Veritas 5.1 SP1 Installation Guide
HP-UX 11i v3

HP Part Number: 5900-1514
Published: September 2011
Edition: 2.0

Contents

About this Document

This document provides information on Veritas 5.1 Service Pack1 (SP1) suite of products on systems running HP-UX 11i v3. Veritas 5.1 SP1 suite of products include Base-VxFS, Base-VxVM, OnlineJFS, Full VxVM, and Cluster Volume Manager (CVM). This document also includes a product overview, system requirements, installation, basic configuration, and removal steps for Veritas 5.1 SP1 suite of products on HP-UX 11i v3.

Intended Audience

This document is intended for system administrators responsible for installing and configuring HP-UX systems with the Veritas suite of products. Readers are expected to have knowledge of the following:

- HP-UX operating system concepts
- System administration concepts
- Veritas Volume Manager concepts
- Veritas File System concepts

Document Organization

Table 1 Document Organization

Chapter Title	Description
Chapter 1: Introduction	Describes the Veritas 5.1 SP1 suite of products for systems running HP-UX 11i v3.
Chapter 2: System Requirements	Describes the OS version, software depot contents, license bundles, disk space requirements, and supported upgrade paths for the Veritas 5.1 SP1 suite of products.
Chapter 3: Installing the Veritas 5.1 SP1 Products	Describes how to install the Veritas 5.1 SP1 suite of products on systems running HP-UX 11i v3.
Chapter 4: Setting up the Veritas 5.1 SP1 Products	Describes how to set up the Veritas 5.1 SP1 suite of products on HP-UX 11i v3.
Chapter 5: Upgrading from Previous Versions of VxFS to VxFS 5.1 SP1	Describes how to upgrade the VxFS disk layout version.
Chapter 6: Upgrading from Previous Versions of VxVM to VxVM 5.1 SP1	Describes how to upgrade the VxVM disk group version.
"Post Upgrade Tasks" (page 64)	Discusses the post upgrade tasks for VxVM.
Chapter 8: Removing Veritas 5.1 SP1 Products	Describes how to remove the Veritas 5.1 SP1 suite of products from an HP-UX 11i v3 system.

Typographic Conventions

This document uses the following typographic conventions:

`monospace`	Computer output, files, directories, software elements such as command options, function names, and parameters. Read tunables from the `/etc/vx/tunefstab` file.
italic	New terms, book titles, emphasis, variables replaced with a name or value See "*Veritas 5.1 SP1 Installation Guide*" for more information.
`%`	C shell prompt
`$`	Bourne/Korn shell prompt
`#`	Superuser prompt (all shells)

\	Continued input on the following line; you do not type this character
`[ ]`	In command synopsis, brackets indicates an optional argument. `ls [ -a ]`
|	In command synopsis, a vertical bar separates mutually exclusive arguments. `mount [ suid | nosuid ]`
blue text	An active hypertext link In PDF and HTML files, click on the links to move to the specified location.

Related Information

Additional information on the Veritas suite of products is available at:

HP Business Support Center.

This website contains the following documents:

- *Veritas File System 5.1 SP1 Release Notes*
- *Veritas File System 5.1 SP1 Administrator's Guide*
- *Veritas Volume Manager 5.1 SP1 Release Notes*
- *Veritas Volume Manager 5.1 SP1 Administrator's Guide*
- *Veritas Volume Manager 5.1 SP1 Troubleshooting Guide*
- *Veritas Enterprise Administrator User's Guide*
- *Veritas Storage Foundation and High Availability Solutions 5.1 SP1 Getting Started Guide*
- *Veritas Storage Foundation 5.1 SP1 Advanced Features Administrator's Guide*
- *Veritas Storage Foundation 5.1 SP1: Storage and Availability Management for Oracle Databases*
- *Veritas Storage Foundation 5.1 SP1 Cluster File System Release Notes*
- *Veritas Storage Foundation 5.1 SP1 Cluster File System Installation Guide*
- *Veritas Storage Foundation 5.1 SP1 for Oracle RAC Release Notes*
- *Veritas Storage Foundation 5.1 SP1 for Oracle RAC Administrator's Guide*

To locate these documents, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Technical Support

For license information, contact:

- Software License Manager:
  http://licensing.hp.com/welcome.slm
- HP Licensing Services:
  http://licensing.hp.com/licenseAdmins.slm
  - (Europe)
    Phone: +353.(0)91.75.40.06
    Email: **codeword_europe@hp.com**
  - (U.S. and Canada)
    Phone: +1 650.960.5111

Email: **hplicense.na@hp.com**

- (Asia Pacific)
  Phone: 0120.42.1231 or 0426-48-9310 (Inside Japan)
  +81.426.48.9312 (Outside Japan)
  Email: **sw_codeword@hp.com**

- For latest information on the available patches, see:
  http://itrc.hp.com
- For technical support, see:
  http://welcome.hp.com/country/us/en/support.html

HP Welcomes Your Comments

HP welcomes your comments concerning this document. HP is committed to providing documentation that meets your needs.

Please send comments to: docsfeedback@hp.com

Please include document title, manufacturing part number, and any comment, error found, or suggestion for improvement you have concerning this document. Also, please include what we did right so we can incorporate it into other documents.

1 Introduction

This chapter introduces the Veritas 5.1 SP1 suite of products. It also describes the features of each product that is included within the Veritas 5.1 SP1 suite of products.

This chapter addresses the following topics:

Overview

Veritas 5.1 SP1 suite of products include Base-VxFS, Base-VxVM, OnlineJFS, VxVM-Full, and CVM. Veritas Volume Manager (VxVM) is a storage management subsystem that enables you to manage physical disks as logical devices called volumes. A volume is a logical device that appears to a data management system as a physical disk. Veritas File System (VxFS) is an extent based, intent logging file system that is designed for use in UNIX environments, which require high performance and availability and deal with large volumes of data.

The Cluster Volume Manager (CVM) allows up to 32 nodes in a cluster to simultaneously access and manage a set of disks under VxVM control (VM disks). The same logical view of disk configuration and any changes to this is available on all the nodes. For more information on CVM, see the *Managing Serviceguard* manual, or the *Veritas Volume Manager Administrator's Guide*.

The HP Serviceguard Storage Management Suite integrates HP Serviceguard with Symantec's Veritas Storage Foundation. This combination provides powerful database and storage management capabilities while maintaining the mission-critical reliability that HP Serviceguard customers have come to expect. For more information about HP Serviceguard Storage Management Suite, see *HP Serviceguard Storage Management Suite Version A.04.00 Release Notes*.

Volume Managers Supported on HP-UX 11i v3

HP-UX 11i v3 supports the following volume managers:

- **HP Logical Volume Manager (HP LVM)**

  The HP LVM is a disk management subsystem that enables you to allocate disk space according to the specific or projected size of your file system or raw data. For more information on HP LVM, see *HP-UX System Administrator's Guide: Logical Volume Management* on HP Business Support Center.

- **Veritas Volume Manager (VxVM)**

  VxVM is a storage management subsystem that enables you to manage physical disks as logical devices called volumes. A volume is a virtual device that appears to a data management system as a physical disk.

Veritas Volume Manager (VxVM)

Introduction

VxVM is a storage management subsystem that removes the physical limitations of disk storage so that you can configure, share, manage, and optimize storage I/O performance online without interrupting data availability. VxVM also provides easy-to-use, online storage management tools to reduce planned and unplanned system downtime, and online disk storage management for

computing environments and Storage Area Network (SAN) environments. Through RAID support, VxVM protects against disk and hardware failure. Additionally, VxVM provides features that offer fault tolerance and fast recovery from disk failure.

VxVM overcomes physical restrictions imposed by hardware disk devices, by providing a logical volume management layer. This enables volumes to span multiple disks. VxVM also dynamically configures disk storage while the system is active.

VxVM Features

Veritas Volume Manager supports the following features:

- **Veritas Enterprise Administrator (VEA)**

  A Java™-based graphical user interface for administering VxVM.

- **Concatenation**

  Concatenation maps data in a linear manner onto one or more subdisks in a plex.

- **Striping**

  Striping maps data, so that data is interleaved among two or more physical disks.

- **Mirroring**

  Mirroring uses multiple mirrors to duplicate information contained in a volume.

- **Mirrored Stripes**

  VxVM supports a combination of mirroring and striping.

- **Striped Mirrors**

  VxVM supports a combination of striping and mirroring.

- **RAID-5**

  RAID-5 provides data redundancy using parity.

- **Online Resizing of Volumes**

  You can dynamically resize VxVM volumes while the data remains available to the user.

- **Hot-relocation**

  The hot-relocation feature in VxVM automatically detects disk failures, and notifies the system administrators of the failure, by email. Hot-relocation also attempts to use spare disks and frees disk space to restore redundancy and to preserve access to mirrored and RAID-5 volumes.

- **Volume Resynchronization**

  Volume resynchronization ensures that all copies of the data match, when mirroring redundant copies of data.

- **Online Relayout**

  Online relayout enables you to convert between storage layouts in VxVM, with uninterrupted data access.

- **Volume Snapshot**

  Volume Snapshots are point in time images of VxVM volumes.

VxVM 5.1 SP1 does not support snapshots of RAID 5 volumes.

- **Dirty Region Logging**

  Dirty Region Logging (DRL) keeps track of the regions that have been changed because I/O writes to a mirrored volume. The DRL uses this information to recover only those portions of the volume that need to be recovered, thereby speeding up recovery after a system crash.

- **SmartMove™ Feature**

  SmartMove reduces the time and I/O required to attach or reattach a plex to an existing VxVM volume, in the specific case where a VxVM volume has a VxFS file system mounted on it. The SmartMove feature uses the VxFS information to detect free extents and avoids copying them.

- **Enhancements to the Dynamic Multipathing Feature**

  This release provides a number of enhancements to the DMP features of VxVM. These enhancements simplify administration and improve display of detailed information about the connected storage. Following are the enhancements to the DMP feature:

  - Dynamic multipathing attributes are now persistent
  - Improved dynamic multipathing device naming
  - Default behavior modified for I/O throttling
  - Specifying a minimum number of active paths
  - Enhanced listing of subpath
  - Enhanced I/O statistics
  - Making DMP restore options persistent
  - New log file location for DMP events
  - Extended device attributes displayed in the `vxdisk list` command
  - Displaying the `use_all_paths` attribute for an enclosure
  - Viewing information about the ASLs installed on the system
  - Displaying the number of LUNs in an enclosure
  - Displaying the LUN serial number
  - Displaying HBA details
  - New `exclude` and `include` Options for the `vxdmpadm` command
  - New command for reporting DMP node information
  - Setting attributes for all enclosures
  - Support for ALUA JBOD devices

- **VxVM Powerfail Timeout (PFTO) feature disabled in the HP-UX Native Multipathing Devices**

  By default, the use of PFTO is now disabled in the HP-UX native multipathing devices. As a result, the native multipathing disk I/O can consume additional service time to complete an I/O successfully. In case of DMP devices, the use of PFTO is enabled by default.

  For information on recommendations to maintain high levels of I/O robustness, refer to "I/O Robustness Recommendations" (page 62).

- **Support for LVM version 2 Volume Groups**

  The LVM version 2 volume groups are now partially supported. VxVM now identifies and protects the LVM version 2 volume groups. However, the LVM version 2 volume groups cannot be initialized or converted.

- **Distributed Volume Recovery**

  In a Cluster Volume Manager (CVM) cluster, upon a node failure, the mirror recovery is initiated by the CVM master. Prior to this release, the CVM master performed all the recovery I/O tasks. Starting from this release, the CVM master can distribute recovery tasks to other nodes in the cluster. Distributing the recovery tasks is desirable in some situations so that the CVM master can avoid an I/O or CPU bottleneck.

- **Campus Cluster enhancements**

  The campus cluster feature provides the capability of mirroring volumes across sites, with hosts connected to storage at all sites through a Fibre Channel network. In this release, the following enhancements have been made to the campus cluster feature:

  - Site Tagging of disks or Enclosures
  - Automatic Site Tagging
  - Site Renaming

- **Estimated Required Time Displayed During Volume Conversion**

  During a volume conversion operation, before the conversion is committed, the `vxvmconvert` command displays the estimated time required.

- **The `vxsited` daemon Renamed to `vxattachd`**

  The `vxsited` daemon is renamed as the `vxattachd` daemon. The `vxattachd` daemon now also handles automatic reattachment and resynchronization for plexes.

- **Automatic Plex Attachment**

  When a mirror plex encounters irrecoverable errors, VxVM detaches the plex from the mirrored volume. By default, VxVM automatically reattaches the affected mirror plexes when the underlying failed disk or LUN becomes visible.

- **Persisted Attributes**

  The `vxassist` command allows you to define a set of named volume allocation rules, which can be referenced in volume allocation requests. The `vxassist` command also allows you to record certain volume allocation attributes for a volume. These attributes are called persisted attributes. You can record the persisted attributes and use them in later allocation operations on the volume, such as increasing the volume.

- **Automatic recovery of volumes during disk group import**

  VxVM allows automatic recovery of volumes during disk group import. After a disk group is imported, disabled volumes can be enabled and started by default.

- **Cross-platform data sharing support for disks greater than 1 TB**

  In releases prior to VxVM 5.1 SP1, the `cdsdisk` format was supported only on disks up to 1 TB in size. Therefore, cross-platform disk sharing (CDS) was limited to disks of size up to 1 TB. VxVM 5.1 SP1 removes this restriction. It introduces CDS support for disks of size greater than 1 TB as well.

NOTE: The disk group version must be at least 160 to create and use the `cdsdisk` format on disks of size greater than 1 TB.

IMPORTANT: VxVM uses the Global Partition Table (GPT) format to initialize disks of size greater than 1TB in the `cdsdisk` format. HP Logical Volume Manager (LVM) and the `diskowner` command do not recognize disks formatted with the GPT layout. So, LVM and the `diskowner` command do not recognize disks of size greater than 1 TB.

For more information, refer to the "Known Problems and Workarounds" section of the *Veritas Volume Manager 5.1 SP1 Release Notes*.

- **Default format for auto-configured disks has changed**

  VxVM will initialize all auto-configured disks with the `cdsdisk` format, by default. To change the default format, use the `vxdiskadm` command to update the `/etc/default/vxdisk` file.

- **Default naming scheme for devices is `Enclosure Based Naming Scheme(ebn)`**

  Starting with the VxVM 5.1 SP1 release, the default naming scheme for devices has changed to the `Enclosure Based Naming Scheme(ebn)`. The following example shows some sample device names on a system running VxVM 5.1 SP1:

  Example 1 Sample device names on a system using the `Enclosure Based Naming Scheme(ebn)` (default in VxVM 5.1 SP1)

```
DEVICE TYPE DISK GROUP STATUS
disk_0 auto:cdsdisk c4t0d0 dg1 online
disk_1 auto:LVM - - LVM
disk_2 auto:LVM - - LVM
disk_3 auto:LVM - - LVM
disk_4 auto:hpdisk rootdisk01 rootdg online
```

  To change the default naming scheme to the `Legacy Device Naming Scheme`, use the following command:

```
# vxddladm set namingscheme=osn mode=legacy
```

  The following example shows some sample device names on a system using the `Legacy Device Naming Scheme`:

  Example 2 Sample device names on a system using the `Legacy Device Naming Scheme`

```
DEVICE TYPE DISK GROUP STATUS
c0t6d0 auto:hpdisk rootdisk01 rootdg online
c3t6d0 auto:LVM - - LVM
c4t0d0 auto:cdsdisk c4t0d0 dg1 online
c4t3d0 auto:LVM - - LVM
c4t9d0 auto:LVM - - LVM
```

  To change the default naming scheme to the `Agile Device Naming Scheme`, use the following command:

```
# vxddladm set namingscheme=osn mode=new
```

The following example shows some sample device names on a system using the `Agile Device Naming Scheme`:

Example 3 Sample device names on a system using the `Agile Device Naming Scheme`

```
DEVICE TYPE DISK GROUP STATUS
disk6 auto:hpdisk rootdisk01 rootdg online
disk7 auto:LVM - - LVM
disk11 auto:cdsdisk c4t0d0 dg1 online
disk10 auto:LVM - - LVM
disk9 auto:LVM - - LVM
```

Only in cases where customers upgrade from an earlier version to this version, `Operating System Native Naming Scheme(osn)` or the setting from the earlier release will override `ebn`. So, customers will continue to see the `osn` naming scheme.

- **Issuing CVM commands from the slave node**

  In releases prior to VxVM 5.1 SP1, CVM required that you issue configuration commands for shared disk groups from the master node of the cluster. Configuration commands change the object configuration of a CVM shared disk group. Examples of configuration changes include creating disk groups, importing disk groups, deporting disk groups, and creating volumes.

  Starting with the VxVM 5.1 SP1 release, you can issue commands from any node, even when the command changes the configuration of the shared disk group.

- **Changing the CVM master online**

  CVM now supports changing the CVM master from one node in the cluster to another node, while the cluster is online. CVM migrates the master node, and re-configures the cluster. After the master change operation starts re-configuring the cluster, other commands that require configuration changes will fail.

  For more information on changing the CVM master while the cluster is online, refer to the *Veritas Volume Manager 5.1 SP1 Administrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

VxVM 5.1 SP1 on HP–UX 11i v3

For more information on features that VxVM 5.1 SP1 supports on HP-UX 11i v3, refer to the *Veritas Volume Manager 5.1 SP1 Release Notes*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Architecture of VxVM

VxVM operates as a subsystem between the HP-UX operating system and other data management systems, such as file systems and database management systems. VxVM is layered on top of the operating system and is dependent on it for the following:

- Physical access to disks
- Device handles
- VM disks
- Multipathing

VxVM Daemons

VxVM relies on the following daemons for its operation:

- `vxconfigd` – The VxVM configuration daemon maintains disk and disk group configuration information, communicates configuration changes to the kernel, and modifies the configuration information stored on the disks.
- `vxiod` – The VxVM I/O daemon provides extended I/O operations without blocking the calling processes.
- `vxrelocd` – The hot-relocation daemon monitors VxVM for events that affect redundancy, and performs hot-relocation to restore redundancy.
- `vxattachd` – The `vxattachd` daemon handles automatic reattachment and resynchronization for plexes.

