VERITAS Volume Manager 3.5 Installation Guide

HP-UX 11i v2

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

Contents

Contents

5. Getting Started with VxVM

6. Uninstalling the VERITAS Volume Manager

Preface

Introduction

This document provides information on installing and initializing VERITAS Volume ManagerTM (VxVMTM) and VERITAS Enterprise AdministratorTM (VEATM)).

Organization

This guide is organized with the following chapters:

Related Documents

The following documents contain related information:

- VERITAS Volume Manager Migration Guide
- VERITAS Volume Manager Release Notes
- VERITAS Volume Manager Hardware Notes

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.

Conventions

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

Table 1   **Typographic Conventions**

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

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

NOTE   "Dynamic Disk Groups" were formerly known as "Disk Groups." They are still sometimes referred to as "Disk Groups" in VxVM documentation, CLI and menu displays.

Getting Help

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

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

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

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

For software updates:

- Email: swupdate@veritas.com

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

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

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

- http://www.veritas.com

Using VRTSexplorer

The `VRTSexplorer` program can help VERITAS Technical Support engineers diagnose the cause of technical problems associated with VERITAS products. You can download this program from the VERITAS FTP site or install it from the VERITAS Installation CD. For more information, consult the README file in the support directory on the VERITAS Installation CD, and the VERITAS *Volume Manager Release Note*

1 Before You Start

The following topics are covered in this chapter:

- "Installation and Upgrade Overview" on page 2

- "Installation and Upgrade Roadmaps" on page 3

- "Installation and Upgrade Requirements" on page 5

- "Preparing to Install VxVM on a System with LVM Volume Groups" on page 11

Installation and Upgrade Overview

The section gives an overview of the installation and upgrade processes.

VxVM Installation

Installing the VxVM for the first time involves the following steps:

Step 1. Familiarize yourself with the section "Installation and Upgrade Requirements" on page 5

Step 2. If you are installing on a system that has existing LVM volume groups, prepare for the installation by following the procedures in "Preparing to Install VxVM on a System with LVM Volume Groups" on page 11.

Step 3. Install the VxVM software packages onto the system. See Chapter 2, "Installing VERITAS Volume Manager," on page 25 for the procedure.

Step 4. For new installations, initialize the VxVM using vxinstall. See Chapter 3, "Setting up the VxVM Environment," on page 35 for the procedure.

Step 5. For upgrading, there are a number of scenarios; including upgrading from HP-UX 11i v2. These are discussed in Chapter 4, "Upgrading the VERITAS Volume Manager," on page 47."

Step 6. If you are planning to use Volume Manager in a cluster environment, follow the instructions in "Enabling Cluster Support in VxVM" on page 57.

Step 7. After installation, add any disk arrays that are unsupported by VERITAS to the JBOD category as described in the section "Using Hot-Relocation" on page 56.

VEA Set Up Overview

The VERITAS Enterprise Administrator, VEA, is the graphical user interface for the VxVM. It is installed by default as part of the Base VxVM 3.5 bundle with all HP 11i v2 Operating Environments.

The VEA has two parts: a server and a client. The client is the graphical user interface. The server must be installed and run on the machine to be administered, which must be an HP-UX 11i v2 machine running the VxVM. The VEA client must be installed and run on a machine that supports the Java 1.1 Runtime Environment. See "Installing VEA" on page 29.

The minimum recommendation for the client is a Hewlett-Packard D-class machine with 512M memory.

If you are planning to run the VEA client (the GUI) on the machine being administered, then all the software you need is installed with the HP Base-VXVM bundle.

You need to start the VEA server on the HP-UX machine on which it is installed (see "Starting the VEA Service" on page 54), and start the client on the machine on which it is installed (see "Starting the VEA Client" on page 55).

If users other than root need to access the VEA, set up security and specify which users can run the VEA. See "Starting the VEA Service" on page 54 for the procedure.

Installation and Upgrade Roadmaps

This section gives an overview of the installation and upgrade procedures:

• "Installing VxVM on HP-UX 11i v2" on page 3

• See "Installation and Upgrade Requirements" on page 5

Installing VxVM on HP-UX 11i v2

VxVM is part of the Operating Environment (OE), and is installed by default on every new OE installation. With the new OEUR(Operating Environment Update Release) release for HP-UX 11 v2, you also have the option of selecting VxVM for the root disk when installing using the HP Ignite-UX installation software. If you choose this option, the root disk group will be initialized by the Ignite-UX installation process, and there will be no need to run the vxinstall script.

If you choose to have an LVM disk as your root disk, or if you are upgrading from a previous version of VxVM, you can use the vxcp_lvmroot(1M) command to create a VxVM root disk. The VxVM root disk can then replace or co-exist with your LVM disk (that is, you can boot from either the LVM or the VxVM root disk).

If you choose to be rooted from LVM disks and want to bring up VxVM, then you need to run vxinstall to initialize the VxVM environment and get it running.

Upgrading to VxVM 3.5

Your upgrade to VxVM 3.5 will fall into one of the following categories:

Step 1. The root disk is controlled by VxVM 3.5. VxVM rootability can be achieved after the upgrade by using the vxcp_lvmroot(1M) command to clone the LVM root disk.

Step 2. The root disk is not controlled by VxVM 3.5.

Upgrading from a HP-UX 11.22 LVM rootable system, with VxVM configured, to HP-UX 11.23 is dealt with in detail in Chapter 4.

VxVM Controlling the Root Disk

Starting with HP-UX 11.22, there are two possibilities for upgrade:

• With LVM root: use `vxcp_lvmroot` to achieve VxVM root

• With VxVM (3.1) root: retain VxVM rootable.

VxVM Not Controlling the Root Disk

The upgrade procedure allows you to retain your existing VxVM configuration without running the `vxinstall` script. After upgrading, you can resume using VxVM with the same configuration as before.

NOTE   On importing a 3.1 VxVM disk group, you are given the option to upgrade the disk group version to the VxVM 3.5 level. Each VxVM 3.5 disk group can also be manually upgraded following the VxVM upgrade. See the chapter "Creating and Administering Disk Groups" in the VERITAS Volume Manager 3.5 Administrator's Guide for information on disk group versioning and upgrading disk groups.

