Veritas Storage Foundation™ Release Notes

HP-UX

5.0.1

symantec™

Veritas Storage Foundation™ Release Notes

Product version: 5.0.1

Document version: 5.0.1.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

www.symantec.com/techsupp/

Contacting Technical Support

Customers with a current maintenance agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/assistance_care.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system

- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

https://licensing.symantec.com

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

Customer Service is available to assist with the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and maintenance contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

clustering_docs@symantec.com

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Storage Foundation Release Notes

This document includes the following topics:

- Overview of this release
- Changes in Storage Foundation
- Storage Foundation for Databases supported features
- System requirements
- Component product release notes
- Software limitations
- Fixed issues
- Known issues
- No longer supported
- Documentation
- Documentation errata

Overview of this release

This document provides release information about the products in the Veritas Storage Foundation 5.0.1 product line:

- Veritas Storage Foundation™ (Standard, Standard HA, Enterprise, and Enterprise HA)
- Veritas™ Volume Replicator (VVR)

Each of these products is activated by a single license key. You must obtain a license key before installing the product.

For the latest information on updates, patches, and known issues regarding this release, see the Late-Breaking News TechNote on the Symantec Technical Support website:

http://entsupport.symantec.com/docs/331560

The hardware TechNote is available at:

http://entsupport.symantec.com/docs/283161

Review this entire document before installing your Veritas Storage Foundation product.

About Veritas Operations Services

Veritas Operations Services (VOS) is a Web-based application that is designed specifically for Veritas Storage Foundation and High Availability products. VOS increases operational efficiency and helps improve application availability.

VOS automates and simplifies the following administrator tasks:

- Determining if systems are ready to install or upgrade Storage Foundation and High Availability products
- Gathering deployment and usage information on Storage Foundation and High Availability products
- Receiving notifications about the latest updates for:
  - Patches
  - Hardware Compatibility Lists (HCLs)
  - Array Support Libraries (ASLs)
  - Array Policy Modules (APMs)
- Determining whether your Storage Foundation and High Availability product configurations conform to best practices
- Managing server and environmental configuration data from a single Web site
- Interpreting Unified Message Identifier (UMI) codes and their solutions
- Identifying and downloading patches for Storage Foundation and High Availability products

To access VOS, go to:

http://go.symantec.com/vos

Changes in Storage Foundation

This section describes the changes in Veritas Storage Foundation 5.0.1.

Installation and upgrade

Storage Foundation installation and upgrade includes the following changes in 5.0.1.

Obsolete packages in Storage Foundation 5.0.1

The following packages were included in previous releases of Storage Foundation but are now obsolete:

```
SYMClma
VRTSsmf
VRTScmcm
VRTSjre
VRTSvsvc
VRTSfsdoc
VRTSvmdoc
VRTSvrdoc
VRTSvcsdc
VRTSdbdoc
VRTScsdoc
VRTScfsdc
VRTSxrptl
VRTSdcli
VRTSmh
VRTScsocw
```

Storage Foundation

Storage Foundation includes the following changes in 5.0.1:

Support for HP-UX Integrity Virtual Machines

Veritas Storage Foundation and Veritas Storage Foundation for Oracle are supported with HP-UX Integrity Virtual Machines (HPIVM) in the Guest OS. You can install and configure Veritas Storage Foundation and Veritas Storage Foundation for Oracle in the Guest OS.

See the *Veritas Storage Foundation and High Availability Application Note: Support for HP-UX Integrity Virtual Machines* at the following URL:

http://entsupport.symantec.com/docs/332614

Storage Foundation and Storage Foundation Cluster File System documentation changes

There are changes to the Storage Foundation and Storage Foundation Cluster File System documentation. The *Veritas Storage Foundation Cluster File System Release Notes* (`sfcfs_notes.pdf`) no longer exists. The Storage Foundation Cluster File System Release Notes content exists in the *Veritas Storage Foundation Release Notes* (`sf_notes.pdf`).

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 avoid copying them.

SmartMove provides the following benefits:

- Less I/O is sent through the host, through the storage network and to the disks/LUNs
- Faster plex creation, resulting in faster array migrations
- Ability to migrate from a traditional LUN to a thinly provisioned LUN, removing unused space in the process

In this release, the feature is set to `thinonly` by default. You can also set the feature to `usefssmartmove=all` or to `usefssmartmove=none` in the `/etc/default/vxsf` file. To use the SmartMove feature, VxVM and VxFS must be version 5.0.1 or later.

Thin Provisioning and Storage Reclamation support

Thin Storage is an array vendor solution for allocating storage to applications only when the storage is truly needed, from a pool of free storage. Thin Storage attempts to solve the problem of under utilization of available array capacity.

Thin Storage Reclamation-capable arrays provide ability to reclaim the unused storage space in the file system back to free storage pool. Storage is allocated from the free pool when the data is written to the file system. However, this storage is not released to the free pool automatically when files get deleted. The administrator must run a reclaim operation to reclaim this unused storage back to the free pool.

Storage Foundation supports reclamation of the free blocks in the VxFS file system on VxVM volumes. The reclamation can be done on a LUN, disk group, enclosure or on a mounted VxFS file system, using the `vxdisk` and `fsadm` commands.