VxVM Objects

VxVM supports the following types of objects:

- **Physical Objects**

  Physical disks or other hardware with block and raw operating system device interfaces that are used to store data.

- **Virtual Objects**

  The virtual objects in VxVM include the following:

  - **Disk Group**

    A group of disks that share a common configuration. A configuration consists of a set of records describing objects (including disks, volumes, plexes, and subdisks) that are associated with one particular disk group. Each disk group has an administrator-assigned name, which can be used by the administrator to reference that disk group. Each disk group also has an internally defined unique disk group ID, which is used to differentiate two disk groups with the same administrator-assigned name.

  - **VM Disks**

    When you place a physical disk under VxVM control, a VM disk is assigned to the physical disk. Each VM disk corresponds to one physical disk. A VM disk is under VxVM control and is usually in a disk group.

  - **Subdisks**

    A VM disk can be divided into one or more subdisks. Each subdisk represents a specific portion of a VM disk, which in turn is mapped to a specific region in a physical disk. VxVM allocates a set of contiguous blocks for a subdisk.

  - **Plexes**

    VxVM uses subdisks to build virtual objects called plexes. A plex consists of one or more subdisks located on one or more physical disks.

  - **Volumes**

    A volume is a virtual disk device that appears like a physical disk device to applications, databases, and file systems. However, VxVM volumes do not have the physical limitations of a physical disk device. A volume consists of one or more plexes, each holding a copy of the selected data in the volume.

Volume Layouts in VxVM

A volume layout is defined by the association of a volume to one or more plexes, each of which maps to a subdisk. VxVM supports two different types of volume layout:

- Non-Layered
- Layered

Non-Layered

In a non-layered volume layout, a subdisk maps directly to a VM disk. This enables the subdisk to define a contiguous extent of storage space backed by the public region of a VM disk.

Layered Volumes

A layered volume is constructed by mapping its subdisks to the underlying volumes. The subdisks in the underlying volumes must map to VM disks, and hence to the attached physical storage.

VxVM Storage Layouts

Data in virtual objects is organized to create volumes by using the following layouts:

- Concatenation and Spanning
- Striping (RAID-0)
- Mirroring (RAID-1)
- Striping Plus Mirroring (Mirrored-Stripe or RAID-0+1)
- Mirroring Plus Striping (Striped-Mirror, RAID-1+0 or RAID-10)
- RAID-5 (Striping with Parity)

Concatenation and Spanning

Concatenation maps data in a linear manner onto one or more subdisks in a plex. To access the data in a concatenated plex sequentially, data is first accessed from the first subdisk from beginning to end and then accessed in the remaining subdisks sequentially from beginning to end, until the end of the last subdisk.

Striping (RAID-0)

Striping maps data so that the data is interleaved among two or more physical disks. A striped plex contains two or more subdisks, spread out over two or more physical disks.

Mirroring (RAID-1)

Mirroring uses multiple mirrors (plexes) to duplicate the information contained in a volume. In the event of a physical disk failure, the plex on the failed disk becomes unavailable.

When striping or spanning across a large number of disks, failure of any one of the disks can make the entire plex unusable. As disks can fail, you must consider mirroring to improve the reliability (and availability) of a striped or spanned volume.

Striping Plus Mirroring (Mirrored-Stripe or RAID-0+1)

VxVM supports combination of mirroring above striping. This combined layout is called a mirrored-stripe layout. A mirrored-stripe layout offers the dual benefits of striping to spread data across multiple disks, while mirroring provides redundancy of data.

Mirroring Plus Striping (Striped-Mirror, RAID-1+0 or RAID-10)

VxVM supports the combination of striping above mirroring. This combined layout is called a striped-mirror layout. Putting mirroring below striping, mirrors each column of the stripe. If there are multiple subdisks per column, each subdisk can be mirrored individually instead of each column.

RAID-5 (Striping with Parity)

Although both mirroring (RAID-1) and RAID-5 provide redundancy of data, they use different methods. Mirroring provides data redundancy by maintaining multiple complete copies of the data in a volume. Data being written to a mirrored volume is reflected in all copies. If a portion of the mirrored volume fails, the system continues to use the other copies of the data. RAID-5 provides data redundancy by using parity. Parity is a calculated value used to reconstruct data, after a failure. If a portion of a RAID-5 volume fails, the data that was on that portion of the failed volume can be recreated from the remaining data and parity information. It is also possible to mix concatenation and striping in the layout.

VxVM Interfaces

VxVM provides the following interfaces:

- Command-Line Interface
- Menu-driven `vxdiskadm` utility
- Veritas Enterprise Administrator

Command-Line Interface

As a superuser, you can administer and configure volumes and other VxVM objects using the supported `vx*` commands.

Menu-driven utility

The `vxdiskadm` utility provides an easy to use menu driven interface for common high-level operations on disks and disk groups.

Veritas Enterprise Administrator

The Veritas™ Enterprise Administrator (VEA) is the graphical user interface for administering disks, volumes, and file systems on local and remote machines.

File Systems Supported on HP-UX 11i v3

Table 2 discusses the file systems that are supported on HP-UX 11i v3.

Table 2 Supported File Systems on HP-UX 11i v 3

File System Type	Description
Hierarchical File System (HFS)	HFS is derived from the UNIX File System, the original BSD file system.
Veritas File System (VxFS)	The Veritas File System is an extent-based, intent logging file system from Symantec Corporation.
Compact Disk File System (CDFS)	The CD file system enables you to read and write to compact disc media.
AutoFS	AutoFS/Automounter mounts directories automatically when users or processes request access to them. AutoFS also unmounts the directories automatically if they remain idle for a specified period of time.
Network File System (NFS)	Network File System (NFS) provides transparent access to files on the network. An NFS server makes a directory available to other hosts on the network by “sharing” the directory.
CacheFS	The Cache File System (CacheFS) is a general purpose file system caching mechanism that improves server performance and scalability by reducing server and network load.

Veritas File System (VxFS)

Introduction

The Veritas File System (VxFS) is a high availability, high performance, commercial grade file system that provides features such as transaction based journaling, fast recovery, extent-based allocation, and online administrative operations, such as backup, resizing, and defragmentation of the file system. It provides high performance and easy manageability required by mission-critical applications, where high availability is critical. It increases the I/O performance and provides structural integrity.

The Veritas File System version (VxFS) is the `vxfs` file system for HP-UX 11i v3 release.

VxFS Features

VxFS supports the following features:

- **Extent-Based Allocation**

  An extent is defined as one or more adjacent blocks of data within the file system. VxFS allocates storage in groups of extents rather than a block at a time, thereby resulting in faster read-write operations.

- **Extent Attributes**

  Extent attributes are the extent allocation policies associated with a file. VxFS allocates disk space to files in groups of one or more extents. VxFS enables applications to control extent allocation.

  The `setext` and `getext` commands enable the administrator to set or view extent attributes associated with a file, as well as to preallocate space for a file.

- **Fast File System Recovery**

  Most file systems rely on full structural verification by the `fsck` utility as the only means to recover from a system failure. For large disk configurations, this involves a time-consuming process of checking the entire structure, verifying that the file system is intact, and correcting any inconsistencies.

  VxFS reduces system failure recovery time by tracking file system activity in the VxFS intent log. This feature records pending changes to the file system structure in a circular intent log.

- **Access Control Lists (ACLs)**

  An Access Control List (ACL) stores a series of entries that identify specific users or groups, and their access privileges for a directory or file.

- **Online Administration**

  A VxFS file system can be defragmented and resized while it remains online and available to users. The online administrations operations supported are backup, resizing, and defragmentation.

- **File Snapshot**

  VxFS provides online data backup using the snapshot feature. An image of a mounted file system instantly becomes an exact read-only copy of the file system at a specific point in time. The original file system is called the snapped file system and the copy is called the snapshot.

- **Expanded Application Interface**

  VxFS supports the following specific features for commercial applications:

  - Pre-allocates space for files
  - Specifies fixed extent size for files

- Bypasses the system buffer cache for file I/O
- Specifies the expected access pattern for a file

- **Extended Mount Options**

  The extended mount options supported by VxFS include the following:

  - Enhanced data integrity modes
  - Enhanced performance modes
  - Temporary file system modes
  - Improved synchronous writes
  - Large file sizes

- **Large File and File System Sizes**

  File systems up to 40 TB and files up to 16 TB in size are supported on HP-UX 11i v3. For more information on files and file system sizes supported by VxFS, see the *Supported File and File System Sizes* white paper on HP Business Support Center.

- **Enhanced I/O Performance**

  VxFS provides enhanced I/O performance by applying an aggressive I/O clustering policy, integrating with VxVM, and allowing application specific parameters to be set on a per-file system basis. However, clustering support is not available with the current release.

- **Storage Checkpoints**

  To increase availability, recoverability, and performance, VxFS offers on-disk and online backup and restore utilities that facilitate frequent and efficient backup of the file system. Backup and restore applications can leverage the Storage Checkpoint, a disk and I/O-efficient copying technology for creating periodic frozen images of a file system. Storage Checkpoints present a view of a file system at a point in time, and subsequently identifies and maintains copies of the original file system blocks. Instead of using a disk-based mirroring method, Storage Checkpoints save disk space and significantly reduce I/O overhead by using the free space pool available to a file system.

- **Quotas**

  VxFS supports quotas, which allocate per-user quotas and limit the use of two principal resources files and data blocks.

- **Multi-Volume Support**

  The Multi-Volume support enables several volumes to be encapsulated into a single virtual object called volume set. This volume set can then be used to create a file system, thereby enabling advanced features such as Dynamic Storage Tiering.

- **SmartMove™ Feature**

  SmartMove reduces the time and I/O required to attach or reattach a plex to an existing VxVM volume, in the specific case where a VxVM volume has a VxFS file system mounted on

it. The SmartMove feature uses VxFS information to detect free extents and avoids copying them.

- **Dynamic Storage Tiering Enhancements**

  The Dynamic Storage Tiering (DST) feature provides the following enhancements:

  - Enhanced DST APIs to provide a new interface for managing allocation policies of storage checkpoints during creation and later, and for managing named data stream allocation policies
  - `fsppadm` support for user ID (UID), group ID (GID), and tagging (TAG) elements in the placement policy XML file
  - Improved scan performance in the `fsppadm` command
  - Suppressed processing of the chosen RULE
  - Parser support for UID, GID, and TAG elements in a DST policy
  - What-if support for analyzing and enforcing without requiring the policy to be assigned
  - Storage Checkpoint data placement support in a DST policy
  - Shared DB thread handle support
  - CPU and I/O throttling support for DST scans
  - New command, `fstag`, for file tagging
  - New command, `fsppmk`, for creating XML policies

- **Availability of the `mntlock` and `mntunlock` Mount Options**

  You can specify the `mntlock` option with the `mount` command to prevent a file system from being unmounted by an application.

- **Autolog replay on mount**

  Starting with the VxFS 5.1 SP1 release, when the `mount` command detects a dirty log in the file system, it will automatically run the VxFS command `fsck` to clean up the intent log. This functionality is only supported on file systems mounted on a Veritas Volume Manager (VxVM) volume.

- **FileSnap**

  FileSnaps provide the ability to snapshot objects that are smaller in granularity than a file system or a volume. This is an Enterprise level feature. It is supported only with DLV 8 filesystems.

- **SmartTier sub-file movement**

  The Dynamic Storage Tiering (DST) feature is now rebranded as SmartTier. With the SmartTier feature, you can now manage the placement of file objects as well as entire files on individual volumes.

- **Tuning performance optimization of inode allocation**

  Starting with the VxFS 5.1 SP1 release, you can optimize the way in which inodes are reused in inode cache by setting the `delicache_enable` tunable parameter. It specifies whether performance optimization of inode allocation and reuse during a new file creation is turned on or not.

- **Veritas File System is more thin-friendly tunable**

  Thin Provisioning is a storage array feature that optimizes storage use by automating storage provisioning. Administrators do not have to estimate how much storage an application requires. Instead, Thin Provisioning lets administrators provision large thin or thin reclaim capable LUNs

to a host. Physical storage capacity is allocated from a thin pool to the thin/thin reclaim capable LUNS only after application I/O writes.

Starting with the VxFS 5.1 SP1 release, you can tune VxFS to enable or disable thin-friendly allocations. This feature is only supported on file systems mounted on a VxVM volume.

- **Partitioned Directories**

  VxFS 5.1 SP1 allows you to create partitioned directories. For every new create, delete, or lookup thread that is created, VxFS searches for the thread's respective hash directory and performs the operation in that directory. This allows uninterrupted access to the parent directory inode and its other hash directories, which significantly improves the read/write performance of cluster file systems.

  This feature is supported only on file systems with DLV 8 or later.

  For more information, refer to the *Veritas File System 5.1 SP1 Administrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

VxFS 5.1 SP1 on HP-UX 11i v3

For more information on features supported with VxFS 5.1 SP1 on HP-UX 11i v3, refer to the *Veritas File System 5.1 SP1 Release Notes*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Architecture of VxFS

HP-UX supports various file systems. In order for the kernel to be able to access these different file system types, there is a layer of indirection above them called Virtual File System (VFS).

Without the VFS layer, the kernel must know the specifics of each file system type and maintain distinct code to handle each.

The VFS layer enables the kernel to possess a single set of routines that are common to all file system types. Handling of the specifics of a file system type are passed down to the file system specific modules. The following sections describe the VxFS file system specific structures.

The following are the VxFS on-disk structures:

- **Superblock**

  A superblock (SB) resides ~8k from the beginning of the storage and tracks the status of the file system. It supports maps of free space and other resources (inodes, allocation units, and so on).

- **Intent Log**

  VxFS reduces system failure recovery time by tracking file system activity in the VxFS intent log. This feature records pending changes to the file system structure in a circular intent log. The intent log recovery feature is not readily apparent to users or a system administrator, except during a system failure. During system failure recovery, the VxFS `fsck` utility performs an intent log replay, which scans the intent log, and nullifies or completes file system operations that were active when the system failed. The file system can then be mounted without completing a full structural check of the entire file system. Replaying the intent log may not completely recover the damaged file system structure if there was a disk hardware failure. Hardware problems may require a complete system check using the `fsck` utility provided with VxFS.

- **Allocation Unit**

  Allocation units are made up of a series of data blocks. Each allocation unit typically consists of 32k contiguous blocks. Several contiguous data blocks make up an extent. The extents are used for file data storage.

Extent Based Allocation

An extent is defined as one or more adjacent blocks of data within the file system. Extent based allocation offers the following advantages:

- Allows large I/Os for efficiency
- Supports dynamic resizing of disk space

2 System Requirements

This chapter discusses the various system requirements for the Veritas 5.1 SP1 suite of products. This chapter addresses the following topics:

- Software Dependency
- OS Version
- Patch Requirements
- Required Packages for VEA
- Software Depot Content
- License Bundles
- Disk Space Requirements

Software Dependency

- VxFS 5.1 SP1 works with both HP LVM and VxVM 5.1 SP1 on HP-UX 11i v3
- VxVM 5.1 SP1 works only when VxFS 5.1 SP1 is installed

OS Version

This release can only be installed on a system running the HP-UX 11i v3 March 2011 Operating Environment Upgrade Release (OEUR) or later on the PA-RISC or Itanium platforms.

To verify the operating system version, use the `swlist` command as follows:

```
# swlist | grep HPUX11i
HPUX11i-DC-OE      B.11.31.1103     HP-UX Data Center Operating Environment
```

Patch Requirements

VxFS 5.1 SP1 and VxVM 5.1 SP1 for HP-UX 11i v3 require certain patches to function correctly. In addition, certain patches are recommended for all installations.

Table 3 lists all the required and recommended patches for VxFS 5.1 SP1 and VxVM 5.1 SP1 for HP-UX 11i v3. You can obtain all these patches from http://itrc.hp.com. However, some of these patches may already be available to you in the HP-UX11i v3 OEUR. Table 3 lists each patch and mentions if the patch or a superseding patch is already included in a given OEUR. If the OEUR does not contain a particular patch, you must download the patch accordingly.

Table 3 Required and Recommended Patches

Patch	Available in HP-UX 11i v3 March 2011 OEUR
Required Patches	
PHKL_38651	Yes
PHKL_38952	Yes
PHKL_40944	Yes
PHKL_41086	Yes
PHSS_39898	Yes
Recommended Patches	
PHCO_41903	No
PHKL_40130	No

Table 3 Required and Recommended Patches *(continued)*

PHKL_40377	No
PHKL_41005	No
PHKL_41083	Yes
PHKL_41087	No
PHKL_41442	Yes

Required Software

In addition to the recommended patches listed in Table 3, the software products listed in Table 4 are required for use with VxFS 5.1 SP1 and VxVM 5.1 SP1 installations. You can download these products from http://www.software.hp.com.

Table 4 Required Software Products

Product Name: Revision
Ignite-UX C.7.12 or later
BaseLVM B.11.31.1104 or later

IMPORTANT: Ignite-UX version C.7.12 does not support archiving of disk groups with a disk group version of 160. Archive creation (when disk groups with a disk group version of 160 are present on the system) does not display any error/warning indicating restore operation will fail.

For more information, refer to the "Known Problems and Workarounds" section in the *Veritas Volume Manager 5.1 SP1 Release Notes*.