Summary of VxVM Products

VxVM consists of several related software packages as shown in the following tables:

Table 1-1   HP-UX 11i v2 September 2004 Release Software Bundle

Product	Name	Availability
Base VxVM	Base VxVM 3.5 for HP-UX 11i v2	Free with HP-UX 11i v2 September 2004 Application Software, HP-UX 11i v2 Internet, all HP-UX 11i v2 Operating Environments
B9116AA	VxVM 3.5 full features for HP-UX 11i v2 September 2004 Release	Purchase license to enable VXVM-FULL; use codeword to unlock from HP-UX 11i v2 Application Software
B9117AA	VERITAS Cluster Volume Manager 3.5 for HP-UX 11i v2 September 2004 Release	Purchase license to enable VXVM-CVM; use codeword to unlock from HP-UX 11i v2 Application Software

NOTE   The HP Base-VXVM bundle consists of the following products: VRTSvxvm, VRTSvmdoc, VRTSvlic, VRTSob, VRTSobgui, VRTSvmpro, and VRTSfspro.:

Table 1-2   VERITAS VxVM Software Packages

Product	Description
VRTSvxvm	VERITAS VxVM 3.5 for HP-UX 11i v2 (Drivers and Utilities)
VRTSvlic	SIG Licensing Package (Licensing Utilities)
VRTSvmdoc	VERITAS VxVM 3.5 documentation
VRTSob	VERITAS Enterprise Administrator Service
VRTSobgui	VERITAS Enterprise Administrator Client
VRTSvmpro	VERITAS Volume Manager Management Services Provider
VRTSfspro	VERITAS File System Management Provider

Installation and Upgrade Requirements

Before you install or upgrade, you need to check that:

Step **1.** The administrator who performs the upgrade has root access and a working knowledge of HP-UX administration.

Step **2.** Sufficient outage time exists for completion of this procedure (which may require several system reboots)

The following actions need to be performed before starting installation or upgrade:

- See "Check Disk Space" on page 5
- See "Install HP-UX Operating System Patches" on page 9

Check Disk Space

Before installing the VxVM software and related patches, verify that there is sufficient disk space to contain all of the software. The SD-UX tools check space available and prevent installation if a space problem is detected.

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

Table 1-3 **Space Requirements**

/ (MB)	/Stand (MB)	/usr (MB)	/opt (MB)	/var (MB)
38	4	42	89	1

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 /usr/share/man.

Examples:

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

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

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

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

Product Licensing

The VERITAS Volume Manager is a licensed product. Before you install the VERITAS Volume Manager 3.5 software packages, you need to 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:

https://vlicense.veritas.com/VLicenseLogin.jhtml

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 one of the phone numbers listed under "Getting Help" on page 8.

Add or Upgrade Licenses

Volume Manager 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:

Table 1-4   **Equivalency Table**

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

Any existing valid VxVM license keys are accepted for backward compatibility on upgrading to Volume Manager 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 Volume Manager 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 Volume Manager 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/sbin/vxliccheck -pv
```

If you need to add new licenses to support features in Volume Manager 3.5, you must add the new licensing package, VRTSvlic.

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.

Upgrading VxVM

You cannot upgrade the VRTSvxvm (or Base-VXVM) and license packages at the same time, since the license package does not support upgrade.

Installing VRTSob, VRTSvmpro, and VRTSfspro

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.

These packages are installed by default when installing HP-UX 11i v2 , or when upgrading to HP-UX 11i v2.

VEA Requirements

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

- VEA server (must run under HP-UX 11i v2). See "Installing the VEA Server and Client on HP-UX" on page 29 on how to install the server.

- VxVM Release 3.5. The vxconfigd and vxsvc daemons must be running on the machine to be administered.

Cluster Environment Requirements

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

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

Step   **2.** Decide where to place the rootdg disk group for each node in the cluster. A system using Volume Manager 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 Volume Manager is running.

Step   **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.

Step   **4.** If you plan to use Dirty Region Logging (DRL) with Volume Manager 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.

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

Check the Operating System Environment

This release of Volume Manager can only be installed on an HP-UX system running the HP-UX 11i v2 operating system. Installation of VxVM does not occur if the system is not up to this release level. For example, installation is not possible on HP-UX 11.0 or HP-UX 11.11i. VxVM 3.5 supports 64-bit HP-UX 11i v2.

To ensure that the operating system is running at release level 11i v2, use the following command on a 64-bit machine:

```
# swlist -l bundle HPUXBaseOS
```

The output must contain:

```
HPUXBaseOS B.11.23 HP-UX Base OS
```

On a 32-bit machine, use the following command:

```
# swlist -l bundle HPUXBase32
```

The output must contain:

```
HPUXBase32 B.11.11HP-UX 32-bit Base OS
```

NOTE	If your system is running at a release level earlier than HP-UX 11i v2, you must upgrade to HP-UX 11.22 before upgrading to HP-UX 1i v2. If your system is running HP-UX 11i v2, you can just install HP-UX 11i v2.

Install HP-UX Operating System Patches

The following table shows all patches required for VxVM, and indicates which patches are mandatory and which are optional. Optional patches can provide functionality without which VxVM does not operate correctly. Lack of these patches does not harm your system in any way; however, the lack of mandatory patches does cause harm. To ensure maximum system efficiency, install all patches. VxVM may also require updates to the HP-UX operating environment in order for VxVM rootability to work properly. These updates are available as patches through standard HP-UX patch distribution facilities on the HP patch distribution sites.

NOTE	The required VxVM rootability patches are installed by default with the HP-UX 0902 Operating Environment Update Release.

The following table shows all patches required for VxVM 3.5.:

Table 1-5 **Required Patches**

Patch Number	Description	Mandatory/Optional
PHKL_26452	scsi3/disc30 patch (includes passthru ioctl)	Mandatory
PHKL_26519	scsi IO Subsystem cumulative patch	Mandatory
PHKL_19246	vm_superpage fix	Mandatory
PHKL_27096	Patch for VxVM rootability changes to HP-UX Kernel	Mandatory (needed for VxVM 3.5 installation)
PHCO_27100	Patch for VxVM rootability changes to HFS mkfs command	Optional
PHCO_27101	Patch for VxVM rootability changes to mkboot commands	Mandatory (needed for VxVM 3.5 installation)

Table 1-5 **Required Patches (Continued)**

Patch Number	Description	Mandatory/Optional
PHCO_27209	Patch for VxVM rootability changes to the insf command	Mandatory (needed for VxVM 3.5 installation)
PHCO_27099	Patch for VxVM rootability changes to LVM commands	Optional
PHCO_27103	Patch for VxVM rootability changes to /sbin/pre_init_rc	Mandatory
PHCO_26736	ioinitrc(iM) patch	Mandatory

Specific patch numbers are correct at the time of this release. If subsequent patches supersede any of these, the required patch numbers will change. Search mechanisms provided through Hewlett Packard's patch distribution sites help you locate the required patches using the patch numbers provided here. In addition, VERITAS provides a list of currently required patches on its web site, support.veritas.com. The table of current patches is located in the Volume Manager, Volume Manager for UNIX, Knowledge Base. However, you should visit the HP patch hub website to obtain the latest patch information.

NOTE   VxVM package scripts confirm that the mandatory HP-UX patches are installed on your system before allowing VxVM installation. You are also warned if any of the optional patches are not found. (VxVM 3.5, whether rootable, or non-rootable, itself requires no patches.)

Preparing to Install VxVM on a System with LVM Volume Groups

You can use the vxvmconvert command to migrate LVM volume groups to VxVM disk groups (apart from the root volume group).

NOTE
You can enable a VxVM rootable system by using the vxcp_lvmroot command to clone the LVM-based root disk to a VxVM root disk.

If any disk that you want to bring under VxVM control is currently under LVM control, execute the following commands before running vxinstall:

Step   1. vgreduce to remove the disk from any LVM volume groups to which it belongs

Step   2. lvremove to remove any LVM volumes using the disk

Step   3. pvremove to erase the LVM disk headers

NOTE
If the disk to be removed is the last disk in the volume group, use vgremove to remove the volume group and then use pvremove to erase the LVM disk headers.

If the disk is not currently in use by any volume or volume group, but has been initialized by pvcreate, you must still use pvremove to remove LVM disk headers.

The pvremove command is an LVM command designed to aid LVM and VxVM coexistence on a system.

As with LVM, VxVM partitions collections of disks into groups. In LVM these are known as volume groups and in VxVM they are referred to as disk groups or dynamic disk groups. For VxVM to be initialized, the default disk group rootdg must be set up and at least one disk must be assigned to it. vxinstall sets up rootdg. You cannot use VxVM until rootdg is set up, and at least one disk must remain in rootdg at all times while VxVM is running.

NOTE
Despite its name, rootdg does not necessarily include your root disk.

2   Installing VERITAS Volume Manager

VxVM is installed by default, when:

1. Upgrading from HP-UX 11.22 to HP-UX 11i v2

2. Installing HP-UX 11i v2

This chapter describes only with installing the VxVM product from the VERITAS CD. It includes the following topics:

- Mounting the CD-ROM

- Installing VxVM

- Installing VEA

- Removing Previously Licensed Packages

Mounting the CD-ROM

Before you install the VxVM or VEA, you need to mount the VxVM CD-ROM on your system as follows:

Step   **1.**   Log in as superuser.

Step   **2.**   Insert the VxVM CD-ROM into your system's CD-ROM drive.

Step   **3.**   Determine the block device file for the CD-ROM drive by entering:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.

Step   **4.**   Run the following commands to start PFS (Portable File System):

