavailable during the upgrade process. If you have a symbolic link from `/opt` to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in `/opt` will not be installed.

3. Bring the system down to single-user mode using the following command.

```
# init s
```

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

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

5. Run the `upgrade_start -check` script to find any problems that exist which could prevent a successful upgrade. This script reports errors if found. Otherwise, it reports success and you can proceed with running the `upgrade_start` script.

```
# ./upgrade_start -check
```

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

```
# /cdrom/volume_manager/scripts/upgrade_start
```

The `upgrade_start` script searches for volumes containing file systems.

7. Reboot to single-user mode.

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

Caution The root disk partitions should not be modified: this will cause `upgrade_start/upgrade_finish` to fail. If you change the root disk partitions, you cannot upgrade VxVM.

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

8. Remove the old VxVM packages by entering:

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

Note If you still have clients running old versions of VxVM, you can skip removing VMSA.

Note VxVM 3.5 requires the 3.5 licensing package. You do not need to remove the previous versions of `VRTSlic`.



9. Reboot to single-user mode.

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

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

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

11. Add the VxVM 3.5 licensing package with the following:

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvlic
```

12. Add the VxVM 3.5 package with the following:

```
# pkgadd -d . VRTSvxvm
```

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

13. Complete the upgrade using the following command:

```
# ../scripts/upgrade_finish
```

14. Perform reconfiguration reboot:

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

15. Using the `pkgadd` command, install the additional packages:

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvmdoc VRTSvmman
# pkgadd -a ../scripts/VRTSobadm -d . VRTSob VRTSobgui
# pkgadd -d . VRTSfspro VRTSvmpro
```

16. Importing a pre-3.5 VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.5 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.

17. See “[Setting Up and Starting VEA](#)” on page 46 to set up and restart VEA.

Upgrading VxVM and Solaris

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

Follow these steps:

1. If you have not already obtained and installed a VxVM 3.5 license key, do so now. See “[Product Licensing](#)” on page 7 for details.
2. If any of the file systems `/`, `/usr`, `/var`, `/opt`, `/export/home`, `/var/adm` or `/usr/kvm` are defined on volumes, make sure that at least one plex for each of those volumes is created from a single subdisk that begins on a cylinder boundary.

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

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

3. Bring the system down to single-user mode using the following command:

```
# init s
```

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

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

5. Run the `upgrade_start -check` script to find any problems that exist which could prevent a successful upgrade. This script reports errors if found. Otherwise, it reports success and you can proceed with running the `upgrade_start` script.

```
# ./upgrade_start -check
```

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



```
# /cdrom/volume_manager/scripts/upgrade_start
```

The `upgrade_start` script searches for volumes containing file systems.

7. Reboot to single-user mode.

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

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

8. Remove the old VxVM packages by entering:

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

Note If you still have clients running old versions of VxVM, you can skip removing VMSA.

Note Licensing in VxVM 3.5 requires the new `VRTSvlic` package. You do not need to remove the existing `VRTSlic` package.

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

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

11. If you are upgrading Solaris, upgrade the operating system to Solaris 2.6, 7, 8, or 9. Refer to the Solaris installation documentation for instructions on how to upgrade Solaris.

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

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

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

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

- 14.** Add the VxVM 3.5 licensing package using the following commands:

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvlic
```

- 15.** Add the VxVM 3.5 package using the following command:

```
# pkgadd -d . VRTSvxvm
```

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

- 16.** Complete the upgrade using the following command:

```
# ../scripts/upgrade_finish
```

- 17.** Perform reconfiguration reboot:

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

- 18.** Using the `pkgadd` command, install additional packages:

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvmdoc VRTSvmman
# pkgadd -a ../scripts/VRTSobadmin -d . VRTSob VRTSobgui
# pkgadd -d . VRTSfspro VRTSvmpro
```

- 19.** Importing a pre-3.5 VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.5 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 VEA](#)” on page 46 to set up and restart VEA.



Upgrading VxVM Only

This section explains how to upgrade VxVM if the root disk is NOT encapsulated. Refer to the “[VxVM and Solaris Release Support Matrix](#)” to determine if upgrading the operating system is required.

The process for upgrading VxVM on a non-encapsulated root disk has changed since VxVM 3.2. You now have two choices of tools to use for upgrading. You can use either:

- ◆ The `pkgadd` command on top of your existing software.
- ◆ The `upgrade_start/upgrade_finish` method to install VxVM 3.5.