Required Packages for Veritas Enterprise Administrator

To use the VEA with VxVM 5.1 SP1, the following software products are required:

- Veritas Enterprise Administrator Service (VRTSob)
- Veritas Enterprise Administrator (VRTSobgui)

The minimum memory requirement for the VEA client is 128 MB. The software products mentioned in this section are installed as part of VxVM.

Software Depot Content

Table 5 and Table 6 lists the composition of the Base-VxFS-51 and Base-VxVM-51 bundles.

Table 5 Base-VxFS-51 Bundle Components

Package	Description
Base-VxFS-51	Veritas File System Bundle 5.1 for HP-UX
Base-VxFS-51.VRTSvlic	Symantec License Utilities
Base-VxFS-51.VRTSvxfs	VERITAS File System

Table 6 Base-VxVM-51 Bundle Components

Package	Description
Base-VxVM-51	Base VERITAS Volume Manager Bundle 5.1 for HP-UX
Base-VxVM-51.VRTSaslapm	Array Support Libraries and Array Policy Modules for Veritas Volume Manager

Table 6 Base-VxVM-51 Bundle Components *(continued)*

Package	Description
Base-VxVM-51.VRTSvxvm	Veritas Volume Manager by Symantec
Base-VxVM-51.VRTSvlic	Symantec License Utilities

Table 7 lists the Base-VxTools-51 bundle components.

Table 7 Base-VxTools-51 Bundle Components

Package	Description
Base-VxTools-51	VERITAS Infrastructure Bundle 5.1 for HP-UX
Base-VxTools-51.VRTSat	Symantec Product Authentication Service
Base-VxTools-51.VRTSvlic	Symantec License Utilities
Base-VxTools-51.VRTSsfmh	Veritas Storage Foundation Managed Host by Symantec
Base-VxTools-51.VRTSperl	Perl 5.10.0 for Veritas
Base-VxTools-51.VRTSobgui	Veritas Enterprise Administrator
Base-VxTools-51.VRTSob	VERITAS Enterprise Administrator Service

Table 8 lists the VxFS-SDK-51 bundle components.

Table 8 VxFS-SDK-51 Bundle Components

Package	Description
VxFS-SDK-51	VERITAS SDK Bundle 5.1 for HP-UX
VxFS-SDK-51.VRTSfssdk	Veritas File System Software Developer Kit

Table 9 (page 25) lists the CVM bundle components.

Table 9 B9117EB Bundle Components

Package	Description
B9117EB	VERITAS Cluster Volume Manager 5.1 for HP-UX
B9117EB.VRTSwl	VERITAS Product Enabler
B9117EB.VRTSvxfen	Veritas Fencing by Symantec
B9117EB.VRTSllt	Veritas Low Latency Transport by Symantec
B9117EB.VRTSgab	Veritas Group Membership and Atomic Broadcast by Symantec

License Bundles

Following license bundles are available for the Veritas 5.1 SP1 suite of products on HP-UX 11i v3:

- OnlineJFS (B3929HB): OnlineJFS for Veritas File System 5.1 SP1 Bundle
- Full VxVM (B9116EB): Full VxVM License for Veritas Volume Manager 5.1 SP1
- CVM License (B9117EB): VERITAS Cluster Volume Manager 5.1 SP1 for HP-UX
- For more information on these licenses and the HP Serviceguard Storage Management Licenses, see *HP Serviceguard Storage Management Suite Version A.04.00 Release Notes*.

Disk Space Requirements

Table 10 and Table 11 list the disk space requirements for VxFS 5.1 SP1 and VxVM 5.1 SP1, respectively.

Disk Space Requirements for VxFS 5.1 SP1

Table 10 lists the disk space required by VxFS 5.1 SP1.

Table 10 Minimum Space Required for Each Directory for VxFS 5.1 SP1

Package	/stand	/sbin	/usr	/opt	/etc	/var	Total
VRTSvxfs	—	—	74044 KB	49791 KB	—	4 KB	123839 KB
VRTSvlic	—	—	3774 KB	1979 KB	—	—	5753 KB

NOTE: Ensure that the /opt directory exists and has write permissions for root.

Disk Space Requirements for VxVM 5.1 SP1

Table 11 lists the disk space required by VxVM 5.1 SP1.

Table 11 Minimum Space Required for Each Directory for VxVM 5.1 SP1

Package	/	/stand	/usr	/opt	/var	Total
Base-VxVM-51.VRTSaslapm	1403 KB	—	4699 KB	—	—	6102 KB
Base-VxVM-51.VRTSvxvm	463872 KB	—	363805 KB	28030 KB	18 KB	855725 KB
Base-VxVM-51.VRTSvlic	4366 KB	—	3774 KB	1979 KB	—	10119 KB
Base-VxTools-51.VRTSat	1371 KB	—	—	207359 KB	29 KB	208759 KB
Base-VxTools-51.VRTSsfmh	—	—	—	142756 KB	—	142756 KB
Base-VxTools-51.VRTSperl	—	—	—	89245 KB	—	89245 KB
Base-VxTools-51.VRTSobgui	—	—	—	126034 KB	—	126034 KB
Base-VxTools-51.VRTSob	205 KB	—	—	126034 KB	—	126239 KB

Disk Space Requirements for CVM

To install CVM, a minimum of 434000 KB of space is required.

3 Installing the Veritas 5.1 SP1 Products

This chapter describes how to install VxFS 5.1 SP1 and VxVM 5.1 SP1 with the *swinstall* (1M) command. This chapter addresses the following topics:

Prior to installing Veritas 5.1 SP1 products, ensure that you have met the following conditions:

- Verify that all the required patches are installed.
- OnlineJFS 4.1(B3929EA), OnlineJFS 5.0(B3929FB) and OnlineJFS 5.0.1(B3929GB) are removed from the system (if installed).

Mounting the HP Serviceguard Storage Management Suite Media

For a media-based installation, you must mount the media before starting the installation process.

To mount the media, complete the following steps:

1. Insert the media into the drive, and log in as root by entering the following command:

   ```
   $ su root
   ```

   For more information on the supported options, see *su* (1).
2. Scan for the device name, as follows:

   If you are using the legacy I/O format, enter the following command:

   ```
   # ioscan -fnC disk
   ```

   For more information on the supported options, see *ioscan* (1M).

   If you are using the new I/O format, enter the following command:

   ```
   # ioscn -fNC disk
   ```
3. Create a mount point for the media, by entering the following command:

   ```
   # mkdir -p /cdrom
   ```

   For more information on the supported options, see *mkdir* (1).
4. Mount the media, by entering the following command:

   ```
   # mount <absolute device-path> /cdrom
   ```

   where `<absolute device-path>` is the device path for the CDROM.

   For more information on the supported options, see *mount_vxfs* (1M).
5. Verify that the media is mounted, by entering the following command:

   ```
   # mount
   ```

   For more information on supported options, see *mount_vxfs* (1M).

Installing Veritas 5.1 SP1 Products

Veritas 5.1 SP1 products are supported with HP-UX 11i v3 March 2011 OEUR and later. If you do not have the March 2011 OEUR, it is recommended that you upgrade the HP-UX version on

your system before installing the 5.1 SP1 version of the Veritas software. As the HP-UX 11i v3 March 2011 OEUR release contains the VxVM 5.0 and VxFS 5.0 products by default, you cannot directly install the 5.1 SP1 products. You must upgrade the Veritas 5.1 SP1 products from 5.0 to 5.1 SP1 after installing HP-UX 11i v3 March 2011 OEUR release.

Installing VxFS 5.1 SP1

VxFS 5.1 SP1 consists of the Base-VxFS-51 and HP OnlineJFS (`B3929HB`) bundles. The following sections discuss the installation of these bundles using the HP-UX Software Distributor (SD) commands.

NOTE: You cannot install the Base-VxFS and the Base-VxVM 5.1 SP1 products on a system where the Veritas 4.1 versions (for HP-UX 11i v2) of HP Serviceguard Storage Management, or HP Serviceguard Cluster File System product suites are installed. The AVXFS and AVXVM products (included in VxFS 5.1 SP1 and VxVM 5.1 SP1, respectively) check the installed products and the cause for an installation failure for the Veritas 5.1 SP1 products on such systems. This issue is applicable only for systems that have been upgraded from HP-UX 11i v2 to HP-UX 11i v3.

Tools such as AVXFS, AVXVM, AVXTOOLS, AFULLVXVM, and AONLINEJFS help pull in dependant products. For example, when you select VxVM 5.1 SP1, these tools automatically pull in VxFS 5.1 SP1. These products also address issues of mutual exclusions.

Installing Base-VxFS-51

You can install Base-VxFS-51 either in the non-interactive mode or in the interactive mode.

Installing Base-VxFS-51 in Non-Interactive Mode

Enter the `swinstall` command in the non-interactive mode, as follows:

```
# swinstall -x autoreboot=true -s <depot-path> Base-VxFS-51
```

For more information on the supported options, see *swinstall* (1M).

NOTE: If you are installing VxFS 5.1 SP1 in a system that contains the HP-UX Data Center Operating Environment (DC-OE), HP-UX Virtual Server Operating Environment (VSE-OE) or HP-UX High Availability Operating Environment (HA-OE), you must either include OnlineJFS 5.1 (B3929HB) in the `swinstall` command line or remove OnlineJFS 5.0.1 (B3929GB) prior to installing VxFS 5.1 SP1.

Installing Base-VxFS-51 in Interactive Mode

To install Base-VxFS-51 in interactive mode, complete the following steps:

1. Enter the `swinstall` command in interactive mode, as follows:

   ```
   # swinstall -x autoreboot=true -s <depot-path>
   ```

   For more information on the supported options, see *swinstall* (1M).
2. Mark the Base-VxFS-51 bundle in the SD Install window.

   NOTE: If you are installing VxFS 5.1 SP1 in a system that includes the HP-UX Data Center Operating Environment (DC-OE), HP-UX Virtual Server Operating Environment (VSE-OE) or HP-UX High Availability Operating Environment (HA-OE), you must either mark OnlineJFS 5.1 (B3929HB) in the SD install dialog box or remove OnlineJFS 5.0.1 (B3929GB) prior to installing VxFS 5.1 SP1.
3. Select **Actions**, and click **Install**. Follow the on-screen instructions to complete the installation.

The system reboots automatically after the installation is complete. You can monitor the installation process for warnings and notes. For more information on installation, see the log file `/var/adm/sw/swagent.log`.

Verifying Base-VxFS-51 Installation

Verify the Base-VxFS-51 installation on your system by entering the following command:

```
# swverify Base-VxFS-51
```

For more information on the supported options, see *swverify* (1M).

Installing HP OnlineJFS (B3929HB)

The HP OnlineJFS (`B3929HB`) license enables additional VxFS functionality not available with Base-VxFS-51.

You can install HP OnlineJFS (`B3929HB`) either in the non-interactive or in the interactive mode.

NOTE: Base-VxFS-51 must be installed on your system before installing HP OnlineJFS. Otherwise the installation of HP OnlineJFS may fail with an error.

Installing HP OnlineJFS (B3929HB) in Non-Interactive Mode

Enter `swinstall` command in non-interactive mode, as follows:

```
# swinstall -s <depot-path> B3929HB
```

For more information on the supported options, see *swinstall*(1M).

Installing HP OnlineJFS (B3929HB) in the Interactive Mode

To install HP OnlineJFS (B3929HB) in the interactive mode, complete the following steps:

1. Enter the `swinstall` command in interactive mode, as follows:

   ```
   # swinstall -s <depot-path>
   ```

2. Mark the `B3929HB` bundle on the SD Install Window.
3. Select **Actions**, and click **Install**. Follow the on-screen instructions to complete the installation.

Verifying HP OnlineJFS (B3929HB) Installation

Verify the installation of OnlineJFS (`B3929HB`) on your system by entering the following command:

```
# swverify B3929HB
```

For more information on the supported options, see *swverify*(1M).

Installing VxVM 5.1 SP1

VxVM 5.1 SP1 consists of the Base-VxVM-51 and Full VxVM (`B9116EB`) bundles. The following sections discuss the installation of these bundles using the HP-UX SD commands.

NOTE: After installing VxVM 5.1 SP1, if there were any prior VxVM versions that had unbundled Array Support Library (ASL) packages installed, the new unbundled ASL packages will not be automatically installed with 5.1 SP1. You must download and install 5.1 SP1 unbundled ASLs separately or `swcopy` VxVM 5.1 SP1 together with VxFS 5.1 SP1 and ASL packages into a single depot and install the depot. The ASLs bundled together with the base product will be updated automatically. For more information on ASLs, see the Hardware Compatibility List (HCL) available on HP Business Support Center.

NOTE: When installing VxVM 5.1 SP1 on a system containing an earlier version of VxVM, any earlier version of VxVM will be removed automatically as part of the post-install configuration.

Installing Base-VxVM-51

You can install Base-VxVM-51 either in the non-interactive mode or in the interactive mode.

NOTE: Ensure that you select the complete `Base-VxVM-51` bundle for installation while installing Full-VxVM bundle `B9116EB` if your system already does not have Base-VxVM-51.

Also, verify that the bundles Base-VxVM-51, Base-VxFS-51, Base-VxTools-51 and `B9116EB` are present on your system after installing `B9116EB`.

Installing Base-VxVM-51 in Non-Interactive Mode

Enter the `swinstall` command in non-interactive mode, as follows:

```
# swinstall -x autoreboot=true -s <depot-path> Base-VxVM-51
```

For more information on the supported options, see *swinstall*(1M).

NOTE: If you are installing VxFS 5.1 SP1 in a system that includes the HP-UX Data Center Operating Environment (DC-OE), HP-UX Virtual Server Operating Environment (VSE-OE) or HP-UX High Availability Operating Environment (HA-OE), you must either include OnlineJFS 5.1 (B3929HB) on the swinstall command line or remove OnlineJFS 5.0.1 (B3929GB) prior to installing VxFS 5.1 SP1.

The installation of VxVM also installs the VEA service and client packages. Installation of VxVM 5.1 SP1 picks up VxFS 5.1 SP1 as a corequisite.

Installing Base-VxVM-51 in the Interactive Mode

To install Base-VxVM-51 in the interactive mode, complete the following steps:

1. Enter the `swinstall` command in the interactive mode as follows:

   ```
   # swintall -x autoreboot=true -s <depot-path>
   ```

   For more information on the supported options, see *swinstall* (1M).
2. Mark the Base-VxVM-51 bundle on the SD Install Window.

   NOTE: If you are installing VxFS 5.1 SP1 in a system that includes the HP-UX Data Center Operating Environment (DC-OE), HP-UX Virtual Server Operating Environment (VSE-OE) or HP-UX High Availability Operating Environment (HA-OE), you must either mark OnlineJFS 5.1 (B3929HB) in the SD Install window or remove OnlineJFS 5.0.1 (B3929GB) prior to installing VxFS 5.1 SP1.
3. Select **Actions**, and click **Install**. Follow the on-screen instructions to complete the installation.

Verifying the Base-VxVM-51 Installation

Verify the Base-VxVM-51 installation on your system by entering the following command:

```
# swverify Base-VxVM-51
```

For more information on the supported options, see *swverify* (1M).

Installing Full VxVM (B9116EB)

The Full VxVM (`B9116EB`) license enables additional VxVM functionality that are not available with Base-VxVM-51.

You can install Full VxVM either in the non-interactive mode or in the interactive mode.

Installing Full VxVM in Non-Interactive mode

Enter the `swinstall` command in the non-interactive mode as follows:

```
# swinstall -x autoreboot=true -s <depot-path> B9116EB
```

For more information on the supported options, see *swinstall*(1M).

Installing Full VxVM in Interactive Mode

To install Full VxVM in the interactive mode, complete the following steps:

1. Enter the `swinstall` command in the interactive mode, as follows:

   ```
   # swinstall -s <depot-path>
   ```

   For more information on supported options, see *swinstall* (1M).
2. Mark the `B9116EB` bundle on the SD Install window.
3. Select **Actions**, and click **Install**. Follow the on-screen instructions to complete the installation.

NOTE: The `B9116EB` bundle contains a license for the Full VxVM 5.1 SP1 functionality.

Verifying the Full VxVM (B9116EB) Installation

Verify the Full VxVM `(B9116EB)` installation by enter the following command:

```
# swverify B9116EB
```

For more information on the supported options, see *swverify*(1M).

Use the `vxlicrep` command to check for the installed licenses.

Installing CVM [B9117EB] on HP-UX 11i v3

To install CVM on HP-UX 11i v3, complete the following steps:

1. Use a clean ignited HP-UX 11i v3 system.
2. If the ignited system has Base-VxVM 4.1 already installed on it, do not clean the system before installing 5.1 SP1 v3
3. Install Base-VxFS-51 , Base-VxTools-51, Base-VxVM-51, and VxFS-SDK-51 by following the instructions discussed in the previous sections.
4. Install CVM [B9117EB] by entering the following command:

   ```
   swinstall -x autoreboot=true -s <-depot_source> B9117EB
   ```

Cold-Installing VxVM 5.1 SP1 and VxFS 5.1 SP1 with HP-UX 11i v3

This section describes the process to cold-install a system with HP-UX 11i v3, VxVM 5.1 SP1, and VxFS 5.1 SP1. A depot with the HP-UX 11i v3 Operating Environment (OE) containing VxFS 5.0 and VxVM 5.0 and a depot with the Veritas VxFS 5.1 SP1 and VxVM 5.1 SP1 software are used for the installation. Veritas VxVM 5.0 and VxFS 5.0 are included on the 11i v3 OE starting in September 2008, and VxFS/VxVM 5.1 SP1 will need to be copied from an Independent media set.