```
# nohup pfs_mountd &
```

```
# nohup pfsd &
```

Step   **5.**   Mount the CD-ROM. For example, to mount the CD-ROM at the mount point /cdrom, enter:

```
# /usr/sbin/pfs_mount -t rrip /dev/dsk/c#t#d# /cdrom
```

where /c#t#d# is the location of the drive.

NOTE   The installation process may generate one or more messages indicating that the media is not writable. You can ignore these messages.

VxVM Packages

The VERITAS CD-ROM for VxVM contains the following packages:

- VRTSvxvm and Volume Replicator software.
- VRTSvlic-VERITAS Licensing Utilities
- VRTSvmdoc-online copies of VxVM guides.

NOTE   The VERITAS Volume Manager Release Notes are not installed with any of the packages. VERITAS recommends that you copy this document to the /opt/VRTS/doc directory so that they are available for future reference.

- VRTSvrdoc-online copies of VERITAS Volume Replicator guides.
- VRTSob-VERITAS Enterprise Administrator Service (required).
- VRTSobgui-VERITAS Enterprise Administrator (required).
- VRTSvmpro-VERITAS Virtual Disk Management Services Provider (required).
- VRTSfspro-VERITAS File System Provider (required, even if you are not installing the VERITAS File System software).

The HP pre-AR0902 CD-ROM contains the following packages:

- Base-VXVM—Base Volume Manager 3.5 bundle for HP-UX.

- B9117AA—Full CVM 3.5 for HP-UX.

- B9116AA—Full Volume Manager 3.5 for HP-UX.

- B9118AA—Volume Manager 3.5 FastResync Option for HP-UX.

Installing VxVM

You can install VxVM using the swinstall utility, as described in the following sections; alternatively, the VERITAS software disc has an automated installation and licensing procedure that lets you install packages using an Installation Menu.

Installing VxVM Using the Installation Menu

The Product License and Installation Guide, included with the VERITAS software disc, provides complete information on using the Installation Menu. You should read the Product License and Installation Guide before using the Installation Menu to install VxVM.

Installing VxVM Non-Interactively with swinstall

To install VxVM for HP-UX non-interactively with swinstall, use the following command:

```
# swinstall -x autoreboot=true -s \
/cdrom/volume_manager3.5/pkgs VRTSvxvm VRTSvmdoc VRTSvlic VRTSob \ VRTSobgui VRTSvmpro
VRTSfspro
```

Your system is then rebooted.

Installing VxVM Interactively with swinstall

To install VxVM for HP-UX interactively with swinstall, do the following steps:

Step 1. Run swinstall:

```
# swinstall -s /cdrom/volume_manager3.5/pkgs
```

Step 2. Mark the software bundle you wish to install, and install them.

Your system is then rebooted.

Installing 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)

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

- VERITAS Enterprise Administrator (VRTSobgui)—the client package for HP-UX.
- VERITAS Enterprise Administrator for Windows® (win32/VRTSobgui.msi)—the client package for Windows.

NOTE   This release of VEA requires Volume Manager Release 3.5. Before you run VEA, you must install Volume Manager and run the vxinstall program to initialize Volume Manager.

The following are minimum system recommendations for the VEA client:

Table 2-1   **VEA Client Requirements**

HP-UX:	64M memory
Windows:	100MHz Pentium with 32M memory

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 See "Installing the VEA Server and Client on HP-UX" on page 29 or "Installing the VEA Client on Windows" on page 30.

Installing the VEA Server and Client on HP-UX

Install the VEA server and client on an HP-UX machine as follows:

Step 1. Log in as root.

Step 2. Install the VEA package as described in See "Mounting the CD-ROM" on page 26.

To install the VEA server and client, answer y at the appropriate prompts.

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

The VEA package for HP-UX contains the following directories:

Table 2-2   **VEA Directories**

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

Table 2-2   **VEA Directories (Continued)**

Directory	Contents
/opt/VRTS/man/	VEA manual pages

Installing the VEA Client on Windows

The VEA client runs on Windows NT, Windows XP, 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 an HP-UX 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:

Step **1.** Log in as administrator.

Step **2.** Insert the CD-ROM containing VEA.

Step **3.** Using Windows Explorer or a DOS Command window, go to the pkgs > win32 directory and execute the vrtsobgui.msi program with Windows installer.

Step **4.** Follow the instructions presented by the vrtsobgui.msi program.

Step **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.

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.
```

Step   **1.** De-install the old package.