Both methods have advantages and disadvantages, as shown in the following:

	Using <code>pkgadd</code>	Using <code>upgrade_start/upgrade_finish</code>
Advantages	Usually requires only 1 reboot	<ol style="list-style-type: none"> 1. There will be only a single VRTSvxvm package displayed by <code>pkginfo</code>. 2. VxVM configuration data is backed up. 3. VRTSlic can be deinstalled, if there are no other packages that require it.
Disadvantages	<ol style="list-style-type: none"> 1. The new package is displayed by <code>pkginfo</code> as VRTSvxvm.2. <p>Note You can avoid having multiple instances by specifying the <code>-a admin</code> option with <code>instance=overwrite</code>. See <code>pkgadd(1M)</code> for more details.</p> <ol style="list-style-type: none"> 2. You may not be able to remove VRTSlic if another installed VRTSvxvm package is still using the license. 	

Caution If your root disk is encapsulated, refer to “[Upgrading VxVM on an Encapsulated Root Disk](#)” on page 58.

These procedures follow.

Upgrade with `pkgadd` command

1. If you have not already obtained and installed a VxVM 3.5 license key, do so now. See “[Product Licensing](#)” on page 7 for details.

2. Bring the system down to single-user mode using the following command:

```
# init s
```

3. Stop the `vxconfigd` daemon:

```
# vxdctl stop
```

4. Stop the `vxiod` daemon:

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

5. (Optional) Remove the VMSA software with the following command:

```
# pkgrm VRTSvmsa
```

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

Note If you still have clients running previous versions of VxVM, you can skip this step.

6. Add the VxVM 3.5 licensing package with the following commands:

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvlic
```

Note VxVM 3.5 requires the 3.5 licensing package. You do not need to remove the previous versions of `VRTSlic`.

7. Add the VxVM 3.5 package with the following:

```
# pkgadd -d . VRTSvxvm
```

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

8. Using the `pkgadd` command, install the additional packages:

```
# pkgadd -d . VRTSvmdoc VRTSvmman
# pkgadd -a ../scripts/VRTSobadmin -d . VRTSob VRTSobgui
# pkgadd -d . VRTSfspro VRTSvmpro
```



Note It is recommended that you install both `VRTSfspro` and `VRTSvmpro` if you intend to use the VEA GUI.

9. Perform a reconfiguration reboot:

```
# cd /; shutdown -g0 -y -i6  
or  
# reboot -- -r
```

10. Importing a pre-3.5 VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.5 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.

11. See “[Setting Up and Starting VEA](#)” on page 46 to set up and restart VEA.

Upgrade with `upgrade_start/upgrade_finish` commands

Follow these steps:

1. If you have not already obtained and installed a VxVM 3.5 license key, do so now. See “[Product Licensing](#)” on page 7 for details.

2. Bring the system down to single-user mode using the following command.

```
# init s
```

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

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

Note You need to mount `/opt` manually if `/opt` is on its own partition.

4. Run the `upgrade_start -check` script to find any problems that exist which could prevent a successful upgrade. This script reports errors if found. Otherwise, it reports success and you can proceed with running the `upgrade_start` script.

```
# ./upgrade_start -check
```

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

```
# /cdrom/volume_manager/scripts/upgrade_start
```

The `upgrade_start` script searches for volumes containing file systems.

6. Remove the old VxVM packages by entering:

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

Note If you still have clients running old versions of VxVM, you can skip removing VMSA.

Note VxVM 3.5 requires the 3.5 licensing package. You do not need to remove the previous versions of VRTSlic.

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

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

8. Add the VxVM 3.5 licensing package with the following:

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvlic
```

9. Add the VxVM 3.5 package with the following:

```
# pkgadd -d . VRTSvxvm
```

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

10. Complete the upgrade using the following command:

```
# ../scripts/upgrade_finish
```

11. Perform reconfiguration reboot:

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

12. Using the `pkgadd` command, install the additional packages:

```
# cd /cdrom/volume_manager/pkgs
# pkgadd -d . VRTSvmdoc VRTSvmman
# pkgadd -a ../scripts/VRTSobadmin -d . VRTSob VRTSobgui
# pkgadd -d . VRTSfspro VRTSvmpro
```



13. Importing a pre-3.5 VxVM disk group does not automatically upgrade the disk group version to the VxVM 3.5 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.
14. See “[Setting Up and Starting VEA](#)” on page 46 to set up and restart VEA.