Preparing the Ignite-UX Server

To prepare the Ignite-UX server, complete the following steps:

1. Enter the `swinstall` command to install Ignite-UX version C.7.12 (available March, 2011) or later on an Ignite-UX server. You can confirm the Ignite software version by entering the command `swlist | grep -i ign`, or by viewing the contents of the `/opt/ignite/Version` file.
2. Create the configuration file for the HP-UX 11i v3 OE depot.
   1. `# cd /var/opt/ignite/data/Rel_B.11.31`
   2. ```
      # /opt/ignite/bin/make_config \
      - s <depot server>:<1131_OE_depot> \
      - c <OE_config_filename>_cfg
      ```
   3. Test the configuration file.

      ```
      # /opt/ignite/bin/instl_adm -T -f <OE_config_filename>_cfg
      ```
      ```

3. Create the configuration file for the VxVM 5.1 SP1 depot.
   1. `# cd /var/opt/ignite/data/Rel_B.11.31`
   2. ```
      # /opt/ignite/bin/make_config \
      - s <1131_Vx*5.1_depot> - c <VX_config_filename>_cfg
      ```
   3. Test the configuration file.

      `# /opt/ignite/bin/instl_adm -T -f <VX_config_filename>_cfg`
4. Combine both configuration files into one Ignite-UX INDEX configuration clause:
   1. `# vi /var/opt/ignite/INDEX`
   2. Append a new configuration clause, such as:

      ```
      cfg "HP-UX B.11.31_Mar2011_OE-Vx51" {
      description "B.11.31_Mar2011_OE-Vx51 - 2 depots"
      "/opt/ignite/data/Rel_B.11.31/config"
      "/var/opt/ignite/data/Rel_B.11.31/<0E_config_filename>_cfg"
      "/var/opt/ignite/data/Rel_B.11.31/<VX_config_filename>_cfg"
      "/opt/ignite/data/Rel_B.11.31/hw_patches_cfg"
      "/var/opt/ignite/config.local"
      }
      ```
   3. Test all the configuration files in the INDEX file.

      `# /opt/ignite/bin/instl_adm -T`
5. If you plan to LAN boot the client, add its boot entry in the appropriate file on the server. This step is not required if you plan to initiate the cold-install from the Ignite-UX server and the client is up.
   1. For an Intel® Itanium®-based client:
      a. `# vi /etc/bootptab`
      b. Upate the System-IPF entry with client_hostname and other relevant values.
   2. For a PA-RISC client:
      a. `# vi /etc/opt/ignite/instl_boottab`
      b. Append an entry for the client using the following format:

         `<client_IP_addr>:<client_MAC_addr>::reserve # <client_hostname>`

Cold-Installing the Client

To cold-install the client, complete the following steps:

1. Boot the client, whose hardware supports B.11.31. For a list of supported systems, see the *HP-UX 11i v3 Release Notes* .
2. Use Ignite-UX to cold-install using the two depots referenced in the Ignite-UX server INDEX configuration clause created above. In the Ignite-UX interface, make your selection in the **Basic** tab, **Configurations**line, such as **HP-UX B.11.31_Mar2011_OE-Vx51**.
3. Use the Ignite-UX interface to select the volume manager. On the **Basic** tab, **File System** line, select the following:

   VERITAS Volume Manager (VxVM) with VxFS
4. Use the Ignite-UX interface to select VxVM 5.1 SP1 and VxFS 5.1 SP1 software. On the **Software** tab, select the following:

Base-VxFS-51	Veritas File System Bundle 5.1 for HP-UX
Base-VxTools-51	VERITAS Infrastructure Bundle 5.1 for HP-UX
Base-VxVM-51	Base VERITAS Volume Manager Bundle 5.1 for HP-UX
5. Use the Ignite-UX interface to deselect the VxVM 5.0 & VxFS 5.0 software, if it is selected. On the **Software tab**, deselect the following:
      ```

Base-VxFS-50(OE)	Veritas File System Bundle 5.0 for HP-UX
Base-VxTools-50(OE)	ERITAS Infrastructure Bundle 5.0 for HP-UX
Base-VxVM-50(OE)	Base VERITAS Volume Manager Bundle 5.0 for HP-UX
B3929FB	OnlineJFS for Veritas File System 5.0 Bundle

6. Make any other configuration choices needed using the Ignite-UX interface, and then select **Go!** to cold-install the client.

Confirming the Client

After the system boots, confirm that the client has the intended software and configuration. To confirm the software and the configuration, complete the following steps:

1. Confirm the correct hostname:

   ```
   # hostname
   ```

2. Confirm the correct OE version:

   ```
   # swlist | grep HPUX
   ```

3. Confirm installation of Base VERITAS Volume Manager Bundle 5.1 for HP-UX 11i v3:

   ```
   # swlist Base-VxVM-51
   ```

4. Confirm installation of Veritas File System Bundle 5.1 for HP-UX 11i v3:

   ```
   # swlist Base-VxFS-51
   ```

The following documents available on HP Business Support Center provide additional information on cold installation:

- *Ignite-UX Administration Guide*
- *HP-UX 11i v3 Installation and Update Guide*
- *HP-UX 11i Version 3 Release Notes*
- *HP-UX System Administrator's Guide*
- *Installation of VxFS and VxVM 5.1 SP1 on HP-UX 11i v3 March 2011 Operating Environment Update Release* whitepaper

Updating HP-UX and Veritas Products Using the `update-ux` Command

This section describes how to update an HP-UX 11i v2 system to HP-UX 11i v3 with Veritas VxVM 5.1 SP1 and VxFS 5.1 SP1 using the `update-ux` command.

To update the operating system from HP-UX 11i v2 to HP-UX 11i v3, complete the following steps:

1. Copy the contents of the HP-UX 11i v3 OE media to an HP-UX depot by entering the `swcopy -s` command:

   For example, to copy the contents of a DVD, such as `/11v3OE_DVD` to a depot called `/var/depots/my_hpux_depot`, enter the following command:

   ```
   # swcopy -s /11iv3OE_DVD \* @ /var/depots/my_hpux_depot
   ```

   If there are multiple HP-UX 11i v3 DVDs, repeat this step for DVD.

2. Alternatively, you can copy the DVD contents in an interactive mode wherein you mount the DVD and use the Find Local CD-ROM option in the Specify Source screen by entering the following command:

   ```
   # swcopy -i \* @ /var/depots/my_hpux_depot
   ```

3. Copy the contents of the HP-UX ServiceGuard and StorageManagement Products media into the depot folder you created:

   ```
   # swcopy -s /11iv3SG_DVD \* @ /var/depots/my_hpux_depot
   ```

For example, copy the `/11iv3SG_DVD` into the same HP-UX depot.

If there are multiple ServiceGuard and StorageManagement DVDs, you must repeat this step for each DVD.

4. If you plan to use this HP-UX depot to update systems other than the system where the depot resides, you must register the HP-UX depot. The `swcopy` command automatically registers the depot when it is first created. However, HP recommends that you register the depot again such that it is visible to other systems. Enter the following command to register the HP-UX depot:

   ```
   # swreg -l depot @ /var/depots/my_hpux_depot
   ```

5. On each system that must be updated to HP–UX 11i v3, enter the following command:

   ```
   # update-ux -s depot_server_name:/var/depots/my_hpux_depot \
   !Base-VxVM-50 B3929HB \
   B9116EB Base-VxFS-51 Base-VxTools-51 Base-VxVM-51 VxFS-SDK-51
   ```

   If the HP-UX depot resides on the system that must be updated, you need not specify the name of the depot server using the `depot_server_name` option. The `depot_server_name` option indicates the hostname of the system where the `my_hpux_depot` resides followed by a colon (`:`).

   This command updates the system to HP-UX 11i v3 and VxVM 5.1 SP1 and VxFS 5.1 SP1. The `!Base-VxVM-50` entries on the command line ensure that VxVM 5.0 is not installed from the depot during the update.

When you update HP-UX 11i v2 systems that include VxVM 3.5, the following message is displayed:

```
WARNING: The following was found while selecting software:
         - There are new revisions of SW available in the source that would
           update SW already installed on the system. The newer revisions are
           not selected for installation. This will leave the incompatible old
           revision of SW on the system. The affected SW is:
           - VRTSfspro.FSPRO-PRG,r=4.1x replaces
             VRTSfspro.FSPRO-PRG,r=3.5_2_ga08.006
           - VRTSfspro,r=4.1x replaces VRTSfspro,r=3.5_2_ga08.006
           - VRTSob.VEAS-FILESET,r=3.2.555.1 replaces
             VRTSob.VEAS-FILESET,r=3.2.532.0.001
           - VRTSob,r=3.2.555.1 replaces VRTSob,r=3.2.532.0.001
           - VRTSobgui.VEA-FILESET,r=3.2.555.1 replaces
             VRTSobgui.VEA-FILESET,r=3.2.532.0.001
           - VRTSobgui,r=3.2.555.1 replaces VRTSobgui,r=3.2.532.0.001
           - VRTSvlic.VLIC-RUN,r=3.02.23.1 replaces
             VRTSvlic.VLIC-RUN,r=3.01.IA.002.010
           - VRTSvlic,r=3.02.23.1 replaces VRTSvlic,r=3.01.IA.002.010
           - VRTSvmdoc.VXVM-DOC,r=4.1.032 replaces
             VRTSvmdoc.VXVM-DOC,r=3.5~IA.014
           - VRTSvmdoc,r=4.1.032 replaces VRTSvmdoc,r=3.5~IA.014
           - VRTSvmpro.VMPRO-PRG,r=4.1.020 replaces
             VRTSvmpro.VMPRO-PRG,r=3.5~IA.014
           - VRTSvmpro,r=4.1.020 replaces VRTSvmpro,r=3.5~IA.014
           - VRTSvxvm.VXMS,r=4.1.032 replaces VRTSvxvm.VXMS,r=3.5~IA.014
           - VRTSvxvm.VXVM-ENG-A-MAN,r=4.1.032 replaces
             VRTSvxvm.VXVM-ENG-A-MAN,r=3.5~IA.014
           - VRTSvxvm.VXVM-HELP,r=4.1.032 replaces
             VRTSvxvm.VXVM-HELP,r=3.5~IA.014
           - VRTSvxvm.VXVM-KRN,r=4.1.032 replaces
             VRTSvxvm.VXVM-KRN,r=3.5~IA.014
           - VRTSvxvm.VXVM-RUN,r=4.1.032 replaces
             VRTSvxvm.VXVM-RUN,r=3.5~IA.014
           - VRTSvxvm,r=4.1.032 replaces VRTSvxvm,r=3.5~IA.014
```

You can ignore this message if it refers only to one or more VRTS products. To ignore this message, press **`Enter`** when prompted, and continue with the update.

For more information on updating your system to HP-UX 11i v3 using the `update-ux` command, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs On this page, select HP-UX 11i v3.

For more information on managing depots, see *Software Distributor Administration Guide: HP-UX 11i v1, 11i v2, and 11i v3*. To locate this document, go to the HP Business Support Center.

4 Setting up the Veritas 5.1 SP1 Products

This chapter describes how to initialize and set up VxVM 5.1 SP1 and VxFS 5.1 SP1. This chapter addresses the following topics:

Configuring Your System after the Installation

To configure your system after the software upgrade, complete the following steps:

1. Reinstate the mount points in the `/etc/fstab` file.
2. Restart the upgraded systems.
3. Restart all the volumes by entering the following command for each disk group:

   ```
   # vxvol -g diskgroup startall
   ```

   You can perform the following optional configuration steps for cluster and Oracle services:

   - Stop the cluster, restore the cluster configuration files, and restart the cluster.
   - To create root volumes that are under VxVM control after installation, use the `vxcp_lvmroot` command. For more information, see "Converting to a VxVM Root Disk" (page 36)" and the *Veritas Volume Manager Administrator's Guide*.
   - To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions specified in sections "Upgrading VxFS Disk Layout Versions" (page 64) and "Upgrading VxVM Disk Group Versions" (page 66).

Converting to a VxVM Root Disk

You can choose VxVM for your root disk when performing a new installation using Ignite-UX. Alternatively, you can use the following procedure to achieve VxVM rootability by cloning your LVM root disk using the `vxcp_lvmroot` command.

To convert to a VxVM root disk from LVM, complete the following steps:

1. Select the disk to be used as your new VxVM root disk.

   HP recommends that this disk is internal to the main computer cabinet. If this is currently an LVM disk, you must remove the disk from LVM control as follows:

   - Use the `vgreduce` command to remove the disk from any LVM volume groups to which it belongs.
   - Use the `lvremove` command to remove any LVM volumes that are using the disk.
   - Use the `pvremove` command to erase the LVM disk headers

If the disk to be removed is the last disk in the volume group, use the `vgremove` command to remove the volume group, and then use the `pvremove` command 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 the `pvcreate` command, you must still use the `pvremove` command to remove LVM disk headers.

If you want to mirror the root disk across multiple disks, ensure that all the disks are free from LVM control.

2. Invoke the `vxcp_lvmroot` command to clone the LVM root disk to the disks you have designated to be the new VxVM root disks, while started on the upgraded LVM root disk. Following is an example `vxcp_lvmroot` command where `c1t0d0` is used for the target VxVM root disk:

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

   To additionally create a mirror of the root disk on `c2t0d0`, use the following command:

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

   HP recommends that you use the verbose (`-v`) option to create the mirror of the root disk. The process to clone the root disk takes a long time to complete and the `-v` option provides a time-stamped progress indication of each copied volume and other major events.

3. Use the `setboot` command to save the hardware path of the new VxVM root disk in the system. You can find the disk hardware paths using the following command:

   ```
   # ioscan -kfnC disk
   ```

   For information on the `setbook` command, see *setboot*(1M)

4. Reboot from the new VxVM root disk. If you have created a mirrored root disk, you need not explicitly take backup of the root disk. The LVM root disk safely co-exists with your VxVM root disk and provides a backup boot target.
5. Convert the original LVM root disk (if required) into a mirror of your VxVM root disk by using the following commands:

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

   After this operation is complete, the system runs on a completely mirrored VxVM root disk.

6. Use the `vxres_lvmroot` command to restore the LVM root disk, if required later.

Starting and Enabling the Configuration Daemon

The VxVM configuration daemon (`vxconfigd`) maintains VxVM disk and disk group configurations. `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 start the `vxconfigd` daemon and to check whether the daemon is enabled and running, and how to enable the daemon:

- **Starting the `vxconfigd` Daemon**

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

  ```
  # vxconfigd
  ```

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

  ```
  mode: disabled
  ```

- **Determining Whether the `vxconfigd` Daemon is Enabled**

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

```
# vxdctl mode
```

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

```
mode: enabled
```

This message indicates that `vxconfigd` is not running:

```
mode: not-running
```

- **Enabling the Volume Daemon**

  To enable the volume daemon, enter the following command:

  ```
  # vxdctl enable
  ```

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

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

Starting the Volume I/O Daemon

This section discusses the procedure to verify whether the `vxiod` daemons are running and to start them if necessary. 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 initial installation and they must be running at all times.

As the `vxiod` daemon is a kernel thread, it is not visible using the `ps` command. Therefore, to verify that `vxiod` daemons are running, enter the following command:

```
# vxiod
```

Following is a sample output if 16 `vxiod` daemons are running:

```
16 volume I/O daemons running
```

where:

`16` specifies the number of `vxiod` daemons currently running

If `vxiod` daemons are not currently running, run the following command to start the daemons:

```
# vxiod set 16
```

where:

`16` is the desired number of `vxiod` daemons that you want to run.

HP recommends that at least one `vxiod` daemon runs for each CPU in the system.

For more information on the `vxiod` daemon, see *vxiod*(1M).

Enabling the Intelligent Storage Provisioning Feature

This section discusses how to enable the Intelligent Storage Provisioning (ISP) feature.

If you load the allocator provider package (VRTSalloc), enter the following commands to restart the VEA service and to enable the Intelligent Storage Provisioning (ISP) feature:

```
# /opt/VRTS/bin/vxsvcctrl stop
# /opt/VRTS/bin/vxsvcctrl start
```

Enabling Cluster Support in VxVM

This release includes an optional cluster feature that enables VxVM to be used in a cluster environment. The cluster functionality in VxVM enables 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. The VxVM cluster feature requires a license that you can obtain from your Customer Support channel.

To enable the cluster functionality in VxVM, complete the following steps:

1. Obtain a license for the VxVM cluster feature.
2. Install the software packages on each node in the cluster.
3. Initialize VxVM.
4. Start VEA.
5. Configure shared disks.

   For information on configuring shared disks, see the *Veritas Volume Manager Administrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

If you are installing VxVM for the first time or adding disks to an existing cluster, you must configure new shared disks. Ensure that you configure shared disks only from one node and also specify if the disk is shared because VxVM cannot inform whether a disk is shared. While you are performing the configuration, ensure that the shared disks are not being accessed from another node. 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. If you are upgrading VxVM from a previous version, verify that your shared disks still exist.

For information on configuring new disks, verifying existing shared disks, and converting existing VxVM disk groups to shared disk groups, see the following sections:

- "Configuring New Shared Disks" (page 39)
- "Verifying Existing Shared Disks" (page 39)
- "Converting Existing VxVM Disk Groups to Shared Disk Groups" (page 40)

Configuring New Shared Disks

If you are installing and setting up VxVM for the first time, you must configure the shared disks.

To configure shared disks, complete the following steps:

1. Start the cluster on at least one node.
2. Run the `vxdiskadm` program on the node and choose option `1` to initialize new disks. When prompted to add these new disks to a disk group, choose `none` to leave the disks for future use.
3. View the newly initialized disks on other nodes in the cluster:

   ```
   # vxdctl enable
   ```

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

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

5. From the master node, use the`vxassist` command or VEA to create volumes in the disk groups. The volumes must be of type `gen`. Do not create RAID-5 volumes. Before creating any log subdisks, read the section on DRL in the *Veritas Volume Manager Adminstrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.
6. If the cluster is running with only one node, bring up the other cluster nodes. Enter the `vxdg list` command on each node to display the shared disk groups.