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

```
WARNING: The VRTSvxfs package depends on the package currently being removed. Dependency
checking failed.

Do you want to continue with the removal of this package [y,n,?,q]
```

Step   **3.** Use swinstall to install the latest VRTSvlic package.

NOTE   VRTSvxfs 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.

3 Setting up the VxVM Environment

Use the `vxinstall` procedure to initialize VxVM. The `vxinstall` initialization procedure:

- Ensures that proper licenses for the VxVM software are in place, given license keys previously obtained from VERITAS customer support

- Prompts to set up the initial VxVM disk group and populate it with disks

- Starts the various configuration, notification, and relocation daemons used by VxVM

Initializing VxVM

After you install the VERITAS Volume Manager, use `vxinstall(1M)` to initialize it. `vxinstall` does the following:

- Sets up the initial VxVM disk group, rootdg, and populates it with disks.

- Starts the configuration, notification, and relocation daemons used by VxVM.

Selecting Disks for VxVM Initialization

The `vxinstall` command initializes the Volume Manager, as described below. Tasks performed through `vxinstall` include initializing disks and bringing them under Volume Manager control. Be aware of the contents of all disks on your system so that you can determine the disposition of each as `vxinstall` is run.

During the initialization process, you must decide whether you wish to place all or just some of the disks on a controller under Volume Manager control. Two types of installation are available, as follows:

- Custom Installation, to initialize certain disks and not others

- Quick Installation, to initialize all disks on a controller together

In order to complete a VxVM initialization through vxinstall, select at least one disk to bring under Volume Manager control. Alternately, use the `vxinstall -C` option to convert an existing LVM Volume Group. Even if you use the conversion option for `vxinstall`, it is recommended that you survey your attached disks.

When you run `vxinstall`, disks that you want to initialize for Volume Manager control cannot be under the control of LVM; `vxinstall` does not allow initialization of LVM disks for VxVM control.

Excluding Disks from VxVM Control

To exclude specific disks from VxVM control, add the names of those disks to the `/etc/vx/disks.exclude` file.

NOTE   The files `/etc/vx/cntrls.exclude`, `/etc/vx/disks.exclude`, and `/etc/vx/enclr.exclude` are used by the `vxinstall` and `vxdiskadm` utilities to automatically exclude controllers, disks or enclosures so that these devices are not configured as Volume Manager devices. These files do not exclude controllers, disks and enclosures from use by any other VxVM commands. See the `vxinstall(1M)` and `vxdiskadm(1M)` manual pages for more information.

LVM disks are automatically excluded from VxVM control and do not need to be added to the `/etc/vx/disks.exclude` file. No Volume Manager utility can be used to initialize an LVM-controlled disk. You must take specific actions to remove a disk from LVM control before attempting to place it under VxVM control.

Volume Manager utilities recognize file systems on raw disks (not under LVM control). Raw disks with file systems cannot be placed under Volume Manager control without administrator confirmation that the disk file system be destroyed. However, to avoid any possible confusion, enter these disks in the `disks.exclude` list.

Raw disks in use by other managing agents, such as databases, are not automatically recognized by VxVM utilities. These disks are good candidates for inclusion in the `disks.exclude` file.

The `disks.exclude` file is not created by VxVM software installation. To add disk names to this file, you must create the file. To list a disk in the file, enter only the base device file name. For example:

```
c0t0d0
c0t2d0
c1t10d0
```

To add disks that are in use by LVM volumes to the `disks.exclude` file, add all disks in the volume group to which the disk belongs. For example:

```
# for i in `vgdisplay -v uservg | egrep dsk | awk '{print $3}'`
> do
> basename $i
> done
```

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

Excluding Controllers from VxVM Control

The `/etc/vx/cntrls.exclude` file is used to exclude controllers. When the name of a disk controller is added to this file, VxVM does not initialize or control all disks on the controller.

The following is an example of the file's contents:

```
c1
c3
c8
```

Excluding Enclosures from VxVM Control

The `/etc/vx/enclr.exclude` file is used to exclude enclosures. When the name of a disk enclosure is added to this file, VxVM does not initialize or control all disks on the controller.

The following is an example of the file's contents:

```
enc2
enc3
```

Using vxinstall

This section describes how to use `vxinstall`.

To place some of your disks in another disk group, wait until after you have used `vxinstall` to set up rootdg. Use `vxdiskadm(1M)` or the VEA to create and populate other groups.

NOTE   You only need to run `vxinstall` once.

Using vxinstall with Unused Disks

NOTE   This procedure describes how to use `vxinstall` to put existing free disks, which are not in use by LVM or other data managers (such as databases or file systems), under VxVM control.

Step   **1.** Log in as root.

Step   **2.** If you want to exclude any disks, controllers, or enclosures from VxVM control, create the file `/etc/vx/disks.exclude`, the file `/etc/vx/cntrls.exclude`, or the file `/etc/vx/enclr.exclude`, respectively, and add the names of those disks, controllers, or enclosures to the appropriate file.

If you have any disks that are in use, explicitly excluding them this way allows you to run the "Quick Installation" option of `vxinstall`.

NOTE   The `/etc/vx/disks.exclude`, `/etc/vx/cntrls.exclude`, and `/etc/vx/enclr.exclude` files are used only by `vxinstall` and `vxdiskadm` to automatically exclude controllers, disks or enclosures so that these devices are not configured as Volume Manager devices. These files do not exclude controllers, disks, and enclosures from use by any other VxVM commands. VEA, does not use these files.

You may want to exclude from VxVM control:

- Raw disks that contain file systems.

- Raw disks in use by other managing agents, such as databases.

When the VEA or a VxVM utility, such as vxinstall, or vxdiskadm, brings a disk under VxVM control, it destroys any data on the disk. VxVM utilities recognize file systems on raw disks, and will ask you to confirm that the data can be destroyed. However, VxVM utilities do not recognize raw disks that are managed by other agents. It is safest to explicitly exclude any disks in use by editing the exclude file.

You must create the exclude files, `/etc/vx/disks.exclude`, `/etc/vx/enclr.exclude`, and `/etc/vx/cntrls.exclude`, if you need them; they are not created automatically. To exclude a disk, add its base device file name on a line by itself in the file. For example:

```
# cat /etc/vx/disks.exclude
c0t0d0
c0t2d0
c1t10d0
```

Adding disk controller names to `/etc/vx/cntrls.exclude` precludes VxVM initialization or control of all disks on that controller. To exclude all disks on a controller, add the controller name on a line by itself in the file. For example:

```
# cat /etc/vx/cntrls.exclude
c1
c3
c8
```

To exclude enclosures, use `/etc/vx/enclr.exclude`. When the name of a disk enclosure is added to this file, VxVM does not initialize or control all disks on the controller. For example:

```
# cat /etc/vx/enclr.exclude
encl2
```

Step 3. Run `vxinstall`:

```
# vxinstall
```

The `vxinstall` program 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 6.