Upgrading Solaris Only

If you are running VxVM 3.5 with Solaris 2.6 or 7, you can upgrade Solaris using the following procedure.

Caution The root disk partitions should not be modified: this will cause `upgrade_start/upgrade_finish` to fail. If you change the root disk partitions, you cannot upgrade VxVM.

1. If the boot disk is mirrored, detach the mirror. You can use VEA and the `vxplex` command.
2. If any of the file systems `/`, `/usr`, `/var`, `/opt`, `/export/home`, `/var/adm` or `/usr/kvm` are defined on volumes, ensure that at least one plex for each of those volumes is formed from a single subdisk that begins on a cylinder boundary. This is necessary because part of the upgrade process involves temporarily converting file systems on volumes back to using direct disk partitions, and Solaris requires that disk partitions start on cylinder boundaries.

The upgrade scripts (located in the `/scripts` directory on the CD) automatically convert file systems on volumes back to using regular disk partitions, as necessary. If the upgrade scripts detect any problems (such as lack of cylinder alignment), they display an explanation of the problem and the upgrade does not proceed.

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

3. Bring the system down to single-user mode using the following command:

```
# init s
```

Note You need to mount `/opt` manually if `/opt` is on its own partition.

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


```
# /etc/init.d/volmgt start
```
5. Run the `upgrade_start -check` script to find any problems that exist which could prevent a successful upgrade. This script reports errors if found. Otherwise, it reports success and you can proceed with running the `upgrade_start` script.



```
# ./upgrade_start -check
```

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

```
# /cdrom/volume_manager/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.

7. Reboot to single user mode:

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

8. Upgrade the operating system to Solaris 7 or 8. Refer to the Solaris installation documentation for instructions on how to upgrade Solaris. After installing Solaris, install the required patches shown in “[Solaris Patches for VxVM](#)” on page 18.

9. Reboot to single-user mode:

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

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

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

11. Complete the upgrade by entering:

```
# /cdrom/volume_manager/scripts/upgrade_finish
```

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

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:

```
root 458 180 10:23:34? 0:00 vxrelocd root
root 414 180 10:23:06? 0:03 vxconfigd
root 478 414 10:23:37? 0:00 vxrelocd root
root 479 478 5 10:23:37? 0:00 vxnotify -f -w 15
```

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:

```
# vxctl 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.

Disabling Hot-Relocation

The hot-relocation feature detects disk failures automatically and notifies you of the nature of the failure, attempts to relocate any affected subdisks that are redundant, and initiates recovery procedures. The hot-relocation feature is enabled by default and it is recommended that you leave it on.

However, you can disable it by preventing the `vxrelocd` daemon from starting up during system startup. This should be done after the VxVM packages have been installed.

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

```
# vxrelocd root &
```

Hot-Relocation Guidelines

Refer to the following guidelines when using hot-relocation.

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



- ◆ Although hot-relocation does not require you to designate disks as spares, designate at least one disk as a spare within each disk group. This gives you some control over which disks are used for relocation. If no spares exist, VxVM uses any available free space within the disk group. When free space is used for relocation purposes, it is possible to have performance degradation after the relocation.
- ◆ After hot-relocation occurs, designate one or more additional disks as spares to augment the spare space. Some of the original spare space may be occupied by relocated subdisks.
- ◆ If a given disk group spans multiple controllers and has more than one spare disk, set up the spare disks on different controllers (in case one of the controllers fails).
- ◆ For a mirrored volume, configure the disk group so that there is at least one disk that does not already contain a mirror of the volume. This disk should either be a spare disk with some available space or a regular disk with some free space and the disk is not excluded from hot-relocation use.
- ◆ For a mirrored and striped volume, configure the disk group so that at least one disk does not already contain one of the mirrors of the volume or another subdisk in the striped plex. This disk should either be a spare disk with some available space or a regular disk with some free space and the disk is not excluded from hot-relocation use.
- ◆ For a RAID-5 volume, configure the disk group so that at least one disk does not already contain the RAID-5 plex (or one of its log plexes) of the volume. This disk should either be a spare disk with some available space or a regular disk with some free space and the disk is not excluded from hot-relocation use.
- ◆ If a mirrored volume has a DRL log subdisk as part of its data plex, you cannot relocate the data plex. Instead, place log subdisks in log plexes that contain no data.
- ◆ Hot-relocation does not guarantee to preserve the original performance characteristics or data layout. Examine the locations of newly-relocated subdisks to determine whether they should be relocated to more suitable disks to regain the original performance benefits.
- ◆ Hot-relocation is capable of creating a new mirror of the root disk if the root disk is mirrored and it fails. The rootdg should contain sufficient contiguous spare or free space to accommodate the volumes on the root disk. (`rootvol` and `swapvol` require contiguous disk space.)
- ◆ Although it is possible to build VxVM objects on spare disks (using `vxmake` or the VEA 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.