Verifying Existing Shared Disks

If you are upgrading from a previous release of VxVM, verify that your shared disk groups still exist.

To verify that your shared disk groups exist, complete the following steps:

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

   ```
   # vxdg -s list
   ```

   This command displays the existing shared disk groups.

Converting Existing VxVM Disk Groups to Shared Disk Groups

This section discusses the procedure to upgrade from VxVM 3.x to VxVM 5.1 SP1 and to convert existing disk groups to shared disk groups.

To convert existing disk groups to shared disk groups, complete the following steps:

1. Ensure that all systems that are running are part of the same cluster.
2. 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.
3. Configure the disk groups:
   - 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 that they are active.
   - 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.
4. 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.

Upgrading in a Clustered Environment with FastResync

In order to upgrade in a clustered environment with the FastResync option, you must perform the procedure discussed in this section in addition to the procedure discussed in "Converting Existing VxVM Disk Groups to Shared Disk Groups" (page 40).

The procedure discussed in this section applies to the following upgrade scenarios:

- Upgrading from VxVM 3.5 to VxVM 5.1 SP1
- Upgrading from VxVM 3.5 Maintenance Pack 4 to VxVM 5.1 SP1

If there are volumes in the shared disk groups with the FastResync set option (`fastresync=on`), reattach each snapshot to its data volume before beginning the upgrade procedure.

To upgrade in a clustered environment when FastResync is set and to reattach each snapshot to its data volume, complete the following steps from the master node:

1. On the master node, list the disk groups that are shared:

   ```
   # vxdg -s list
   ```

NOTE: To find out if you are on the master node, enter the following command:

```
# vxdctl -c mode
```

2. Using the diskgroup names displayed in the previous command, list the disk groups that have volumes on which FastResync is set:

   ```
   # vxprint -g diskgroup -F "%name" -e "v_fastresync"
   ```

3. Reattach each snapshot:

   ```
   # vxassist -g diskgroup -o nofmr snapback snapshot_volume
   ```

4. If you are upgrading from VxVM 3.5 Maintenance Patch 3 or from VxVM 3.2 Maintenance Patch 5, set FastResync to `off` for each volume:

   ```
   # vxvol -g diskgroup set fastresync=off volume
   ```

Setting Up VxVM 5.1 SP1

This section discusses the procedure to initialize VxVM, add and manage disks under VxVM, and to set up a VEA client and server. This section addresses the following topics:

Initializing VxVM Using the `vxinstall` Utility

You can use the `vxinstall` utility to initialize VxVM. The `vxinstall` initialization procedure enables you to do the following:

- Specify the name of the default disk group to be used by the commands if the `-g` option is not used to specify a disk group explicitly.
- Choose whether to use enclosure-based naming for disks. This naming type enables you to associate more meaningful disk-access names to disks in the different arrays.

To complete a typical `vxinstall` session, complete the following steps:

1. Set up Veritas Volume Manager using `vxinstall`:

   ```
   # vxinstall
   ```

   For more information on supported options, see *vxinstall* (1M).

   The following output is displayed:

   ```
   VxVM uses license keys to control access. If you have not yet installed
   a VxVM license key on your system, you will need to do so if you want to
   use the full functionality of the product.
   Licensing information:
   System host ID:<hostid>
   Host type:<server-id>
   Are you prepared to enter a license key [y,n,q] (default: n) n
   ```

2. To use enclosure-based names, enter **<y|n>** when prompted by the `vxinstall` utility:

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

   For more information on the supported options, see *vxinstall* (1M).

NOTE: Disks use the traditional naming format, usually `c#t#d#`. Enclosure-based naming enables disk devices to be named for enclosures rather than for the controllers through which they are accessed. In a Storage Area Network (SAN) that uses Fibre Channel hubs or fabric switches, information about the disk location provided by the operating system may not correctly indicate the physical location of the disks. Enclosure-based naming enables VxVM to access enclosures as separate physical entities. By configuring redundant copies of your data on separate enclosures, you can safeguard your data against failure of one or more enclosures.

3. To set up a system wide default disk group, enter **<y|n>** when prompted by the `vxinstall` utility:

   ```
   Do you want to setup a system wide default disk group ?
   [y,n,q,?] (default: y)
   ```

   If you know the name of the disk group to be used as the default disk group, enter `y`, and type the name of the disk group at the following prompt, or else use the `list` option and make a selection. Enter **n** if you do not want to define a default disk group.

   NOTE: In releases prior to VxVM 4.1, the default disk group was `rootdg` (the root disk group). For VxVM 3.5 to function, the `rootdg` disk group had to exist, and it had to contain at least one disk. This is no longer required in VxVM 5.1 SP1. However, you can still create a root disk group as a system wide default disk group. For operations that require a disk group, the system wide default disk group is used if the VxVM command does not explicitly mention the disk group with the `-g` option. The main advantage of creating a default disk group is that VxVM commands default to the default disk group, and you do not need to use the `-g` option.

4. Verify the default disk group that is created by entering the following command:

   ```
   # vxdg defaultdg
   ```

   For more information on the supported options, see *vxdg* (1M).

   You cannot use the following names for the default disk group because they are reserved words:

   `bootdg`, `defaultdg` and `nodg`.

5. Define or change the name of the default disk group at a later time by entering the following command:

   ```
   # vxdctl defaultdg diskgroup
   ```

   For more information on the supported options, see *vxdctl* (1M).

This completes the installation of VxVM. You can now use the `vxdiskadm` command and the VEA GUI to create disk groups, and to populate disks.

Moving Disks Under VxVM Control

You can add and manage disks under VxVM by using the menu driven `vxdiskadm` utility.

To add and initialize disks, complete the following steps:

1. Detect new disks, by entering the following command:

   ```
   # ioscan -fnC disk
   ```

   For more information on the supported options, see *ioscan* (1M).

2. Install device files by entering the following command:

   ```
   # insf -C disk
   ```

   For more information on the supported options, see *insf* (1M).

3. For the volume manager to detect these new disks, enter the following command:

   ```
   # vxdisk scandisks
   ```

   For more information on supported options, see *vxdisk* (1M).

4. Invoke the `vxdiskadm` menu driven utility, by entering the following command:

   ```
   # vxdiskadm
   ```

   For more information on the supported options, see *vxdiskadm* (1M).

   Follow the on-screen instructions to complete adding and initializing the disk.

The steps to be followed in moving disks under VxVM control depends on the following:

- If the disk is new, it must be initialized and placed under VxVM control. You can use the menu-based `vxdiskadm` utility to do this as mentioned previously.
- If the disk is not needed immediately, it can be initialized (but not added to a disk group) and reserved for future use. Enter **`none`** when asked to name a disk group.
- If the disk was previously initialized for future use by VxVM, it can be re-initialized and placed under VxVM control.
- If the disk was previously used for a file system, VxVM prompts you to confirm that you really want to destroy the file system.
- If the disk was previously in use by the LVM subsystem, you can preserve the existing data while still letting VxVM take control of the disk. This is accomplished using the conversion utility `vxvmconvert`.
- If the disk was previously in use by the LVM subsystem, and you do not want to preserve the data on it, use the LVM command, `pvremove`, before attempting to initialize the disk for VxVM.

  Multiple disks on one or more controllers can be placed under VxVM control simultaneously. Depending on the scenario, all the disks may not be processed the same way.

For more information on VxVM administration, see the *Veritas Volume Manager 5.1 SP1 Administrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Setting Up a Veritas Enterprise Administrator Server

The VEA server is automatically started at system reboot. However, you can also set up the VEA server later by completing the following procedure.

To set up a VEA server, complete the following steps:

1. Check the status of the VEA service by entering the following command:

   ```
   # /opt/VRTSob/bin/vxsvcctrl status
   ```

   For more information on the supported options, see *vxsvcctrl*(1M).

2. Start the VEA service by entering the following command:

   ```
   # /opt/VRTSob/bin/vxsvcctrl start
   ```

   For more information on the supported options, see *vxsvcctrl*(1M).

   NOTE: To stop the VEA service later, enter the following command:

   ```
   # /opt/VRTSob/bin/vxsvcctrl stop
   ```

Setting Up a Veritas Enterprise Administrator Client

To run VEA and administer a local or remote system, ensure that you have sufficient privileges. Also, ensure that the VxVM and the VEA server are installed on the same system to be administered and that the `vxconfigd` daemon and the VEA server are running.

Following are the commands to administer a local and remote HP-UX system:

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

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

- To administer a remote system, enter the following command:

  ```
  # /opt/VRTSob/bin/vea -host <remote_machine_name> &
  ```

For more information on using VEA, see the *Veritas Enterprise Administrator (VEA) User's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Setting up and Managing VxFS 5.1 SP1

This section discuss the various operations that can be performed on a VxFS file system. It addresses the following topics:

Creating a VxFS File System

To create a file system, enter the following command:

```
# newfs -F vxfs <character_device>
```

For more information on the supported options, see *newfs_vxfs*(1M).

The `newfs` command builds a `vxfs` file system by invoking `mkfs`. You can also use the `mkfs` command, instead of the `newfs` command to create the VxFS file system.

Identifying the Type of File System

To determine the file system type, enter the following command:

```
# /usr/sbin/fstyp -v special
```

For more information on supported options, see *fstyp*(1M).

Converting a File System to VxFS

To convert an HFS file system to a VxFS file system, enter the following command:

```
# /sbin/fs/vxfs/vxfsconvert [-l logsize] [-s size] [-efnNvyY] special
```

For more information on the supported options, see *vxfsconvert*(1M).

Mounting a VxFS File System

To mount a `vxfs` file system, enter the following command:

```
# mount -F vxfs [generic_options] [-r] [-o specific_options] special
<mount_point>
```

For more information on the supported options, see *mount_vxfs*(1M).

Displaying Information on Mounted File System

To view the status of mounted file systems, enter the following command:

```
# mount -v
```

For more information on the supported options, see *mount*(2).

Unmounting a VxFS File System

Enter the `umount` command to unmount a file system, as follows:

```
umount special | mount_point
```

Specify the file system to be unmounted as a value for `mount_point` or for `special` (the device on which the file system resides). For more information on the `umount` command, see *umount*(2).

NOTE: Symbolic links to all VxFS 5.1 SP1 command executables are installed in the `/opt/VRTS/bin` directory. Add this directory to the end of your `PATH` environment variable to access the commands.

If you add an entry for the file system to the `/etc/fstab` file, the file system is mounted automatically.

For more information on VxFS administration, see the *Veritas File System 5.1 SP1 Administrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Setting Environment Variables

This section discusses the path that you must set for the `PATH` environment variable to access relevant files and directories.

While you are installing the Veritas 5.1 SP1 products, you will need to run many commands from the command line. Most of the commands used in the installation are available in the `/sbin` or in `/usr/sbin` directory. You can add these directories to your `PATH` environment variable so that you can access the commands in those directories from the command line. If you install the Veritas man pages, you must set the path of your `MANPATH` environment variable to include the relevant directories.

Cluster Environment Requirements

If you have a cluster environment, you must complete the following steps in addition to setting environment variables discussed in "Setting Environment Variables" (page 45):

1. If you plan to place the root disk group under the control of VxVM, decide into which disk group you want to configure it for each node in the cluster. The root disk group, usually aliased as `bootdg`, contains volumes that are used to start the system. VxVM sets `bootdg` to the

appropriate disk group if it takes control of the root disk. Alternatively, `bootdg` is set to `nodg`. To check the name of the disk group, enter the command:

```
# vxdg bootdg
```

2. Decide on the layout of shared disk groups (as one or more shared disk groups may exist) and determine how many you wish to use.
3. If you plan to use the Dirty Region Logging (DRL) feature with VxVM in a cluster, ensure that some space is available on the disk for these logs. The log size is proportional to the volume size and the number of nodes. For information on DRL, see the *Veritas Volume Manager 5.1 SP1 Administrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.
4. Install the license that supports the clustering feature on every node in the cluster.

5 Upgrading from Previous Versions of VxFS to VxFS 5.1 SP1

This chapter discusses various VxFS upgrade paths based on the version of your OS. This chapter addresses the following topics:

△ **CAUTION:** You must ensure that:

- In the file system upgrade procedures discussed in this chapter, you do not upgrade `/` and `/stand` file systems to DLV 6, DLV 7 or DLV 8. The HP-UX boot loader does not support DLV 6, DLV 7 or DLV 8.

Upgrading from VxFS 3.3 or 3.5 on HP-UX 11i v1 to VxFS 5.1 SP1 on HP-UX 11i v3

To upgrade from VxFS 3.5 on HP-UX 11i v1 to VxFS 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. To determine the volume names, enter the following command:

   ```
   # bdf
   ```

2. To check for DLV 2 or 3 partitions present in the system (because these file systems cannot be mounted after upgrading to VxFS 5.1 SP1), enter the following command:

   ```
   # fstyp -v <partition_name>
   ```

3. To upgrade file systems with DLV 2 or 3, before installing 5.1 SP1, complete the following steps:

   a. To upgrade a mounted VxFS file system from DLV 2 to DLV 3, enter the following command:

      ```
      # vxupgrade -n 3 <dir_name>
      ```

   b. To upgrade a mounted VxFS file system from DLV 3 to DLV 4, enter the following command:

      ```
      # vxupgrade -n 4 <dir_name>
      ```

   NOTE: You cannot upgrade from DLV 2 to DLV 4 directly. DLV4 is the lowest disk layout that is supported by VxFS 5.1 SP1.

4. Upgrade to HP-UX 11i v3. For more information on upgrading to HP-UX 11i v3, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.
5. Install VxFS 5.1 SP1. For more information on installing VxFS 5.1 SP1, see Chapter 3 (page 27).

6. If there are any VxFS file systems with disk layout 2 or 3 or non-VxFS file systems on your system, after upgrading to VxFS 5.1 SP1, that must be converted to VxFS. To upgrade to DLV 7 and convert an HFS file system to VxFS, enter the following command:

   ```
   # opt/VRTS/bin/vxfsconvert <vol_name>
   ```

   The `vxfsconvert` command works with unmounted file systems. To upgrade mounted file system, use the `vxupgrade` command, as described in step 8. You must run `fsck` after `vxfsconvert` because `vxfsconvert` does not create all metadata files.

7. To run the VxFS-specific full `fsck` on the converted file system, enter the following command:

   ```
   # fsck -F vxfs -y -o full <vol_name>
   ```

   During pass 4, `fsck` displays several error messages that require a `yes` response to complete the conversion process.

8. VxFS file systems with disk layout 4 can still be mounted with VxFS 5.1 SP1. However, to upgrade to disk layout versions 5, 6, or 7, use `vxupgrade`, as follows:

   a. To upgrade a mounted VxFS file system from disk layout 4 to disk layout 5, enter the following command:

      ```
      # vxupgrade -n 5 <mount_point>
      ```

   b. To upgrade a mounted VxFS file system from disk layout 4 to disk layout 6, enter the following command:

      ```
      # vxupgrade -n 5 <mount_point>
      # vxupgrade -n 6 <mount_point>
      ```

   c. To upgrade a mounted VxFS file system from disk layout 4 to disk layout 7, enter the following commands:

      ```
      # vxupgrade -n 5 <mount_point>
      # vxupgrade -n 6 <mount_point>
      # vxupgrade -n 7 <mount_point>
      ```

9. To verify the conversion, enter the following command:

   ```
   # fstyp -v <vol_name>
   ```

Upgrading from VxFS 3.5 on HP-UX 11i v2 to VxFS 5.1 SP1 on HP-UX 11i v3

To upgrade from VxFS 3.5 on HP-UX 11i v2 to VxFS 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. To determine the volume names, enter the following command:

   ```
   # bdf
   ```

2. To determine the current disk layout version for system partitions, enter the following command:

   ```
   # fstyp -v <vol_name>
   ```

3. To upgrade file systems with DLV 2 or 3, before installing 5.1 SP1, complete the following steps:

NOTE: You cannot upgrade from DLV 2 to DLV 4 directly. DLV4 is the lowest disk layout that is supported by VxFS 5.1 SP1.

a. To upgrade a mounted VxFS file system from DLV 2 to DLV 3, enter the following command:

```
# vxupgrade -n 3 <dir_name>
```

b. To upgrade a mounted VxFS file system from DLV 3 to DLV 4, enter the following command:

```
# vxupgrade -n 4 <dir_name>
```

4. Upgrade to HP-UX 11i v3. For more information on upgrading, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.
5. Install VxFS 5.1 SP1. For more information on installing VxFS 5.1 SP1, see Chapter 3 (page 27).
6. To convert any non-vxfs file system on your system to a `vxfs` file system, and upgrade the DLV from 2 and 3 to 7, enter the following command:

   ```
   # /opt/VRTS/bin/vxfsconvert <vol_name>
   ```

   The `vxfsconvert` command works with unmounted file systems. To upgrade mounted file system use the `vxupgrade` command as described in step 6, you must run `fsck` after `vxfsconvert` because `vxfsconvert` does not create all metadata files.
7. To run the VxFS-specific full `fsck` on the converted file system, enter the following command:

   ```
   # fsck -F vxfs -y -o full <vol_name>
   ```

   During pass 4, `fsck` displays several error messages that require a `yes` response to complete the conversion process.
8. VxFS file systems with disk layout 4, 5 and 6 can still be mounted with VxFS 5.1 SP1. However, to upgrade the disk layout version of mounted file systems to 7, use the `vxupgrade` command as follows:

   a. To upgrade a mounted VxFS file system from disk layout 4 to disk layout 7, enter the following command:

   ```
   # vxupgrade -n 5 <mount_point>
   # vxupgrade -n 6 <mount_point>
   # vxupgrade -n 7 <mount_point>
   ```

   b. To upgrade a mounted VxFS file system from disk layout 5 to disk layout 7, enter the following command:

   ```
   # vxupgrade -n 6 <mount_point>
   # vxupgrade -n 7 <mount_point>
   ```

   c. To upgrade a mounted VxFS file system from disk layout 6 to disk layout 7, enter the following command:

   ```
   # vxupgrade -n 7 <mount_point>
   ```

   d. To verify the conversion, enter the following command:

   ```
   # fstyp -v <vol_name>
   ```

Upgrading from VxFS 4.1 on HP-UX 11i v2 or HP-UX 11i v3 to VxFS 5.1 SP1 on HP-UX 11i v3

To upgrade from VxFS 4.1 on HP-UX 11i v2 or HP-UX 11i v3 to VxFS 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. To determine the volume names, enter the following command:

   ```
   # bdf
   ```

2. To determine the current disk layout version for system partitions, enter the following command:

   ```
   # fstyp -v <vol_name>
   ```

3. To upgrade file systems with DLV 2 or 3, before installing 5.1 SP1, complete the following steps:

   NOTE: You cannot upgrade from DLV 2 to DLV 4 directly. DLV4 is the lowest disk layout that is supported by VxFS 5.1 SP1.

   a. To upgrade a mounted VxFS file system from DLV 2 to DLV 3, enter the following command:

      ```
      # vxupgrade -n 3 <dir_name>
      ```

   b. To upgrade a mounted VxFS file system from DLV 3 to DLV 4, enter the following command:

      ```
      # vxupgrade -n 4 <dir_name>
      ```

4. Upgrade to HP-UX 11i v3. For more information on upgrading to HP-UX 11i v3, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

   If your system already contains the HP-UX 11i v3 operating system, you can skip this step.