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

Step 4. 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.

Enter `y` to assign a new name to the enclosures. This helps to identify their physical location. Enclosure names cannot exceed 25 characters in length. Enter the name of the enclosure that you wish to rename, and then enter its new name.

Step 5. 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
```

Step 6. If you selected `n` in the previous step, the `vxinstall` program displays a brief introduction to the installation process, as follows:

Hit RETURN to continue

Step 7. Press **Return** to continue

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

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 VERITAS Volume Manager 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.

Step 8. Select either 1 for Quick Installation or 2 for Custom Installation.

- Quick Installation uses default options and brings every disk on your system under VxVM control, except for disks under LVM control (including your root/boot disk) and disks listed in the files `/etc/vx/disks.exclude` and `/etc/vx/cntrls.exclude`. See "Quick Installation" on page 40 below for details.

- Custom Installation allows you to control whether and how each disk will be brought under VxVM control. See "Custom Installation" on page 41 for details.

Quick Installation

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

Quick Installation allows you to initialize disks in a disk array, even if valid data exist on those disks. Avoid selecting the initialize option for any disks that contain valid data.

Quick Installation brings every disk attached to your system under VERITAS Volume Manager control, unless you have chosen to exclude those disks. Disks could be chosen for exclusion either through, 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 `/etc/vx/disks.exclude`, `/etc/vx/ntrls.exclude` or `/etc/vx/enclr.exclude` files, the option "Prevent multipathing/Suppress devices from VxVM's view" in `vxinstall`, or the Custom Installation option (which allows you to leave certain disks alone).

NOTE   `vxinstall` shows default responses in parentheses; to choose the default, press **Return**. You can quit the initialization at any time by typing q at any of the `vxinstall` prompts. Disks are initialized at the end of the `vxinstall` procedure. If you quit `vxinstall` before it initializes the disks, all disks are left as they were before vxinstall was started.

Step 1. The `vxinstall` program prompts for 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.

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

CAUTION   Initializing a disk destroys all data on that disk.

Hit RETURN to continue.

Step 2. You are then prompted to indicate if you want to initialize all of the disks on the disk array.

- If you enter n, the `vxinstall` program assumes that you want to leave all the disks on the disk array as they are.

  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:

If you enter y, vxinstall automatically assigns and lists default disk names for each disk.

If you enter n, vxinstall prompts you for a disk name for each disk on the disk array individually. For each disk, respond with the desired disk name and press Return.

Step **3.** 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.

- If you enter n, the vxinstall program prompts you for a disk name for each disk on the disk array individually.

Step **4.** For each disk, respond with the desired disk name and press **Return**.

Step **5.** When all of the disks on the current disk array have been named, press **Return** and repeat Step 3 and Step 4 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) on each controller.

If you enter y, vxinstall initializes (destroying any existing data) all disks listed as "New Disk."

If you enter n, vxinstall prompts you for the name of a disk to be removed from the list and excluded from VxVM control. To remove a disk from the list, enter the name of the disk to be removed from the list and press **Return**. Repeat for each disk to be removed and then press **Return** once more. vxinstall displays an updated summary. If you are satisfied with the summary, enter y.

Depending on your input during the vxinstall program, you may be requested to shut down and reboot your machine. If this happens, shut down and reboot your machine immediately or at the earliest convenient opportunity.

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, successful completion of the installation is shown as follows:

```
Installation is complete.
```

Custom Installation

Custom installation allows you to control which disks are added to VxVM control and how they are added. This is important if you do not want all your disks under VxVM control. Custom Installation also gives you the option of initializing all disks on a controller or initializing just some disks on a controller. The following session example illustrates a Custom Installation.

Note the following disk initialization issues with Custom Installation:

- Custom Installation does not allow you to initialize disks on an controller if LVM logical volumes are found on those disks. Avoid selecting the initialize option for any disks containing valid data.

- Custom Installation allows you to initialize disks in a disk array, even if valid data exists on those disks. Avoid selecting the initialize option for any disks that contain valid data.

NOTE	vxinstall shows default responses in parentheses; to choose the default, press Return. You can quit the initialization at any time by typing q at any of the vxinstall prompts. Disks are initialized at the end of the vxinstall procedure. If you quit vxinstall before it initializes the disks, all disks are left as they were before vxinstall was started.

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

Step   **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**.

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

Step   **2.** 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 controller and generating a list of its disks.

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

Step   **3.** Press **Return** to continue.

Step   **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 new disks.(discards data on disks!)
2  Install one disk at a time.
3  Leave these disks alone.

?  Display help about menu
?? Display help about the menuing system
q  Exit from menusSelect an operation to perform:
```

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

CAUTION	Initializing a disk destroys all data on the disk.

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

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

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

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

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

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

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

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

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

- To install one disk at a time, select 2.

  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.

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

  Options 1 and 2 allow you initialize the disk or to omit it from installation.

  Repeat this step for all disks on the current disk array.

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

- To leave all disks in the disk array unaltered, select 3.

  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.

Step 5. Repeat Step 4 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) on each controller.

If you enter y, the vxinstall program proceeds to initialize (destroying any existing data) all disks listed with New Disk.

If you enter n, the `vxinstall` program prompts you for the name of a disk to be removed from the list and excluded from Volume Manager control, as follows:

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

[<name>,q,?]
```

Step 6. 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.

Step 7. 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.

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.
```

Preventing Multipathing/Suppressing Devices from VxVM's View

This section describes how to exclude a device which is under Volume Manager or Dynamic Multipathing control. Follow these steps to disable multipathing.

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

The following message displays:

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

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

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

Step **2.** Enter `y`.