RAID-5 Guidelines

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

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

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



- ◆ The RAID-5 plex must be derived from at least two subdisks on two or more physical disks. If any log plexes exist, they must belong to disks other than those used for the RAID-5 plex.
- ◆ RAID-5 logs can be mirrored and striped.
- ◆ If the volume length is not explicitly specified, it is set to the length of any RAID-5 plex associated with the volume; otherwise, it is set to zero. If you specify the volume length, it must be a multiple of the stripe unit size of the associated RAID-5 plex, if any.
- ◆ If the log length is not explicitly specified, it is set to the length of the smallest RAID-5 log plex that is associated, if any. If no RAID-5 log plexes are associated, it is set to zero.
- ◆ Sparse RAID-5 log plexes are not valid.

Protecting Your System and Data

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

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

The following are suggestions for protecting your system and data:

- ◆ Place the disk containing the root file system (the root or boot disk) under VxVM control through encapsulation. Encapsulation converts the `root` and `swap` devices to volumes (`rootvol` and `swapvol`).
- ◆ Mirror the root disk so that an alternate root disk exists for booting purposes. By mirroring disks critical to booting, you ensure that no single disk failure leaves your system unbootable and unusable.

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

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

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

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

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

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

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

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

- ◆ Use logging to prevent corruption of recovery data in RAID-5 volumes. Make sure that each RAID-5 volume has at least one log plex.
- ◆ Perform regular backups to protect your data. Backups are necessary if all copies of a volume are lost or corrupted. Power surges can damage several (or all) disks on your system. Also, typing a command in error can remove critical files or damage a file system directly. Performing regular backups ensures that lost or corrupted data is available to be retrieved.



Uninstalling VxVM

Before uninstalling, you should move volumes from VxVM control. If the root volume is encapsulated, use the procedures in [“Moving Volumes From an Encapsulated Root Disk”](#) on page 82. If your root disk is not encapsulated, you can start with the procedures in [“Moving Volumes To Disk Partitions.”](#)

Moving Volumes From an Encapsulated Root Disk

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

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

For information about the plex, use the following command:

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

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

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

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

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

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

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

Continue to [“Moving Volumes To Disk Partitions”](#) on page 82 to complete the removal.

Moving Volumes To Disk Partitions

1. All volumes must be moved to disk partitions. This can be done using one of these procedures:
 - Back up the system fully onto tape and then recover from it.

- Back up each file system individually and then recover them all after creating new file systems on disk partitions.
- Use VxVM to move volumes incrementally onto disk partitions. Use the following instructions to do this.

Move Volumes onto Disk Partitions Using VxVM

1. Evacuate disks using the `vxdiskadm` program, VEA, or the `vxevac` script. You should consider the amount of target disk space required for this before you begin. Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.
2. Remove the evacuated disks from VxVM control using the following commands:


```
# vxdg rmdisk disk_name
# vxdisk rm devname
```
3. Decide which volume to move first. If the volume to be moved is mounted, unmount it.
4. If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that data on the volume is synced.
5. Create a partition on free disk space of the same size as the volume. If there is not enough free space for the partition, a new disk must be added to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this volume.
6. Copy the data on the volume onto the newly created disk partition using a command similar to the following:


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

 where `c2t2d2` is the disk outside of VxVM and `s7` is the newly created partition.
7. Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly created partition.
8. Mount the disk partition if the corresponding volume was previously mounted.
9. Remove the volume from the VxVM using the following command:


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



10. Remove any disks that have become free (those having no subdisks defined on them) by removing the volumes from VERITAS Volume Manager control.