5. Install VxFS 5.1 SP1. For more information on installing VxFS 5.1 SP1, see Chapter 3 (page 27).

6. To convert any non-VxFS file system on your system to a VxFS file system, and to upgrade the DLV from 2 and 3 to 7, enter the following command:

   ```
   # /opt/VRTS/bin/vxfsconvert <vol_name>
   ```

   The `vxfsconvert` command works with unmounted file systems. To upgrade mounted file systems use the `vxupgrade` command as described in step 6. You must run `fsck` after `vxfsconvert`, because `vxfsconvert` does not create all metadata files.

7. To run the VxFS-specific full `fsck` on the converted file system, enter the following command:

   ```
   # fsck -F vxfs -y -o full <vol_name>
   ```

   During pass 4, `fsck` displays several error messages that require a `yes` response to complete the conversion process.

8. VxFS file systems with disk layout 4, 5, and 6 can still be mounted with VxFS 5.1 SP1. However, to upgrade the disk layout version of mounted file systems to 7, use `vxupgrade`, as follows:

   a. To upgrade a mounted VxFS file system from disk layout 4 to disk layout 7, enter the following commands:

      ```
      # vxupgrade -n 5 <mount_point>
      # vxupgrade -n 6 <mount_point>
      # vxupgrade -n 7 <mount_point>
      ```

   b. To upgrade a mounted VxFS file system from disk layout 5 to disk layout 7, enter the following commands:

      ```
      # vxupgrade -n 6 <mount_point>
      ```

```
# vxupgrade -n 7 <mount_point>
```

c. To upgrade a mounted VxFS file system from disk layout 6 to disk layout 7, enter the following command:

```
# vxupgrade -n 7 <mount_point>
```

9. To verify the conversion, enter the following command:

```
# fstyp -v <vol_name>
```

Upgrading from VxFS 5.0 on HP-UX 11i v2 or HP-UX 11i v3 to VxFS 5.1 SP1 on HP-UX 11i v3

To upgrade from VxFS 5.0 on HP-UX 11i v2 or HP-UX 11i v3 to VxFS 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. To determine the volume names, enter the following command:

```
# bdf
```

2. To determine the current disk layout version for system partitions, enter the following command:

```
# fstyp -v <vol_name>
```

3. To upgrade file systems with DLV 2 or 3, before installing 5.1 SP1, complete the following steps:

NOTE: You cannot upgrade from DLV 2 to DLV 4 directly. DLV4 is the lowest disk layout that is supported by VxFS 5.1 SP1.

a. To upgrade a mounted VxFS file system from DLV 2 to DLV 3, enter the following command:

```
# vxupgrade -n 3 <dir_name>
```

b. To upgrade a mounted VxFS file system from DLV 3 to DLV 4, enter the following command:

```
# vxupgrade -n 4 <dir_name>
```

4. Upgrade to HP-UX 11i v3. For more information on upgrading to HP-UX 11i v3 March 2011 OEUR or later, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

   If your system already contains the HP-UX 11i v3 operating system, you can skip this step.

5. Install VxFS 5.1 SP1. For more information on installing VxFS 5.1 SP1, see Chapter 3 (page 27).

6. To convert any non-VxFS file system on your system to a VxFS file system, and to upgrade the DLV from 2 and 3 to 7, enter the following command:

```
# /opt/VRTS/bin/vxfsconvert <vol_name>
```

   The `vxfsconvert` command works with unmounted file systems. To upgrade mounted file systems use the `vxupgrade` command as described in *Step 8*. You must run `fsck` after `vxfsconvert`, because `vxfsconvert` does not create all metadata files.

7. To run the VxFS-specific full `fsck` on the converted file system, enter the following command:

```
# fsck -F vxfs -y -o full <vol_name>
```

   During pass 4, `fsck` displays several error messages that require a `yes` response to complete the conversion process.

8. VxFS file systems with disk layout 4, 5, and 6 can still be mounted with VxFS 5.1 SP1. However, to upgrade the disk layout version of mounted file systems to 7, use the `vxupgrade` command, as follows:

a. To upgrade a mounted VxFS file system from disk layout 4 to disk layout 7, enter the following commands:

```
# vxupgrade -n 5 <mount_point>
# vxupgrade -n 6 <mount_point>
# vxupgrade -n 7 <mount_point>
```

b. To upgrade a mounted VxFS file system from disk layout 5 to disk layout 7, enter the following commands:

```
# vxupgrade -n 6 <mount_point>
# vxupgrade -n 7 <mount_point>
```

c. To upgrade a mounted VxFS file system from disk layout 6 to disk layout 7, enter the following command:

```
# vxupgrade -n 7 <mount_point>
```

9. To verify the conversion, enter the following command:

```
# fstyp -v <vol_name>
```

Upgrading from VxFS 5.0.1 on HP-UX 11i v3 to VxFS 5.1 SP1 on HP-UX 11i v3

To upgrade from VxFS 5.0.1 on HP-UX 11i v3 to VxFS 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. To determine the volume names, enter the following command:

```
# bdf
```

2. To determine the current disk layout version for system partitions, enter the following command:

```
# fstyp -v <vol_name>
```

3. To upgrade file systems with DLV 2 or 3, before installing 5.1 SP1, complete the following steps:

NOTE: You cannot upgrade from DLV 2 to DLV 4 directly. DLV4 is the lowest disk layout that is supported by VxFS 5.1 SP1.

a. To upgrade a mounted VxFS file system from DLV 2 to DLV 3, enter the following command:

```
# vxupgrade -n 3 <dir_name>
```

b. To upgrade a mounted VxFS file system from DLV 3 to DLV 4, enter the following command:

```
# vxupgrade -n 4 <dir_name>
```

4. Install VxFS 5.1 SP1. For more information on installing VxFS 5.1 SP1, see "Installing the Veritas 5.1 SP1 Products" (page 27).
5. To convert any non-VxFS file system on your system to a VxFS file system, and to upgrade the DLV from 2 and 3 to 7, enter the following command:

```
# /opt/VRTS/bin/vxfsconvert <vol_name>
```

The `vxfsconvert` command works with unmounted file systems. To upgrade mounted file systems use the `vxupgrade` command as described in *Step 8*. You must run `fsck` after `vxfsconvert`, because `vxfsconvert` does not create all metadata files.

6. To run the VxFS-specific full `fsck` on the converted file system, enter the following command:

```
# fsck -F vxfs -y -o full <vol_name>
```

During pass 4, `fsck` displays several error messages that require a `yes` response to complete the conversion process.

7. VxFS file systems with disk layout 4, 5, 6 and 7 can still be mounted with VxFS 5.1 SP1. However, to upgrade the disk layout version of mounted file systems to 8, use the `vxupgrade` command, as follows:
   a. To upgrade a mounted VxFS file system from disk layout 4 to disk layout 8, enter the following commands:

      ```
      # vxupgrade -n 5 <mount_point>
      # vxupgrade -n 6 <mount_point>
      # vxupgrade -n 7 <mount_point>
      # vxupgrade -n 8 <mount_point>
      ```

   b. To upgrade a mounted VxFS file system from disk layout 5 to disk layout 8 enter the following commands:

      ```
      # vxupgrade -n 6 <mount_point>
      # vxupgrade -n 7 <mount_point>
      # vxupgrade -n 8 <mount_point>
      ```

   c. To upgrade a mounted VxFS file system from disk layout 6 to disk layout 8 enter the following command:

      ```
      # vxupgrade -n 7 <mount_point>
      # vxupgrade -n 8 <mount_point>
      ```

   d. To upgrade a mounted VxFS file system from disk layout 7o disk layout 8 enter the following command:

      ```
      # vxupgrade -n 8 <mount_point>
      ```

8. To verify the conversion, enter the following command:

   ```
   # fstyp -v <vol_name>
   ```

6 Upgrading from Previous Versions of VxVM to VxVM 5.1 SP1

This section discusses the requirements to update from the pervious version of Veritas Volume Manager to the 5.1 SP1 version of Veritas Volume Manager and also the various supported disk groups. It addresses the following topics:

Determining VxVM Disk Group Version

All Veritas Volume Manager disk groups have an associated version number. 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 work only on disk groups with the current disk group version, so you need to upgrade existing disk groups before you can perform the tasks. When you upgrade from a previous release, you can import disk groups that were created with a supported disk group. For example, VxVM 5.1 SP1 supports all disk group versions that VxVM 4.1 supported (90 and 120). Consequently, all disk groups that were created using VxVM 4.1 can be imported after upgrading to VxVM 5.1 SP1.

The default disk group version for VxVM 5.1 SP1 is version 160. VxVM 5.1 SP1 supports disk groups with certain earlier versions. However, some new features for VxVM 5.1 SP1 require the latest disk group version. VxVM supports shared disk groups only on disk group version 140 or later.

Table 12 summarizes the disk group versions that correspond to each VxVM release on HP-UX from 3.0 to the current version.

Table 12 VxVM Disk Group Version

VxVM Release Number	Disk Group Version	Supported Disk Group Version
3.0	60	20–40, 60
3.1	70	20–70
3.1.1	80	20–80
3.5	90	90
4.1	120	90, 120
5.0	140	90, 120, and 140

Table 12 VxVM Disk Group Version *(continued)*

5.0.1	140	90,120, and 140
5.1 SP1	160	90,120, 140 and 160

Use the following command to determine the version of a disk group:

```
# vxdg list <diskgroup>
```

To upgrade a disk group to version 160, use the following command:

```
# vxdg upgrade <diskgroup>
```

After you upgrade from a previous release of VxVM, HP recommends that you upgrade to the latest disk group version.

For more information on disk groups, see the *Veritas Volume Manager 5.1 SP1 Administrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Native Multipathing with Veritas Volume Manager

By default, VxVM uses DMP to manage multiple paths to I/O devices. It is possible to configure VxVM to use HP-UX native multipathing instead by defining I/O devices as foreign devices, which are bypassed by DMP. To revert to using DMP, you must remove the I/O devices from the foreign device list. This section outlines a system-wide procedure to switch between using DMP and using HP-UX native multipathing using the `vxddladm addforeign` and `vxddladm rmforeign` commands. Using the procedure you can migrate all devices in the `/dev/disk` and `/dev/rdisk` directories. To migrate an individual device or only a subset of devices, see *vxdiskadm3*(1M).

Migrating from DMP to HP-UX native multipathing for systems that do not have VxVM root disks

△ **CAUTION:** Before switching between DMP and HP-UX native multipathing, ensure that no application is accessing the VxVM volumes and no file system is mounted on the VxVM volumes. Otherwise, you cannot complete the following procedure. The following procedure cannot be used on a system that has VxVM root disks.

Prior to migrating from DMP to HP-UX native multipathing, ensure that there are no applications accessing VxVM volumes and no file system is mounted on VxVM volumes.

To migrate from DMP to HP-UX native multipathing, complete the following steps:

1. Stop all volumes in each disk group on the system:

   ```
   # vxvol -g diskgroup stopall
   ```

2. Initiate the migration:

   ```
   # vxddladm addforeign blockdir=/dev/disk chardir=/dev/rdisk
   # vxconfigd -kr reset
   ```

3. Restart the volumes in each group:

   ```
   # vxvol -g diskgroup startall
   ```

The output from the `vxdisk list` command now displays a list of devices following the agile naming scheme of the operating system instead of the DMP metanode names

For example:

```
# vxdisk list
DEVICE     TYPE           DISK         GROUP        STATUS
disk155    auto:LVM         -            -           LVM
disk156    auto:LVM         -            -           LVM
disk224    auto:cdsdisk     -            -          online
disk225    auto:cdsdisk     -            -          online
disk226    auto:cdsdisk     -            -          online
```

```
disk227    auto:hpdisk      -            -           online
disk228    auto:hpdisk      -            -           online
disk229    auto:hpdisk      -            -           online
```

This procedure essentially instructs DMP to bypass all auto-configured devices. As a result, no DMP metanodes are created under `/dev/vx/[r]dmp` directory. VxVM names these devices according to their corresponding agile or persistent device special files (DSFs) assigned by the operating system.

CAUTION: Foreign devices are not supported as cluster-sharable disks. Do not perform this procedure on systems that require clustering features of VxVM (CVM).

NOTE: Irrespective of the mode set, if you use the `vxddladm` or the `vxdisk list` commands, only the new style DSFs are displayed after migrating to HP-UX native multipathing. Enclosure-based naming scheme is not supported for foreign devices.

NOTE: Administrative commands, such as `vxdmpadm iostat` or `vxdmpadm setattr`, which operate on the level of enclosures, controllers or dmpnodes, are not supported and will not work with foreign devices

Migrating from DMP to HP-UX native multipathing for systems that have VxVM root disks

Foreign devices are not supported as cluster-sharable disks. You cannot use this procedure on systems requiring clustering features of VxVM (CVM).

To migrate from DMP to HP-UX native multipathing, complete the following procedure:

1. To migrate to agile addressing mode, enter the following command:

   ```
   # vxddladm set namingscheme=osn mode=new
   ```

2. To initiate the migration, enter the following command:

   ```
   # vxddladm addforeign blockpath=/dev/disk/<root disk>
   charpath=/dev/rdisk/<rootdisk>
   ```

   Following are the message that are displayed:

   ```
   VxVM vxddladm INFO V-5-1-0 Transactions are disabled until the
   reboot.
   VxVM vxddladm INFO V-5-1-0 Root device has been added to foreign
   device category changes will be effective only after reboot
   VxVM vxddladm INFO V-5-1-0 Please reboot the system
   ```

3. Restart the system.
4. Complete the procedure listed in the *Migrating from DMP to HP-UX native multipathing for systems that do not have VxVM root disks* section.

NOTE: Administrative commands which operate on the level of enclosures, controllers or dmpnodes, such as, `vxdmpadm iostat`, `vxdmpadm setattr`, are not supported and will not work correctly on foreign devices.

Migrating from HP-UX native multipathing to DMP for systems that do not have VxVM root disks

Prior to running this procedure, ensure that there are no applications accessing VxVM volumes and no file system is mounted on VxVM volumes. To migrate from HP-UX native multipathing to DMP, complete the following steps:

1. To stop all volumes in each disk group on the system, enter the following command:

   ```
   # vxvol -g diskgroup stopall
   ```

2. To initiate the migration, enter the following commands:

   ```
   # vxddladm rmforeign blockdir=/dev/disk chardir=/dev/rdisk
   ```

```
# vxconfigd -kr reset
```

3. To restart all the volumes in each disk group, enter the following command:

```
# vxvol -g diskgroup startall
```

The output from the `vxdisk list` command shows the DMP metanode names according to the current naming scheme.

Output in the default old naming scheme:

```
# vxdisk list
DEVICE       TYPE          DISK        GROUP        STATUS
c2t0d0     auto:LVM         -            -           LVM
c3t2d0     auto:LVM         -            -           LVM
c89t0d0    auto:cdsdisk     -            -           online
c89t0d1    auto:cdsdisk     -            -           online
c89t0d2    auto:cdsdisk     -            -           online
c89t0d3    auto:hpdisk      -            -           online
c89t0d4    auto:hpdisk      -            -           online
c89t0d5    auto:hpdisk      -            -           online
```

Output in agile naming scheme.

```
# vxdisk list
DEVICE       TYPE          DISK         GROUP        STATUS
disk155     auto:LVM         -            -           LVM
disk156     auto:LVM         -            -           LVM
disk224     auto:cdsdisk     -            -           online
disk225     auto:cdsdisk     -            -           online
disk226     auto:cdsdisk     -            -           online
disk227     auto:hpdisk      -            -           online
disk228     auto:hpdisk      -            -           online
disk229     auto:hpdisk      -            -           online
```

Migrating from HP-UX native multipathing to DMP for systems that have VxVM root disks

Complete the following procedure to switch to DMP from HP-UX native multipathing for systems that have VxVM root disks:

1. To initiate the migration, enter the following commands:

```
# vxddladm rmforeign blockdir=/dev/disk/<root disk>
chardir=/dev/rdisk/<root disk>
```

   The following messages are displayed:

```
VxVM vxddladm INFO V-5-1-0 Transactions are disabled until the reboot
VxVM vxddladm INFO V-5-1-0 Root device has been removed from foreign
device category, changes will be effective only after reboot.
VxVM vxddladm INFO V-5-1-0 Please reboot the system
```

2. Restart the system.
3. To ensure that the remaining disks are migrated after the restarting the machine, complete the steps listed in the *Migrating from HP-UX native multipathing to DMP for systems that do not have VxVM root disks* section.