Step **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:
```

The following list briefly summarizes each operation. Additional explanatory information about each operation is displayed when you enter a choice, and you are always given the opportunity to exit from the menu before executing the choice.

- Option 1 prevents VxVM from using all paths through a specified controller. The specified controller is disabled.

- Option 2 prevents VxVM from using one or more specified paths. The specified path is excluded from VxVM control

- Option 3 prevents VxVM from using one or more disks specified by a VID:PID combination, where VID stands for a Vendor ID, and PID stands for Product ID. (The command `/etc/vx/diag.d/vxdmping` can be used to obtain these IDs. This command displays the standard SCSI inquiry data returned from the disk controller.)

- For example, the VID:PID combination for an ACME Z1 disk array would be specified as ACME:Z1.

- Option 4 prevents VxVM from defining a set of paths, or pathgroup, for disks that are not multipathed by VxVM. (A pathgroup specifies a set of paths that VxVM should recognize as a single path. Only one path is made visible from a pathgroup.)

- Option 5 excludes devices by preventing VxVM from applying dynamic multipathing (DMP) to all disks on a controller. The controller you name is excluded from DMP control.

- Option 6 excludes a specified path from dynamic multipathing. For the path you specify, the corresponding disks are claimed in the OTHER_DISKS category and are not multipathed.

- Option 7 excludes devices by preventing VxVM from applying dynamic multipathing to disks specified by a VID:PID combination. All disks returning a VID:PID combination are claimed in the OTHER_DISKS category and are not multipathed.

- Option 8 lists devices that have either been suppressed from the view of VxVM, or that have been excluded from dynamic multipathing.

If you select any one of options 1 through 7, you must reboot the system for Device Exclusion to take effect, and then re-run `vxinstall`.

4   Upgrading the VERITAS Volume Manager

The following topics are covered in this chapter:

Upgrading to the Latest VxVM Version

We recommend that you upgrade to VxVM 3.5 on any and all of the VxVM products you have previously installed. Version VxVM 3.5 is more robust than earlier versions of the software.

Upgrading to the latest VxVM software is automated when installing HP-UX 11i v2.

Do not run vxinstall(1M) after upgrading.

License keys from older versions of B9116AA and B9118AA continue to work when you upgrade to VxVM Version 3.5.

Upgrading Disk Group Versions

All disk groups have a version number associated with them. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks only work on disk groups with the current disk group version, so you need to upgrade existing disk groups before you can perform these tasks. The following table summarizes the disk group versions that correspond to each VxVM release:

Table 4-1   **Disk Group Versions**

Volume Manager Release	Disk Group Version	Supported DIsk Group Versions
3.1	70	20-70
3.2, 3.5	90	20-90

You can get a disk group version listing by specifying a disk group with this command:

```
# vxdg list diskgroup
```

You can also determine the disk group version by using the vxprint(1M) command with the -l format option.

To upgrade a disk group, use the following command:

```
# vxdg [-T version] upgrade diskgroup
```

If a disk group version is not specified using the -T option, this command upgrades the disk group to the highest version supported by the VxVM version on your system.

For shared disk groups, the latest disk group version is only supported by the latest cluster protocol version. See "Upgrading Cluster Functionality" in the Volume Manager Administrator's Guide for more information on changing the cluster protocol version.

Supported Upgrade Scenarios

If you are already running a previous version of VxVM, you can upgrade to VxVM 3.5. In some cases, you do not have to do a complete removal of the old VxVM packages.

With the release of VxVM 3.5, you also have the choice of having your root disk placed under VxVM control. With this in mind, the following upgrade procedures are broken down into:

- Procedures needed to upgrade to the VxVM 3.5 release from the VxVM and HP-UX release that you are currently running.

- Procedures needed to upgrade to the VxVM 3.5 release from the VxVM and HP-UX release that you are currently running, plus putting the root disk under VxVM control.

Upgrading to a VxVM Root Disk

With the advent of the HP-UX 11i AR0902 release, it is now possible to select VxVM as a choice for your root disk when performing a new installation using Ignite-UX. However, if you are a current HP-UX user, and you are performing an upgrade from a prior release of HP-UX, or if you are a current HP-UX user also running a prior release of VxVM, you cannot invoke Ignite-UX to convert your system to be VxVM rootable. You can use the following procedure to achieve VxVM rootability by cloning your LVM root disk using the vxcp_lvmroot (1M) command.

Step 1. Upgrade your system to HP-UX 11i v2 if your current version of HP-UX and/or VxVM is not already at this level.

Step 2. Select the disk to be used as your new VxVM root disk. It is recommended, that this disk is internal to the main computer cabinet. If this is currently an LVM disk, then it must be removed from LVM control (see "Preparing to Install VxVM on a System with LVM Volume Groups" on page 11). If you have two spare disks, and you want to set up a mirrored root disk, make sure that the second disk is also free from LVM control.

Step 3. While booted on the newly upgraded LVM root disk, invoke the vxcp_lvmroot command to clone the LVM root disk to the disk(s) you have designated to be the new VxVM root disk(s). In the following examples, c1t0d0 and c2t0d0 are used for the target VxVM root disk and the (optional) root mirror.

```
# /etc/vx/bin/vxcp_lvmroot -v c1t0d0
```

or; for a mirrored root:

```
# /etc/vx/bin/vxcp_lvmroot -v -m c2t0d0 c1t0d0
```

The-v option is highly recommended. The cloning of the root disk is a lengthy operation, and the -v option gives a time-stamped progress indication as each volume is copied, and other major events.

NOTE   The vxcp_lvmroot command checks the OS revision number, and will not allow upgrade to proceed unless the OS revision is at the required level.

Step 4. You can use the setboot (1M) command to save the hardware path of the new VxVM root disk in the system NVRAM. The disk hardware paths can be found using the command:

```
# ioscan -kfnC disk
```

Step 5. Reboot from your new VxVM root disk. If you created a mirrored root disk, then there is nothing more to do. The LVM root disk can now safely co-exist with your VxVM root disk, giving you an emergency backup boot target.

Step 6. If you wish, you can convert your original LVM root disk into a mirror of your VxVM root disk using the following commands:

```
# /etc/vx/bin/vxdestroy_lvmroot -v c2t0d0
# /etc/vx/bin/vxrootmir -v c2t0d0
```

Once this operation is complete, you will be running on a completely mirrored VxVM root disk.

Step 7. If you later choose to return to an LVM root, you can use the vxres_lvmroot (1M) command to restore the LVM root disk.

Upgrade Scenarios

Three upgrade scenarios are considered; these include upgrading with, and without, VxVM rootability:

Step 1. Upgrade from a prior version of HP-UX 11.22 LVM rootable system, with VxVM configured, to HP-UX 11i v2 LVM rootable system with VxVM 3.5.

Step 2. Upgrade from a prior version of HP-UX 11.22 LVM rootable system, with VxVM configured, to HP-UX 11i v2 VxVM rootable system with VxVM 3.5.

Step 3. Upgrade from HP-UX 11.22 LVM rootable system with VxVM 3.1 to HP-UX 11i v2 LVM rootable system with VxVM 3.5.

Table 4-2   **Rootability Upgrade Paths**

	If Your Current Configuration is ...	**... And You Want to Upgrade to ...**	**... Then You Need to ...**	**Refer to**
Scenario 1	HP-UX 11.22 LVM rootable, with VxVM 3.1 configured	HP-UX 11i v2 LVM rootable with VxVM 3.5	Upgrade using the HP-UX 11i v2 CD. VxVM 3.5 is installed by default	"VxVM Not Controlling the Root Disk" on page 3
Scenario 2	HP-UX 11.22 LVM rootable with VxVM 3.1 configured	HP-UX 11i v2 VxVM rootable with VxVM 3.5	Upgrade using the HP-UX 11i v2 CD.	"Installing VERITAS Volume Manager" on page 25
Scenario 3	An HP-UX 11.22 VxVM rootable	HP-UX 11i v2 VxVM rootable with VxVM 3.5	Upgrade using the HP-UX 11i v2 CD	

5 Getting Started with VxVM

The following topics are covered in this chapter:

- "Starting the VEA Service" on page 54
- "Starting the VEA Client" on page 55
- "Using Hot-Relocation" on page 56
- "Enabling Cluster Support in VxVM" on page 57
- "Configuring Shared Disks" on page 58
- "System Setup Guidelines" on page 59
- "Protecting Your System and Data" on page 63
- "Working With VxVM 3.5 Root Disks" on page 64

Starting the VEA Service

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, Volume Manager and the VEA server must be installed on the machine to be administered. The Volume Manager vxconfigd daemon and the VEA server must be running on the machine to be administered.

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

On an HP-UX System

To administer the local HP-UX system, use the following command:

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

To administer a remote UNIX system, use the following command:

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

Confirming that VxVM Processes are Running

After VxVM has been successfully installed or upgraded, use the following command to confirm that key VxVM processes (vxconfigd, vxnotify, and vxrelocd) are running:

```
# ps -ef | egrep vx
```

This command displays output similar to the following:

```
root   18    0    0    Apr 19    ?      412:53  vxiod
root   20    0    0    Apr 19    ?      0:00    vxfsd
root   324  316  0    Apr 19    ?      0:00    vxnotify -f -w 15
root   316  0    0    Apr 19    ?      0:00   vxrelocd - /etc/vx/bin/vxrelocd
root   253  1    0    Apr 19    ?      0:00    vxconfigd
root   4033 3997 1    14:40:50  pts/0  0:00    egrep vx
```

NOTE   If you have disabled Hot-Relocation, vxrelocd should not be shown as running.

Using 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:

Step **1.** Leave the Volume Manager 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.

Step **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.

Disabling Hot-Relocation

The hot-relocation feature is enabled by default. The associated daemon, vxrelocd, is automatically started during system startup. If you decide to disable hot-relocation, you need to prevent vxrelocd from running after you load the Volume Manager software.

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

```
# nohup vxrelocd root &
```

Refer to the Volume Manager Administrator's Guide for more information on hot-relocation.

Enabling Cluster Support in VxVM

NOTE   Cluster management software such as MC/ServiceGuard must be installed on the nodes of your cluster system.

VxVM 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.

Enable the cluster functionality in VxVM as follows:

Step   **1.** Obtain a license for the Volume Manager cluster feature from your Customer Support channel.

Step   **2.** Install the software packages onto each system (node) to be included in the cluster as described in Chapter 2, "Installing VERITAS Volume Manager," on page 25.

Step   **3.** Initialize VxVM using the `vxinstall` procedure.

Step   **4.** Set up and start the VEA as described in "Starting the VEA Service" on page 54.

Step   **5.** Configure shared disks as described in "Configuring Shared Disks" on page 58.

Configuring Shared Disks

When installing VxVM for the first time or adding disks to an existing cluster, you must configure new shared disks. When upgrading VxVM, you must verify that your shared disks still exist.

Shared disks are configured from one node only. Since the VxVM software is unable to determine whether or not a disk is shared, you must specify which disks are shared.

When configuring disks, ensure that no one else is accessing shared disks from another node. Start the cluster on the node where the configuration is performed in order to prevent disk accesses from other nodes. Quorum control reserves the disks for the single node.

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 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

Step **1.** Place disks under VxVM control.

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

Step **3.** Create volumes.

Step **4.** Put file systems on volumes.

- Options

    — 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.

Dirty Region Logging (DRL) Guidelines

Dirty Region Logging (DRL) can speed up recovery of mirrored volumes following a system crash. When DRL is enabled, Volume Manager keeps track of the regions within a volume that have changed as a result of writes to a plex. Volume Manager 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

NOTE   The VxVM Full license is required to use this feature.

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

NOTE   The VxVM Full license is required to use this feature.

Refer to the following guidelines when using RAID-5. For more information, see the 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

VxVM can protect your system from loss of data due to disk failure.

The following are suggestions for protecting your data:

- 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 the *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 Dirty Region Logging (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   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.

Working With VxVM 3.5 Root Disks

This section describes:

- "Adding Dump Volumes Using Volume Manager Disks" on page 64
- "Changing the Boot Disk to be the New Volume Manager Root Disk" on page 64
- "Removing a Mirrored Volume Manager Root Disk" on page 64

Adding Dump Volumes Using Volume Manager Disks

Volume Manager volumes can be used for additional dump volumes in configurations with LVM or VxVM root disks.

Step 1. Remove a disk from LVM control. Execute the following:

```
# pvcreate -f /dev/rdsk/c#t#d#
# pvremove /dev/rdsk/c#t#d#
```

Step 2. Put a disk under Volume Manager control and place in rootdg:

Use vxdiskadm option 1, or VEA.

Step 3. Add a volume to this disk:

Use vxassist or VEA.

Step 4. Add the new volume as a target for dumps:

First make sure you have the PHKL_27096 patch, then use crashconf as you would under LVM.

For example;

```
# crashconf /dev/vx/dsk/rootdg/your_dump_volume
```

Changing the Boot Disk to be the New Volume Manager Root Disk

Step 1. Get the current boot disk:

```
# cat /stand/bootconf
```

Step 2. Extract resultant device path:

```
# ioscan -k node_path_from_bootconf
```

Step 3. Change primary boot path back to current boot device:

```
# setboot -p device_path_from_above
```

(Use -a if you want to change the alternate bootpath.)

Removing a Mirrored Volume Manager Root Disk

Use the following commands to remove the mirrored Volume Manager Root Disk:

```
# daname=<c#t#d#>             /* of mirrored root disk */
# dmname=`vxprint -g rootdg -F %dmname -e sd_da_name~/$daname/|head -1`
# plxnames=`vxprint -g dg -F "%assoc" -e sd_dm_name==\"$dmname\"`
# vxplex -o rm dis $plxnames
# vxdg -g rootdg rmdisk $dmname
```

Adding Swap Space Using VxVM Volumes

The HP System Administration Manager (SAM) currently does not have the capability to add swap space using VxVM volumes. You need to use the VxVM CLI; for example:

Step **1.** Run the command:

```
# /etc/vx/bin/vxdisksetup -i device
```

Step **2.** If this command comes back with:

vxvm:vxdisk: ERROR: Device device is already in use by LVM.

then, go to Step 3; otherwise, go to Step 4.

Step **3.** Run the following commands (it is necessary to run pvcreate first, since /etc/lvmtab may not exist):

```
# pvcreate -f /dev/rdsk/ device
# pvremove /dev/rdsk/ device
# vxdctl enable
# /etc/vx/bin/vxdisksetup -i device
```

Step **4.** Run the following commands:

```
# vxdg adddisk swapdisk=device
# vxprint swapdisk

Disk group: rootdg
TY NAME          ASSOC KSTATE LENGTH PLOFFS STATE TUTIL0 PUTIL0
dm swapdisk   c2t6d0 - 4191378  - - - -
```

Step **5.** Run the commands:

```
# vxassist make swapvol2 4191378 swapdisk
# echo "/dev/vx/dsk/rootdg/swapvol2 ... swap pri=1 0 0" >>/etc/fstab
```

Step **6.** Use SAM to:

- Increase the maxswapchunks tunable as required by swapon.

- Build a the kernel.

- Reboot your system.

6   Uninstalling the VERITAS Volume Manager

This chapter conves the following topics:

- "Preparing to Remove VxVM" on page 68

    — "Preparatory Steps" on page 68

    — "Moving Data Out of VxVM Control" on page 69

- "Shutting Down and Removing VxVM" on page 70

    — "Shutting Down VxVM" on page 71

    — "Removing the VxVM Package" on page 71

Preparing to Remove VxVM

Before removing the VxVM software with swremove, you must backup your data, move data out of VxVM control, and stop the VxVM.

CAUTION	If you try to swremove the VxVM software from your system without doing the following preparatory steps, you will likely lose data and leave your system in an uncertain state.

Preparatory Steps

NOTE	If VxVM rootability has been installed using either an Ignite_UX install, or by using vxcp_lvmroot, VxVM cannot be uninstalled until the root disk is under LVM control. If you used vxcp_lvmroot to become VxVM rootable, then you can use the vxres_lvmroot command to restore the root disk to LVM control.

Consider which of the following steps apply to your system, and complete them.

NOTE	Because reorganizing your system in preparation for removing VxVM software is complicated, it is important to do a full system backup before beginning.

- All file systems and other applications that use volume devices must be modified to use the underlying disks or logical volumes, and so on.

- Any file systems that have been created since the VxVM was installed must be made accessible through a disk or LVM logical volume.

- All data from any volumes that were created from multiple regions of storage, including striped or spanned volumes, must be moved onto a single disk or appropriate LVM logical volume.

- Any mirror copies of file systems and other databases should be removed so that only one copy remains, to free up as much space as possible.

- If any volume contains more than one plex (mirror), all but one plex should be removed.

  To display a list of all volumes, use the command:

```
# vxprint -Ath
```

  To remove a named plex, use the command:

```
# vxplex -o rm disk plex
```

  If any remaining plex has more than one subdisk, consolidate those subdisks into a single subdisk. A simple way to do this is to use the commands:

```
# vxassist mirror volume layout=contig
# vxplex -o rm dis plex_name
```

NOTE	This operation does not work if you do not have sufficient space on another disk.

- Modify /etc/fstab to mount all file systems using a partition or logical volume instead of a VxVM volume. Update /etc/fstab for volumes that are not on rootdg.

Moving Data Out of VxVM Control

Before removing the VxVM, all the VxVM volumes must be moved to LVM disks. This can be done in one of the following ways:

- Incrementally create LVM volumes and move the data from each VxVM volume to an LVM volume one at a time. This procedure is described in "Preparing to Remove VxVM" on page 68.

- Back up the data from each VxVM volume, then remove the VxVM volumes and disk groups, use vxdiskunsetup(1M) to decommission the disks, put the disks under LVM control, create LVM logical volumes of the same size as the VxVM volumes, and restore the data from the backup.

- If you have raw data that is not on a file system, you can do a full backup of the raw data, then remove VxVM, then recover the backup.

Shutting Down and Removing VxVM

To uninstall VxVM, you must move all VxVM volumes to LVM disks. This can be done using one of the following three methods:

- Back up the system to tape or other media and recover the system from this.

- Back up each file system individually, create new file systems on LVM logical volumes, and recover all volumes.

- Move volumes incrementally ("evacuate") onto logical volumes. Evacuation moves subdisks from the source disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to LVM volumes. Use the following steps to do this:

Step 1. Evacuate as many disks as possible in one of the following ways:

- use the "Remove a disk" option in vxdiskadm and follow through the prompted steps

- use VxVM

- use the vxevac script

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

```
# vxdg rmdisk diskname
```

```
# vxdisk rm devname
```

Step 3. Decide which volume to move first. If the volume to be moved is mounted, unmount it. If the volume is being used as a raw partition for database applications, ensure that the application is not updating the volume and that data on the volume has been synchronized.

Step 4. On the free disk space, create an LVM logical volume that is the same size as the VxVM volume. If there is not enough free space for the logical volume, add a new disk to the system for the first volume to be removed. For subsequent volumes, you can use the free space generated by the removal of the first volume.

Step 5. Copy the data on the volume onto the newly created LVM logical volume using the following command:

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

where vgvol is a newly created LVM volume group and lhome is a new logical volume.

Step 6. Replace the entry (if any) for the volume in /etc/fstab with an entry for the newly created LVM logical volume.

Step 7. Mount the logical volume if the corresponding volume was previously mounted.

Step 8. Remove the volume from Volume Manager control using the following command:

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

Step 9. Remove any free disks (without subdisk definitions) by removing volumes from Volume Manager control. To check if subdisks remain on a disk, use the following command:

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

If the output is other than 0, some subdisks remain on this disk and must be removed. If the output is 0, remove the disk from Volume Manager control using the following commands:

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

Step 10. Copy the data in the next volume to be removed to the newly created free space.

Step 11. Reboot the system after all volumes have been converted successfully. Verify that no open volumes remain after the system reboot using the following command:

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

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

Shutting Down VxVM

Use the following commands to shut down the VxVM, use the following commands:

```
# vxdctl stop
```

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

Removing the VxVM Package

If you installed VxVM from an HP package, you can remove VxVM using the following command:

```
# swremove -x autoreboot=true VRTSvxvm VRTSvmdoc
```

The host will now reboot.

If you installed VRST packages, remove them using commands of the form:

```
# swremove -x autoreboot=true VRTSvxvm
```

(This will remove the VRSTvxvm package.)

The host will now reboot.

Errors that occur during the software removal process are due to the following:

- the system contains open volumes

- the root disk is under VxVM control

For more information, select the Logfile button in the user interface to check the log file /var/adm/sw/swagent.log .