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

```
# vxprint -F '%sdnum' diskname
```

If the output is not 0, there are still some subdisks on this disk that you need to remove. If the output is 0, remove the disk from VERITAS Volume Manager control as follows:

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

Use the free space created for adding the data from the next volume you want to remove.

11. After you successfully convert all volumes into disk partitions, reboot the system.
12. After the reboot, make sure none of the volumes are open by using the following command:

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

If any volumes remain open, repeat the steps 4 to 12, above.

“[Moving Volumes onto Partitions - Example](#)” on page 75 is an example showing how to move the data on a volume to a disk partition.

Shutting Down VxVM

To shut down VxVM, use the following commands:

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

Removing the VERITAS Volume Manager Packages

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

To check for other installed VERITAS products:

```
pkginfo | grep -i vrts
```

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

Removing the VERITAS Volume Manager

To remove the VxVM installed packages, use the following command:

```
# pkgrm VRTSvmdoc VRTSvmman VRTSvxvm VRTSvlic
```

If VxVM 2.x is still installed and you want to remove VxVM 3.5, 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 VRTSvmdoc VRTSvmman \
    VRTSvxvm.2 VRTSvxvm VRTSlic VRTSvlic
```

Removing the VERITAS Provider Packages

To remove the VERITAS Virtual Disk Management Services Provider and VERITAS File System Provider packages, use the following commands:

```
# pkgrm VRTSvmpro
# pkgrm VRTSfspro
```

Removing the VERITAS Enterprise Administrator

To remove the VERITAS Enterprise Administrator packages, use the following commands:

```
# pkgrm VRTSob
# pkgrm VRTSobgui
```

Removing the Windows Client Software Package

To remove the Windows Client software, you:

Click **Start > Settings > Control Panel > Add/Remove Software**

Choose **VERITAS Enterprise Administrator** for removal.

Verifying Removal

To verify the removal of the packages:

```
# pkginfo | grep VRTS
```



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

This should complete the removal of VERITAS Volume Manager.

Restoring VERITAS Volume Manager 3.5 Configuration

You can restore VxVM configuration information from the backup directory `/VXVM-CFG-BAK` after VxVM 3.5 has been removed by `pkgrm`.

Note This feature applies only to VxVM 3.5. If you used `pkgrm` to remove versions VxVM prior to 3.5, the configuration information is not saved by `pkgrm`.

Note This feature is available only when VxVM is removed using the stand-alone `pkgrm` command. That is, it is not available when `pkgrm` is used as part of the `upgrade_start` script.

The following files are available for restoration (using `pkgadd`) from directory `/VXVM-CFG-BAK`:

- `/etc/vx/volboot`
- `/etc/vx/array.info`
- `/kernel/drv/vxdmp.conf`
- `/kernel/drv/vxio.conf`
- `/etc/vx/jbod.info`
- `/etc/vx/ddl.support`
- `/etc/vx/vxddl.exclude`
- `/etc/vx/vxvm.exclude`
- `/etc/vx/vxdmp.exclude`
- `/etc/vx/vvrports`

Moving Volumes onto Partitions - Example

This example shows how to move the data on a volume to a disk partition. In the example, there are three disks: `disk1` and `disk2` are subdisks on volume `vol01` and `disk3` is a free disk. The data on `vol01` is copied to `disk3` using `vxevac`.

1. Diskgroup `voldg` content before the data on `vol01` is copied to `disk3`.

```
# vxprint -g voldg -ht
DG NAME      NCONFIG      NLOG      MINORS      GROUP-ID
DM NAME      DEVICE      TYPE      PRIVLEN     PUBLEN     STATE
RV NAME      RLINK_CNT    KSTATE    STATE       PRIMARY     DATAVOLS  SRL
RL NAME      RVG          KSTATE    STATE       REM_HOST    REM_DG     REM_RLNK
V NAME       RVG          KSTATE    STATE       LENGTH      READPOL    PREFPLEX
UTYPE
PL NAME      VOLUME      KSTATE    STATE       LENGTH      LAYOUT     NCOL/WID
MODE
SD NAME      PLEX        DISK      DISKOFFS    LENGTH      [COL/]OFF  DEVICE
MODE
SV NAME      PLEX        VOLNAME   NVOLLAYR    LENGTH      [COL/]OFF  AM/NM
MODE
DC NAME      PARENTVOL    LOGVOL
SP NAME      SNAPVOL      DCO