The output from the `vxdisk list` command shows the DMP metanode names according to the current naming scheme.

Output in the default old naming scheme:

```
# vxdisk list
DEVICE       TYPE          DISK        GROUP        STATUS
c2t0d0     auto:LVM         -            -           LVM
c3t2d0     auto:LVM         -            -           LVM
c89t0d0    auto:cdsdisk     -            -           online
c89t0d1    auto:cdsdisk     -            -           online
c89t0d2    auto:cdsdisk     -            -           online
```

```
c89t0d3     auto:hpdisk      -            -            online
c89t0d4     auto:hpdisk      -            -            online
c89t0d5     auto:hpdisk      -            -            online
```

Output in agile naming scheme.

```
# vxdisk list
DEVICE       TYPE          DISK         GROUP        STATUS
disk155      auto:LVM        -            -           LVM
disk156      auto:LVM        -            -           LVM
disk224      auto:cdsdisk    -            -           online
disk225      auto:cdsdisk    -            -           online
disk226      auto:cdsdisk    -            -           online
disk227      auto:hpdisk     -            -           online
disk228      auto:hpdisk     -            -           online
disk229      auto:hpdisk     -            -           online
```

For information on migrating an individual device or only a subset of devices, see `vxdiskadm(1M)`. For information on foreign devices, see *Veritas Volume Manager 5.1 SP1 Administrator's Guide*. For more information on foreign devices and naming schemes, see the *Migration from Legacy to Agile Naming Model in VxVM 5.0.1 on HP-UX 11i v3* whitepaper. To locate these documents, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Upgrading from VxVM 3.5 on HP-UX 11i v1 to VxVM 5.1 SP1 on HP-UX 11i v3

You must consider the following guideline prior to upgrading from VxVM 3.5 on HP-UX 11i v1 to VxVM 5.1 SP1 on HP-UX 11i v3:

- VxVM 3.5 and VxVM 5.1 SP1 both support disk group version 90.

All upgrades from HP-UX 11i v1 to HP-UX 11i v3 require an intermediate upgrade to HP-UX 11i v2.

If you upgrade to HP-UX 11i v2 May 2005, then the default VxVM version is 3.5. If you upgrade to HP-UX 11i v2 December 2005 and later, the default version is 4.1.

To Upgrade from VxVM 3.5 on HP-UX 11i v1 to VxVM 5.1 SP1 on HP-UX 11i v3

To upgrade from VxVM 3.5 on HP-UX 11i v1 to VxVM 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
2. Upgrade from HP-UX 11i v1 to HP-UX 11i v2. This upgrade will install Base-VXVM 4.1 by default.
3. Upgrade from HP-UX 11i v2 to HP-UX 11i v3. This upgrade will install Base-VXVM 5.0 by default.

   For more information on upgrading to HP-UX 11i v3, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.
4. If patches to VxVM 5.1 SP1 on HP-UX 11i v3 are required, the patches must be applied prior to upgrading the product.
5. Install the VxVM 5.1 SP1 package.

   ```
   # swinstall -x autoreboot=true -s <depot_path> Base-VxVM-51
   ```

Upon successful completion of the upgrade, VxVM 5.0 will be replaced by VxVM 5.1 SP1. Disk groups with older disk group versions (90, 120 or 140) will be accessible after the upgrade.

Upgrading from VxVM 3.5 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3

You must consider the following guideline prior to upgrading from VxVM 3.5 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3:

- VxVM 3.5 and VxVM 5.1 SP1 both support disk group version 90.

To Upgrade from VxVM 3.5 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3

To upgrade from VxVM 3.5 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
2. Upgrade from HP-UX 11i v2 to HP-UX 11i v3. For more information on upgrading to HP-UX 11i v3, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.
3. If patches to VxVM 5.1 SP1 on HP-UX 11i v3 are required, the patches must be applied prior to upgrading the product.
4. Install the VxVM 5.1 SP1 package.

   ```
   # swinstall -x autoreboot=true -s <depot_path> Base-VxVM-51
   ```

Upon successful completion of the upgrade, VxVM 5.0 will be replaced by VxVM 5.1 SP1. Disk groups created in VxVM 3.5 on HP-UX 11i v2 are accessible after the upgrade to VxVM 5.1 SP1.

Upgrading from VxVM 4.1 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3

Following are some of the guidelines that you must consider prior to upgrading from VxVM 4.1 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3:

- Disk group version 120 is the default disk group version for VxVM 4.1, whereas disk group version 140 is the default disk group version for VxVM 5.1 SP1.
- All disk group versions supported by VxVM 4.1 (90, 110, and 120) are also supported by VxVM 5.1 SP1. Consequently, all the disk groups created using VxVM 4.1 can be imported after upgrading to VxVM 5.1 SP1. Some new features and tasks work only on disk groups with the current disk group versions. To use new features, the disk group versions need to be upgraded.
- VxVM 5.1 SP1 on HP-UX 11i v3 supports shared disk groups only on disk group version 140. VxVM 4.1 on HP-UX 11i v3 does not support shared disk groups.

To Upgrade from VxVM 4.1 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3

To upgrade from VxVM 4.1 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
2. Upgrade from HP-UX 11i v2 to HP-UX 11i v3. For more information on upgrading to HP-UX 11i v3, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

   After the OS upgrade is complete, Base-VXVM 5.0 will be the default VxVM software available in the system.

3. If patches to VxVM 5.1 SP1 for HP-UX 11i v3 are required, the patches must be applied prior to upgrading the product.
4. Install the VxVM 5.1 SP1 package.

```
# swinstall -x autoreboot=true -s <depot_path> Base-VxVM-51
```

Upon successful completion of the upgrade, VxVM 5.0 will be replaced by VxVM 5.1 SP1. Disk groups created in VxVM 4.1 on HP-UX 11i v2 will be accessible after you upgrade the system to VxVM 5.1 SP1.

Upgrading from VxVM 4.1 on HP-UX 11i v3 to VxVM 5.1 SP1 on HP-UX 11i v3

Following are the guidelines that you must consider prior to upgrading to VxVM 5.1 SP1 on HP-UX 11i v3:

- Disk group version 120 is the default disk group version for VxVM 4.1, whereas disk group version 160 is the default disk group version for VxVM 5.1 SP1.
- All disk group versions supported by VxVM 4.1 (90 and 120) are also supported by VxVM 5.1 SP1. Consequently, all the disk groups created using VxVM 4.1 can be imported after upgrading to VxVM 5.1 SP1. Some new features and tasks work only on disk groups with the current disk group versions. To use new features, the disk group versions need to be upgraded.
- VxVM 5.1 SP1 on HP-UX 11i v3 supports shared disk groups only on disk group version 140.

To Upgrade from VxVM 4.1 on HP-UX 11i v3 to VxVM 5.1 SP1 on HP-UX 11i v3

To upgrade from VxVM 4.1 on HP-UX 11i v3 to VxVM 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. Stop activity to all VxVM volumes. For example, stop applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
2. If patches to VxVM 5.1 SP1 for HP-UX 11i v3 are required, the patches must be applied prior to upgrading the product.
3. Install the VxVM 5.1 SP1 package.

```
# swinstall -x autoreboot=true Base-VxVM-51
```

After this procedure, VxVM 4.1 is replaced by VxVM 5.1 SP1. Disk groups created in VxVM 4.1 will be accessible after you upgrade to VxVM 5.1 SP1. For information on upgrading a disk group to version 160, see "Upgrading VxVM Disk Group Versions" (page 66).

Upgrading from VxVM 5.0 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3

You must consider the following guideline prior to upgrading to VxVM 5.1 SP1 on HP-UX 11i v3:

- All disk group versions supported by VxVM 5.0 (90, 120, and 140) on HP-UX 11i v2 are also supported by VxVM 5.1 SP1 on HP-UX 11i v3.

As VxVM 5.0 and VxFS 5.0 are bundled in the HP-UX 11i March 2009 OEUR or later release, a fresh installation is not possible. You must perform an upgrade to move to the 5.1 SP1 versions of the Veritas products.

After you install the 5.1 SP1 software, the Veritas Volume Manager 5.0 software that was bundled with the HP-UX operating system will no longer be available. To reinstall Veritas Volume Manager 5.0, you must use your HP-UX software CD/DVD. As only one version of the Veritas Volume Manager can reside on your system, if you choose to reinstall Veritas Volume Manager 5.0, the 5.1 SP1 software will no longer be available.

Installing the 5.1 SP1 Veritas software updates Veritas File System 5.0. After installing the 5.1 SP1 version of Veritas File System, to revert to the 5.0 version, you must remove the 5.1 SP1 version and then re-install the 5.0 version. The 5.0 and 5.1 SP1 versions of Veritas File System cannot coexist.

To Upgrade from VxVM 5.0 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3

To upgrade from VxVM 5.1 SP1 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3, complete the following steps:

1. Stop activity to all VxVM volumes. For example, stop applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
2. Convert the VxVM root device volumes into LVM root device volumes.
3. Upgrade from HP-UX 11i v2 to HP-UX 11i v3. For more information on upgrading to HP-UX 11i v3, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.
4. If patches to VxVM 5.1 SP1 on HP-UX 11i v3 are required, the patches must be applied prior to upgrading the product.
5. Install the VxVm 5.1 SP1 package.

```
# swinstall -x autoreboot=true Base-VxVM-51
```

Upon successful completion of the upgrade, VxVM 5.0 (for HP-UX 11i v2) is replaced by VxVM 5.1 SP1 (for HP-UX 11i v3). Disk groups created in VxVM 5.0 for HP-UX 11i v2 are accessible after the upgrade. If you have converted the VxVM root volumes into LVM root devices, then you must convert them back into VxVM root volumes.

NOTE: During the OS upgrade, the following errors will be reported. However these errors can be ignored and the OS upgrade will complete successfully.

```
ERROR: A later revision (one with a higher revision number) of fileset "VRTSvxvm.VXVM-KRN,r=4.1.032" has already
 been installed.

ERROR: A later revision (one with a higher revision number) of fileset "VRTSvxvm.VXVM-RUN,r=4.1.032" has already
 been installed.
```

Alternatively you can upgrade from VxVM 5.0 on HP-UX 11i v2 to VxVM 5.1 SP1 on HP-UX 11i v3 without removing the VxVM 5.0 package on HP-UX 11i v2.

To upgrade to VxVM 5.1 SP1 on HP-UX 11i v3 without removing VxVM 5.0 on HP-UX 11i v2

To upgrade to VxVM 5.1 SP1 on HP-UX 11i v3 without removing the VxVM 5.0 package on HP-UX 11i v2, complete the following steps:

1. Copy the depot to a depot server.
2. Copy VxVM 5.1 SP1 on HPUX 11i v3 into the ignite depot using the `swcopy` command.
3. Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
4. Remove VxVM 5.0 (Base-VxVM-50 bundle) version which is the default VxVM version on HP-UX 11i v3 from the ignite depot using the `swremove` command.
5. Upgrade from HP-UX 11i v2 to HP-UX 11i v3. For more information on upgrading to HP-UX 11i v3, see the *HP-UX 11i v3 Installation and Update Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

After the OS upgrade is complete, VxVM 5.1 SP1 will be installed. Upon successful completion of the upgrade, disk groups created in VxVM 5.0 for HP-UX 11i v2 will be accessible.

Upgrading From VxVM 5.0 on HP-UX 11i v3 to VxVM 5.1 SP1 Using Integrated VxVM 5.1 SP1 Package for HPUX 11i v3

You can upgrade from VxVM 5.0 HP-UX 11i v3 to VxVM 5.1 SP1 using the integrated VxVM 5.1 SP1 package for HP-UX 11i v3 from the ignite depot.

To upgrade from VxVM 5.0 HP-UX 11i v3 to VxVM 5.1 SP1 on HP-UX 11i v3 using the integrated VxVM 5.1 SP1 package for HP-UX 11i v3 from the ignite depot, complete the following steps:

1. Integrate the VxVM 5.1 SP1 package with the latest HP-UX fusion.

   For information on the procedure to integrate the VxVM 5.1 SP1 package, see the *Installation of VxFS and VxVM 5.1 SP1 on HP-UX 11i v3 March 2011 Operating Environment Update Release* white paper. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

2. Upgrade the HP-UX 11i v3 operating system to the latest fusion using the integrated VxVM 5.1 SP1 package. This installs the VxVM 5.1 SP1 package with the operating system. All disk groups created using VxVM 5.0 on HP-UX 11i v3 will be accessible.

Upgrading From VxVM 5.0.1 on HP-UX 11i v3 to VxVM 5.1 SP1 Using Integrated VxVM 5.1 SP1 Package for HPUX 11i v3

You can upgrade from VxVM 5.0.1 HP-UX 11i v3 to VxVM 5.1 SP1 using the integrated VxVM 5.1 SP1 package for HP-UX 11i v3 from the ignite depot.

To upgrade from VxVM 5.0.1 HP-UX 11i v3 to VxVM 5.1 SP1 on HP-UX 11i v3 using the integrated VxVM 5.1 SP1 package for HP-UX 11i v3 from the ignite depot, complete the following steps:

1. Integrate the VxVM 5.1 SP1 package with the latest HP-UX fusion.

   For information on the procedure to integrate the VxVM 5.1 SP1 package, see the *Installation of VxFS and VxVM 5.1 SP1 on HP-UX 11i v3 March 2011 Operating Environment Update Release* white paper. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

2. Upgrade the HP-UX 11i v3 operating system to the latest fusion using the integrated VxVM 5.1 SP1 package. This installs the VxVM 5.1 SP1 package with the operating system. All disk groups created using VxVM 5.0.1 on HP-UX 11i v3 will be accessible.

I/O Robustness Recommendations

To take full advantage of I/O retry mechanisms in HP-UX and ensure the highest levels of robustness in SAN environments, HP recommends the following steps:

- Disable `VxVM Powerfail Timeout` (PFTO) feature on all disk groups using either of the following commands:

  ```
  vxdisk -g <dg_name> set <disk_name> pftostate = {enabled | disabled}
  vxpfto -g <dg_name> -o pftostate = {enabled | disabled}
  ```

  This essentially allows HP-UX's I/O stack to perform sufficient retries on I/Os and not prematurely timeout the I/O request. Any PFTO setting currently existing is ignored after disabling the PFTO feature. You can use the commands mentioned above to re-enable PFTO, if necessary. For more information, see *pfto(7)*, *vxpfto(1M)*, *vxdisk(1M)* man pages.

- Enable DMP time-based LUN retries using the following command:

  ```
  vxdmpadm settune dmp_lun_retry_timeout=120
  ```

This system-wide tunable instructs DMP to retry failed I/O requests on a LUN for the time in seconds specified. The *vxdmpadm* manpage has not been updated for this release. A brief description of the new parameter, *`dmp_lun_retry_timeout`* is described as follows:

`dmp_lun_retry_timeout`

When all paths to a disk fail, there may be certain paths that have a temporary failure and are likely to be restored soon. If I/Os are not retried for a period of time, the I/Os may be failed to the application layer even though some paths are experiencing a transient failure. For more robust handling of such transient errors, the `dmp_lun_retry_timeout` tunable provides a time-based retry mechanism at the LUN level. If the tunable is set to a non-zero value, I/Os to a disk with all failed paths will be retried until the `dmp_lun_retry_timeout` interval expires or until the I/O succeeds on one of the paths, whichever happens first.

Use the following command to set and retrieve the tunable:

```
vxdmpadm gettune dmp_lun_retry_timeout
```

```
vxdmpadm settune dmp_lun_retry_timeout=<interval>
```

Where the value for <interval> is in seconds.

The default value of this tunable is 0, which means that the paths are probed only once. HP recommends a value of 120 seconds.

You need to set the `pftostate` and `dmp_lun_retry_timeout` tunables only once as the values you set are persistent across reboots.

If these tunables are not set as recommended, certain transient I/O error conditions could result in interruptions to volume availability.

7 Post Upgrade Tasks

This chapter discusses the tasks that you must complete after upgrading the VxVM software. This chapter addresses the following topics:

Optional Configuration Steps

After the upgrade is complete, you can optionally perform the steps discussed in this section based on the requirement, in addition to the steps discussed in "Upgrading From VxVM 5.0 on HP-UX 11i v3 to VxVM 5.1 SP1 Using Integrated VxVM 5.1 SP1 Package for HPUX 11i v3" (page 62).

To perform the optional configuration steps, complete the following procedure:

1. To encapsulate and mirror the boot disk, follow the procedures in the "Administering Disks" chapter of the *Veritas Volume Manager Administrator's Guide*.
2. To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions discussed in "Upgrading VxVM Disk Group Versions" (page 66).

Upgrading Disk Layout Versions

You must upgrade your older disk layout versions to make use of the extended features available in VxFS 5.1 SP1.

For information on the new features, see the *Veritas File System 5.1 SP1 Release Notes*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Use the `vxfsconvert` or `vxupgrade` utilities to upgrade older disk layout versions to disk layout Version 7.

⚠ **WARNING!** Never upgrade the `/` and `/stand` file systems to disk layout Version 7. The HP-UX bootloader does not support disk layout Version 7.