dg voldg      default      default     115000
1017856044.1141.hostname.veritas.com

dm disk1      c1t12d0s2    sliced     2591        17900352   -
dm disk2      c1t14d0s2    sliced     2591        17899056   -
dm disk3      c1t3d0s2     sliced     2591        17899056   -

v vol1        -            ENABLED    ACTIVE      4196448    ROUND      -
fsgen
pl pl1        vol1         ENABLED    ACTIVE      4196448    CONCAT     -
RW
sd sd1        pl1          disk1      0           2098224    0          c1t12d0
ENA
sd sd2        pl1          disk2      0           2098224    2098224    c1t14d0
ENA
```

2. Evacuate `disk1` to `disk3`.

```
# /etc/vx/bin/vxevac -g voldg disk1 disk3
```



```
# vxprint -g voldg -ht
```

DG NAME	NCONFIG	NLOG	MINORS	GROUP-ID		
DM NAME	DEVICE	TYPE	PRIVLEN	PUBLEN	STATE	
RV NAME	RLINK_CNT	KSTATE	STATE	PRIMARY	DATAVOLS	SRL
RL NAME	RVG	KSTATE	STATE	REM_HOST	REM_DG	REM_RLNK
V NAME	RVG	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
UTYPE						
PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID
MODE						
SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE
MODE						
SV NAME	PLEX	VOLNAME	NVOLLAYR	LENGTH	[COL/]OFF	AM/NM
MODE						
DC NAME	PARENTVOL	LOGVOL				
SP NAME	SNAPVOL	DCO				

```
dg voldg          default      default  115000
1017856044.1141.hostname.veritas.com
```

dm disk1	c1t12d0s2	sliced	2591	17900352	-	
dm disk2	c1t14d0s2	sliced	2591	17899056	-	
dm disk3	c1t3d0s2	sliced	2591	17899056	-	

v voll	-	ENABLED	ACTIVE	4196448	ROUND	-
fsngen						
pl pll	voll	ENABLED	ACTIVE	4196448	CONCAT	-
RW						
sd disk3-01	pll	disk3	0	2098224	0	c1t3d0
ENA						
sd sd2	pll	disk2	0	2098224	2098224	c1t14d0
ENA						

3. Evacuate disk2 to disk3.

```
# /etc/vx/bin/vxevac -g voldg disk2 disk3
```

```
# vxprint -g voldg -ht
```

DG NAME	NCONFIG	NLOG	MINORS	GROUP-ID		
DM NAME	DEVICE	TYPE	PRIVLEN	PUBLEN	STATE	
RV NAME	RLINK_CNT	KSTATE	STATE	PRIMARY	DATAVOLS	SRL
RL NAME	RVG	KSTATE	STATE	REM_HOST	REM_DG	REM_RLNK
V NAME	RVG	KSTATE	STATE	LENGTH	READPOL	PREFPLEX
UTYPE						
PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID
MODE						
SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE
MODE						
SV NAME	PLEX	VOLNAME	NVOLLAYR	LENGTH	[COL/]OFF	AM/NM
MODE						
DC NAME	PARENTVOL	LOGVOL				
SP NAME	SNAPVOL	DCO				

```
dg voldg          default      default  115000
```



```
1017856044.1141.hostname.veritas.com
```

```
dm disk1      c1t12d0s2    sliced  2591    17900352 -
dm disk2      c1t14d0s2    sliced  2591    17899056 -
dm disk3      c1t3d0s2     sliced  2591    17899056 -

v  vol1       -                ENABLED ACTIVE  4196448  ROUND  -
fsgen
pl pl1        vol1            ENABLED ACTIVE  4196448  CONCAT -
RW
sd disk3-01   pl1            disk3    0        2098224  0        c1t3d0
ENA
sd disk3-02   pl1            disk3    2098224  2098224  2098224  c1t3d0
ENA
```

4. Remove the evacuated disks from VERITAS Volume Manager control.

```
# vxdisk -g voldg list
DEVICE      TYPE      DISK      GROUP      STATUS
c1t3d0s2    sliced   disk3     voldg      online
c1t12d0s2    sliced   disk1     voldg      online
c1t14d0s2    sliced   disk2     voldg      online

# vxvg rmdisk disk1
# vxvg rmdisk disk2
# vxdisk rm c1t12d0
# vxdisk rm c1t14d0
```

5. Verify that the evacuated disks have been removed from VERITAS Volume Manager control.

```
# vxdisk -g voldg list
DEVICE      TYPE      DISK      GROUP      STATUS
c1t3d0s2    sliced   disk3     voldg      online
```

6. Check to see whether the volume you want to move first is mounted.

```
# mount | grep vol1
/vol1 on /dev/vx/dsk/voldg/vol1
read/write/setuid/log/nolargefiles/dev=12dc138 on Wed Apr  3
10:13:11 2002
```

7. Create a partition on free disk space of the same size as the volume. In this example, a 2G partition is created on disk1 (c1t12d0s1).

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
```



```

0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
  /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
1. clt3d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
  /sbus@1f,0/SUNW,fas@2,8800000/sd@3,0
2. clt9d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
  /sbus@1f,0/SUNW,fas@2,8800000/sd@9,0
3. clt10d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
  /sbus@1f,0/SUNW,fas@2,8800000/sd@a,0
4. clt11d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
  /sbus@1f,0/SUNW,fas@2,8800000/sd@b,0
5. clt12d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
  /sbus@1f,0/SUNW,fas@2,8800000/sd@c,0
6. clt14d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
  /sbus@1f,0/SUNW,fas@2,8800000/sd@e,0
7. clt15d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
  /sbus@1f,0/SUNW,fas@2,8800000/sd@f,0
Specify disk (enter its number): 5
selecting clt12d0
[disk formatted]

```

FORMAT MENU:

```

disk      - select a disk
type      - select (define) a disk type
partition - select (define) a partition table
current   - describe the current disk
format    - format and analyze the disk
repair    - repair a defective sector
label     - write label to the disk
analyze   - surface analysis
defect    - defect list management
backup    - search for backup labels
verify    - read and display labels
save      - save new disk/partition definitions
inquiry   - show vendor, product and revision
volname   - set 8-character volume name
!<cmd>    - execute <cmd>, then return
quit
format> p

```

PARTITION MENU:

```

0      - change `0' partition
1      - change `1' partition
2      - change `2' partition
3      - change `3' partition
4      - change `4' partition
5      - change `5' partition
6      - change `6' partition
7      - change `7' partition
select - select a predefined table
modify - modify a predefined partition table
name   - name the current table
print  - display the current table
label  - write partition map and label to the disk
!<cmd> - execute <cmd>, then return

```



```

quit
partition> 1
Part      Tag      Flag      Cylinders      Size      Blocks
  1 unassigned    wm          0          0      (0/0/0)          0
Enter partition id tag[unassigned]:
Enter partition permission flags[wm]:
Enter new starting cyl[0]:
Enter partition size[0b, 0c, 0.00mb, 0.00gb]: 2.00gb
partition> 1
Ready to label disk, continue? y

partition> p
Current partition table (unnamed):
Total disk cylinders available: 13814 + 2 (reserved cylinders)

Part      Tag      Flag      Cylinders      Size      Blocks
  0 unassigned    wm          0          0      (0/0/0)          0
  1 unassigned    wm      0 - 3236      2.00GB      (3237/0/0)    4195152
partition> q

```

8. Copy the data on vol01 to the newly created disk partition.

```
# dd if=/dev/vx/dsk/voldg/vol01 of=/dev/dsk/c1t12d0s1
```

9. In the /etc/vfstab file, remove the following entry.

```
/dev/vx/dsk/voldg/vol1 /dev/vx/rdisk/voldg/vol1 /vol1 vxfs 4 yes rw
```

10. Replace it with an entry for the newly created partition.

```
/dev/dsk/c1t12d0s1 /dev/rdisk/c1t12d0s1 /vol01 vxfs 4 yes rw
```

11. Mount the disk partition.

```
# mount -F vxfs /dev/dsk/c1t12d0s1 /vol01
```

12. Remove vol01 from VERITAS Volume Manager.

```
# vxedit -rf rm /dev/vx/dsk/voldg/vol01
```

13. Remove any disks that have become free (those having no subdisks defined on them) by removing the volumes from VERITAS Volume Manager control.

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

```
# vxprint -F '%sdnum' diskname
```

If the output is not 0, there are still some subdisks on this disk that you need to remove. If the output is 0, remove the disk from VERITAS Volume Manager control as follows:

```
# vxdg rmdisk diskname
```



```
# vxdisk rm devname
```

Use the free space created for adding the data from the next volume you want to remove.