Upgrading VxFS Disk Layout Versions

This section discusses the additional steps required to prepare the file system for being mounted on all nodes of the cluster. Veritas File System 5.1 SP1 supports Versions 5, 6, 7 and 8 for locally mounted file systems and disk layout Versions 6, 7 and 8 for cluster mounted file systems. If you have cluster-mounted file systems with disk layout versions lower than 6, upgrade the disk layout after upgrading to VxFS 5.1 SP1. Disk layout Versions 1, 2, and 3 are not supported by VxFS 5.1 SP1. All file systems created on VxFS 5.1 SP1 use disk layout Version 7 by default. For more information on disk layout versios, see the *Veritas File System 5.1 SP1 Administrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

To upgrade VxFS disk layout versions, complete the following steps:

1. Select one of the nodes of the cluster and mount the file system locally on this node. Use the `mount` command but without the `-o` cluster option. Following is a sample `mount` command:

   ```
   # mount -F vxfs /dev/vx/dsk/sharedg/vol1 /mnt1
   ```

2. Find the current disk layout version on a file system:

   ```
   # fstyp -v | grep version | \ awk '{print $2}'
   ```

3. On the node selected in Step 1, incrementally upgrade the disk layout of this file system to layout Version 6 or 7.

   For example, if you had a cluster mounted file system of disk layout Version 4 running with previous version of VxFS, after upgrading to VxFS 5.1 SP1, you must upgrade the disk layout to Version 6 or 7. The incremental upgrade is as follows:

   ```
   # vxupgrade -n 5 /mnt1
   # vxupgrade -n 6 /mnt1
   # vxupgrade -n 7 /mnt1
   ```

4. On the node selected in Step 1, unmount the file system after the disk layout has been successfully upgraded:

   ```
   # umount /mnt1
   ```

5. You can mount this file system on all nodes of the cluster.

Using the `vxfsconvert` Command

You can use the following `vxfsconvert` command to convert an unmounted HFS file system to a Veritas file system with disk layout Version 7:

```
# vxfsconvert /device_name
```

For more information on the `vxfsconvert` command and `fsadm_vxfs` command, see the *vxfsconvert*(1M) and *fsadm_vxfs*(1M) manpages.

Using the `vxupgrade` Command

You can use the following `vxupgrade` command to upgrade old VxFS disk layouts to disk layout Version 7 while the file system remains mounted::

```
# vxupgrade -n 7 /mount_point
```

For more information on the `vxupgrade` and `fsadm_vxfs` commands, see *vxupgrade*(1M) and *fsadm_vxfs*(1M).

⚠ **WARNING!** The contents of intent logs created on a previous disk layout version cannot be used after the disk layout version is upgraded.

Requirements for Upgrading to Disk Layout Version 7

Converting a previous disk layout to a Version 7 disk layout requires adequate free space. The space and time required to complete the upgrade increases with the number of files, extended attributes, and hard links in the file system. The maximum space required is at least two additional inodes with one block for every inode. A minimum of 10 minutes if required to upgrade every million inodes in the file system.

Upgrading the VxVM Cluster Protocol Version

If you are upgrading a cluster and want to take advantage of the new features in this release, you must upgrade the version of the VxVM cluster protocol. To upgrade the protocol to version 100, enter the following command on the master node of the cluster:

```
# vxdctl upgrade
```

Upgrading VxVM Disk Group Versions

All Veritas Volume Manager disk groups have an associated version number. 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 work only on disk groups with the current disk group version. Before you can perform the tasks, you need to upgrade existing disk groups.

After upgrading from Storage Foundation 4.x to 5.1 SP1, you must upgrade any existing disk groups that are organized by ISP. Without the version upgrade, configuration query operations continue to function properly. However, configuration change operations do not function correctly.

For 5.1 SP1, the Veritas Volume Manager disk group version is the same as it was for the VxVM 5.0 release. Upgrading the disk group version is only required if you upgraded from a version earlier than 5.0.

Run the following command to find the version of a disk group:

```
# vxdg list diskgroup
```

Run the following command to upgrade a disk group to the current disk group version:

```
# vxdg upgrade diskgroup
```

For more information about disk group versions, see the *Veritas Volume Manager 5.1 SP1 Administrator's Guide*. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

Updating Variables

In the `/etc/profile` file, update the `PATH` and `MANPATH` variables as required.

The `MANPATH` variable can include the `/opt/VRTS/man` directory and the `PATH` variable can include the `/opt/VRTS/bin` directory.

Setting the Default Disk Group

In releases prior to Veritas Volume Manager 4.0, the default disk group was the root disk group (`rootdg`). For the Volume Manager to function, the `rootdg` disk group had to exist and it had to contain at least one disk.

This requirement no longer exists; however, you may find it convenient to create a system-wide default disk group. The main benefit of creating a default disk group is that VxVM commands default to the default disk group. You need not use the `-g` option.

You can set the name of the default disk group after installation by running the following command on a system:

```
# vxdctl defaultdg diskgroup
```

For more information on setting default disk groups, see the *Veritas Volume Manager Administrator's Guide*.

If you want to confirm that the root disk is encapsulated, enter the command:

```
# vxdg bootdg
```

Upgrading the Array Support Library

VxVM provides support for new disk arrays in the form of Array Support Library (ASL) software packages.

You can obtain ASL packages from the following locations:

- The VxVM release package
- The disk array provided by the vendor
- The Symantec Technical Support site

Converting from QuickLog to Multi-Volume Support

The 4.1 release of the Veritas File System is the last major release to support QuickLog. The Version 6 or Version 7 disk layout does not support QuickLog. The functionality provided by the Veritas Multi-Volume Support (MVS) feature replaces most of the functionality provided by QuickLog.

This section discusses the procedure to convert from QuickLog to MVS. The QuickLog device need not be related to the file system. For MVS, the log volume and the file system volume must be in the same disk group. Unlike QuickLog, which enabled logging of up to 31 VxFS file systems to one device, MVS enables intent logging of only one file system per device. Therefore, the following procedure must be performed for each file system that is logged to a QuickLog device if Version 6 or Version 7 disk layout is used.

To convert Quicklog to MVS, complete the following steps:

1. Select a QuickLog-enabled file system to convert to MVS and unmount the file system:

   ```
   # umount myfs
   ```

2. Detach one of the QuickLog volumes from the QuickLog device that the file system had been using. This volume can be used as the new intent log volume for the file system.

   ```
   # qlogdetach -g diskgroup log_vol
   ```

3. Create the volume set:

   ```
   # vxvset make myvset myfs_volume
   ```

4. Mount the volume set:

   ```
   # mount -F vxfs /dev/vx/dsk/rootdg/myvset /mnt1
   ```

5. Upgrade the file system of the volume set to Version 6 or Version 7 disk layout. Following is a sample command to upgrade the file system:

   ```
   # vxupgrade -n 6 /mnt1
   ```

   For more information on upgrading disk layout versions, see "Upgrading Disk Layout Versions" (page 64).

6. Add the log volume from step 2 to the volume set:

   ```
   # vxvset addvol myvset log_vol
   ```

7. Add the log volume to the file system. The size of the volume must be specified.

   ```
   # fsvoladm add /mnt1 log_vol 50m
   ```

8. Move the log to the new volume.

   ```
   # fsadm -o logdev=log_vol,logsize=16m /mnt1
   ```

8 Removing Veritas 5.1 SP1 Products

This chapter discusses how to remove the Veritas 5.1 SP1 suite of products from an HP-UX 11i v3 system, using the HP-UX SD commands.

This chapter addresses the following topics:

△ **CAUTION:** If VxVM 5.1 SP1 is installed on a system that contained earlier versions of VxVM previously, the earlier versions of VxVM will be removed from the system as part of the installation. Removing VxVM 5.1 SP1 later leaves the system without any versions of VxVM installed on it. Any existing VxVM volumes will be inaccessible after such removal. If a system relies on the VxVM volumes for reboot, the system may be left unbootable.

NOTE: You must remove the SMO bundles before VxFS and VxVM bundles.

Removing VxVM

VxVM 5.1 SP1 can be removed from systems only if the root disk is under HP Logical Volume Manager (LVM) control. If you try to remove VxVM 5.1 SP1 software from your system without completing the following steps, you will lose data and your system will be in an unusable state.

The following steps provide an overview of the VxVM removal procedure:

- Back-up data on your system.
- Modify file systems and other applications that use volumes to use disks or LVM volumes.
- Remove copies of file systems to free up as much space as possible.
- Remove plexes, except one.
- Striped or spanned volumes must be moved to a single disk or to an appropriate LVM volume.
- Shut down VxVM.

This section addresses the following topics:

Moving VxVM Volumes to LVM Volumes

To move VxVM volumes to LVM, complete the following steps:

1. Evacuate as many disks as possible. Evacuation moves sub-disks from the source disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to LVM volumes. Disks can be evacuated in one of the following ways:
   - Using the `vxdiskadm` command with the **Remove a disk** option.
   - Using the `vxevac` script.

2. Remove the evacuated disks from VxVM control by entering the following command:

   ```
   # vxdg -g <dg> rmdisk <diskname
   # vxdisk rm <devname>
   ```

   For more information on supported options, see *vxdisk* (1M) and *vxdg* (1M).
3. Unmount the disk to be removed first. If the volume is being used as a raw partition for database applications, ensure that the application is not updating the volume and the data on the volume is synchronized.
4. Create an LVM volume of the same size as the plex on 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.
5. Copy the data on the volume to the newly created LVM logical volume by entering the following command:

   ```
   # dd if=</dev/vx/dsk/<logical_vol>> of=</dev/<vol_gr>/<logical_vol>>
   ```

   where **<vol_gr>** is a newly created LVM volume group and **<logical_vol>** is a new logical volume.

   For more information on the supported options, see *dd* (1).
6. Replace the entry (if any) for the volume in `/etc/fstab` file with an entry for the newly created LVM logical volume.
7. Mount the logical volume only if the VxVM volume was mounted before.
8. Remove the volume from Veritas Volume Manager control:

   ```
   # vxedit -g <dg> -rf rm <volume_name>
   ```

   For more information on the supported options, see *vxedit* (1M).
9. Check whether any sub-disks remain on a disk:

   ```
   # vxprint -g <dg> -F "%sdnum" <diskname>
   ```

   If the output is 0, remove the disk from Volume Manager control:

   ```
   # vxdg -g <dg> rmdisk <diskname>
   ```

   For more information on the supported options, see *vxdg* (1M).

   ```
   # vxdisk rm devname
   ```

   For more information on the supported options, see *vxdisk* (1M).

   If the output is not `0`, remove sub-disks from the disk and then remove the disk.
10. Copy the data in the next volume to be removed to the newly created free space.
11. Reboot the system after all volumes are converted successfully.
12. Verify that no open volumes remain on the system:

    ```
    # vxprint -g <dg> -Aht -e <v_open>
    ```

    For more information on the supported options, see *vxprint* (1M).

If any volumes remain open, repeat the above steps.

Removing Plexes

To remove plexes, complete the following steps:

1. Display a list of all the volumes:

   ```
   # vxprint -g <dg> -Ath
   ```

   For more information on the supported options, see*vxprint* (1M).

2. Remove a named plex:

   ```
   # vxplex -g <dg> -o rm disk <plex_name>
   ```

   If the remaining plex has more than one sub-disk, enter the following sequence of commands to consolidate the sub-disks into a single sub-disk:

   ```
   # vxassist -g <dg> mirror volume layout=config
   # vxplex -g <dg> -o rm disk <plex_name>
   ```

   For more information on the supported options, see *vxassist*(1M) and *vxplex*(1M).

NOTE: The consolidation operation fails if you do not have sufficient space on a single sub-disk.

Shutting Down VxVM

To shut down VxVM, enter the following commands:

```
# vxdctl stop
```

For information on the supported options, see *vxdctl* (1M).

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

For information on the supported options, see *vxiod* (1M).

Removing Full VxVM (B9116EB)

You must remove Full VxVM (`B9116EB`) before you can remove Base-VxVM or you must remove them together.

To remove Full VxVM (`B9116EB`), enter the following command:

```
# swremove B9116EB
```

For information on the supported options, see *swremove*(1M).

Removing Base-VxVM-51

To remove Base-VxVM-51, enter the following command:

```
# swremove -x autoreboot=true Base-VxVM-51
```

For more information on the supported options, see *swremove*(1M).

Removing VxFS

This section discusses how to remove the HP Online JFS (B3929HB) and the Base-VxFS-51 products.

Removing HP OnlineJFS (B3929HB)

NOTE: You must remove HP OnlineJFS 5.1 SP1 before removing Base-VxFS-51, or you must remove them together. Removing OnlineJFS 5.1 SP1 does not automatically remove Base-VxFS-51.

You must remove the HP OnlineJFS (`B3929HB`) license before attempting to remove the Base-VxFS-51 bundle.

To remove HP OnlineJFS (`B3929HB`), enter the following command:

```
# swremove B3929HB
```

For more information on the supported options, see *swremove*(1M).

Removing Base-VxFS-51

To remove Base-VxFS-51, enter the following command:

```
# swremove -x autoreboot=true Base-VxFS-51
```

For more information on the supported options, see *swremove*(1M).

△ **CAUTION:** Do not remove the `VRTSvlic` package if there are any other Veritas products running on your system.

NOTE: VxFS 4.1 is activated after the removal of Base-VxFS-51.

Removing CVM

To remove CVM, enter the following command:

```
# swremove -x autoreboot=true B9117EB
```

Removing the Veritas Enterprise Administrator (VEA) Client

You must remove the VEA client software from any system that you used to access the Veritas software.

To remove the VEA client from an HP-UX system other than the server, complete the following steps:

1. Stop the VEA Service:

   ```
   # /opt/VRTS/bin/vxsvcctrl stop
   ```

2. Use the `swremove` command to remove the VEA client packages.

   ```
   # swremove VRTSobgui VRTSat VRTSpbx VRTSicsco
   ```

To remove the VEA client from a Windows system, complete the following steps:

1. Log in as the database administrator.
2. Select **Start > Settings > Control Panel**.
3. Double-click **Add/Remove Programs** to display a list of installed products.
4. Select **VeritasEnterpriseAdministrator** from the list, and click the **Remove** button.
5. Click **Yes** when a dialog box appears asking you to confirm the removal.

A Files Added and Modified After VxFS Installation

This appendix discusses the files that are added and modified after installing VxFS 5.1 SP1.

Files Added After VxFS Installation

Table 13 lists the files added after installing VxFS 5.1 SP1.

Table 13 Files Added After VxFS Installation

File	Description
/usr/conf/lib/libvxfs50.a	Kernel library for VxFS 5.1 SP1
/sbin/lib/mfsconfig.d/vxfs50	VxFS 5.1 SP1 command configuration file
/usr/lib/libxdsm.a	DMAPI library
/usr/lib/libvxfsutil.a	VxFS 5.1 SP1 APIs library
/usr/conf/mod/fdd	Quick I/O module
/sbin/fs/vxfs5.0/ [bcheckrc, extendfs, fsck, fsdb, fstyp, mkfs, mount, newfs, umount, vxdump, vxenablef, vxfsconvert, vxfsstat, vxrestore, vxtunefs,vxupgrade, vxumount]	VxFS 5.1 SP1 commands
/usr/lbin/fs/vxfs5.0/ [df, fcladm, ff, fsadm,fsapadm, fscat, fscdsadm, fscdsconv, fscdstask, fsckpt_restore, fsckptadm, fsenvadm, fsvoladm, getext, ncheck, quot, quotacheck, setext, vxdump, vxlsino, vxrestore]	VxFS 5.1 SP1 commands
/usr/sbin/ [fsclustadm, qiomkfile, qioadmin, qiostat vxfsckd, vxgetmsg]	VxFS 5.1 SP1 commands and Quick I/O commands
/usr/lib/ [vxckptpriv.so, vxfspriv.so, vxfsutil.so, .libvxfsutil.so]	Dynamic libraries for VxFS 5.1 SP1 APIs
/usr/lib/pa20_32/ [vxckptpriv.so, vxfspriv.so, vxfsutil.so, .libvxfsutil.so]	32-bit dynamic libraries for VxFS 5.1 SP1 APIs
/usr/lib/pa20_64/ [vxckptpriv.so, vxfspriv.so, vxfsutil.so, .libvxfsutil.so]	64-bit dynamic libraries for VxFS 5.1 SP1 APIs
/usr/lib/hpux32/ [vxckptpriv.so, vxfspriv.so, vxfsutil.so, .libvxfsutil.so]	32-bit dynamic libraries for VxFS 5.1 SP1 APIs. These are installed only on Integrity systems.
/usr/lib/hpux64 [vxckptpriv.so, vxfspriv.so, vxfsutil.so,.libvxfsutil.s]	64-bit dynamic libraries for VxFS 5.1 SP1 APIs. These are installed only on Integrity systems.
/opt/VRTS/bin/*	Symbolic links to VxFS 5.1 SP1 and Quick I/O commands
/opt/VRTS/vxfs5.0/man/	Manpages for VxFS 5.1 SP1
/opt/VRTS/vxfs5.0/include/	Header files for VxFS 5.1 SP1
/opt/VRTSvxms/lib/map/libvxfs. sl	32-bit VxMS plugin for the Veritas File System
/opt/VRTSvxms/lib/map/pa20_64/ libvxfs.sl	64-bit VxMS plug-in for the Veritas File System

Files Modified After VxFS Installation

Table A-2 lists the files modified after installing VxFS 5.1 SP1.

Table 14 Files Modified After VxFS Installation

File	Modifications
`/stand/system`	Deactivates JFS and OnLineJFS and configures VxFS 5.1 SP1 and Quick I/O.
`/etc/MANPATH`	Inserts `/opt/VRTS/vxfs5.1/man` before the `/usr/share/man` directory in the `MANPATH` environment variable to ensure that the VxFS 5.1 SP1 manpages are displayed, instead of HP OnLineJFS/JFS 5.1 SP1 manpages located in the `/usr/share/man` directory.

Printed in the US