14. After you successfully convert all volumes into disk partitions, reboot the system.
15. After the reboot, make sure none of the volumes are open by using the following command:

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

Index

- A**
 - A5x00 devices 13
 - AP (alternate pathing) 14
 - arrays 5
- B**
 - boot disk 25
 - encapsulated 29
 - encapsulation requirements 11
 - restrictions 10
- C**
 - cluster functionality
 - enabling 50
 - environment requirements 14
 - shared disks 51
 - cluster requirements 14
 - configuration daemon (vxconfigd), starting 73
 - configuring
 - new disks 51
 - shared disks 51
 - controllers, mirroring guidelines 78
 - custom installation method 35
- D**
 - daemons, checking 72
 - data
 - protecting 80
 - deleting VxVM packages 84
 - devices
 - suppress devices 40
 - disk arrays 5
 - disk drives
 - variable geometry 77
 - disk groups 3
 - disks
 - initializing all 29
 - initializing selected 35
 - placing in other disk group 25
 - rootdg 25
 - stripe unit size 77
 - types of 10
- DMP**
 - prevent multipathing 40
- DMP (dynamic multipathing)**
 - AP co-existence 14
 - preventing 40
 - SENA category devices 14
- driver, VxDMP 20**
- DRL guidelines 78**
- E**
 - encapsulated disks 10
 - encapsulating
 - all disks 36
 - boot disks 29
- G**
 - guidelines
 - DRL 78
 - mirroring 78
 - mirroring and striping 79
 - RAID-5 79
- H**
 - hot-relocation 12, 75
- I**
 - I/O daemon (vxiod), starting 74
 - initialized disks 10
 - initializing
 - all disks 37
 - installing
 - individual disks 38
 - VEA client (Windows) 48
 - VEA server (UNIX) 47
- L**
 - layouts, types of 4
 - license keys, obtaining 7



-
- license, viewing 9
 - log subdisks 78
 - logical objects 3
 - M**
 - MANPATH variable 7
 - mirroring
 - and striping guidelines 79
 - guidelines 78
 - root disk 80
 - mirroring controllers 78
 - N**
 - new disks, configuring 51
 - O**
 - overlapped seeks 78
 - P**
 - patches
 - A5x00 devices 13
 - PATH variable 7
 - pkgadd 58, 59, 66
 - pkgrm command 84
 - plexes 4
 - Prevent Multipathing/Suppress Devices
 - from VxVM's view 40
 - preventive maintenance 80
 - private region 3
 - processes, checking 72
 - protecting data 80
 - public region 2
 - Q**
 - quick installation method 29
 - R**
 - RAID-5
 - guidelines 79
 - removing VxVM packages 84
 - restoring VxVM configuration 86
 - root disk
 - alternate 80
 - rootdg 25
 - rootdg (root disk group) 3
 - S**
 - SANs, configuring 15
 - shared disks, configuring 51
 - shutting down VxVM 84
 - Solaris upgrade 57
 - starting vxconfigd configuration daemon 73
 - starting vxiod daemon 74
 - storage layouts 4
 - storage management 1
 - stripe unit size recommendations 77
 - striping guidelines 77
 - subdisks 3
 - physical disk placement 79
 - RAID-5 plex, configuring 80
 - suppress devices 40
 - swap relocation 11
 - U**
 - upgrade_finish 66, 71
 - upgrade_start 60, 71
 - upgrading
 - VxVM and Solaris 58, 63
 - VxVM only 66
 - VxVM with AP 20
 - upgrading to 3.5 54
 - V**
 - VEA 46
 - client, starting 49
 - VEA client (UNIX), installing 47
 - VEA client (Windows), installing 48
 - VM disks 2
 - volume length, RAID-5 guidelines 80
 - volumes 4
 - DRL 78
 - VRTSexplorer ix
 - vxassist mirror command 81
 - vxconfigd configuration daemon,
 - starting 73
 - vxctl mode command 73
 - VxDMP driver 20
 - vxinstall program 26
 - vxinstall program, running 25
 - vxiod I/O daemon, starting 74
 - VxVM
 - obtaining system information ix
 - VxVM configuration
 - restoring 86
 - VxVM packages, removing 84