Check the hardware compatibility list (HCL) for the supported storage arrays for Thin Provisioning and Storage Reclamation. The HCL Technote is:

http://seer.entsupport.symantec.com/docs/283161.htm

Veritas Volume Manager

Veritas Volume Manager includes the following changes in 5.0.1:

Dynamic Multipathing (DMP) device id limits

Dynamic Multipathing (DMP) supports the identifier of the LUN (SCSI device id) up to a maximum length of 64 characters. Previously, the maximum length was 32 characters. The LUN identifier is a unique SCSI identifier for each storage disk, which consists of the Vendor ID, Product ID, Cabinet Serial Number and LUN Serial Number.

For HP-UX Integrity Virtual Machine (IVM) version 3.5, the total length of the device identifier as supported by DMP in the Guest, is limited to 64 characters. When exporting HPVM Host VxVM volumes as the HPVM Guest backing store, this device identifier includes the VxVM volume path. Hence in this case, the disk group and volume name together cannot be more than 40 characters. When exporting a Host Veritas File System (VxFS) file to the Guest, the length of the full path of the file cannot exceed 54 characters, including slash '/' characters.

Estimated required time displayed during volume conversion

During a volume conversion operation, before the conversion is committed, vxvmconvert displays the estimated time required as follows:

```
VxVM INFO V-5-2-4906
The expected time for convert is: 0 hrs 0 mins 7 secs.
```

Support for iSCSI devices

Veritas Volume Manager now provides support for the use of iSCSI devices in DMP. VxVM provides the same functionality available on Fibre Channel devices for iSCSI devices. This feature is available on all the supported iSCSI arrays. In addition, VxVM provides new interfaces to obtain Fibre Channel and iSCSI configuration information visible to the host. Wherever support from the operating system is available, you can also view and set iSCSI session parameters to improve the performance of iSCSI devices.

Enhancements to the Dynamic Multipathing feature

This release provides a number of enhancements to the Dynamic Multipathing (DMP) features of VxVM. These enhancements simplify administration, and improve display of detailed information about the connected storage.

Dynamic Multipathing path attributes are now persistent

The Dynamic Multipathing (DMP) path attributes, that are set using the `vxdmpadm setattr` command, are now persistent across reboots. These attributes include `standby`, `active`, `preferred`, `primary`, `secondary`, `nopreferred`, and `nomanual`. The attributes are stored in the file `/etc/vx/dmppolicy.info`.

Improved Dynamic Multipathing device naming

The DMP device naming feature has been enhanced to provide a more consistent and user friendly approach for naming the DMP devices.

The following enhancements apply regardless of the specified naming scheme:

- DMP now enables you to assign customized names for DMP devices. You can specify customized names for individual devices, or you can use a file containing user-defined names to assign multiple names.
- You can specify a DMP device name to commands using the name of any of its subpaths. The output displays the DMP device name assigned.
- In a symmetric cluster, the DDL-generated enclosure-based names for DMP devices are now consistent across all the nodes in the cluster.
- Device names can be made persistent. This is the default for the enclosure-based naming (EBN) scheme.

The following enhancements apply to the EBN naming scheme:

- DDL generates the device name in the format *enclosure_index*. If you specify the `use_avid` argument, the name is generated with the Array Volume ID for the index number to provide a more meaningful name.

Default behavior for I/O throttling

By default, DMP is now configured with no I/O throttling. In previous releases, I/O throttling was set to on. Use the `vxdmpadm setattr` command with the `recoveryoption` keyword to configure I/O throttling for DMP.

Specifying a minimum number of active paths

You can now configure a minimum redundancy level, which is the minimum number of paths for the devices under an enclosure. Use the `redundancy` option of the `vxdmpadm getdmpnode` command to display any devices that have fewer

than the minimum number of paths. You can also configure DMP to notify you when the number of active paths falls below the configured minimum.

Enhanced subpaths listing

The `vxdmpadm getsubpaths` command now provides the ability to list all subpaths known to DMP, subpaths of an enclosure, or subpaths through an array port or pwwn. To list the paths through an array port, specify either a combination of enclosure name and array port id, or the array port WWN.

The default listing of the `vxdmpadm getsubpaths` command is sorted by enclosure name, then by DMP node and within that by pathname. The new option `-s` enables you to sort the output based on path name, DMP node name, enclosure name, or host controller name.

Enhanced I/O statistics

The following enhancements have been made to I/O statistics:

Queued and Erroneous I/O counts

The `vxdmpadm iostat show` command now provides options to display queued I/O counts (`-q` option) and erroneous I/O counts (`-e` option). These options are applicable for DMP node, path and controller.

Filter zero entries

The `vxdmpadm iostat show` command now provides the `-z` option to filter out entities for which all data entries are zero. This option is especially useful in a cluster environment, when many paths are required for failover capabilities, but the paths are not being used for I/O.

Specifying units for statistics data

You can now specify the units in which the statistics data is displayed. The `-u` option accepts `k`, `m` and `g` arguments to display throughput in kilo-, mega-, and giga- system blocks. The `us` argument displays average read/write time in microseconds. By default, the read/write times are displayed in milliseconds up to 2 decimal places. The throughput data is displayed in terms of 'BLOCKS' and the output is scaled, meaning that the small values are displayed in small units and the larger values are displayed in bigger units, keeping significant digits constant. The `bytes` argument to `-u` option can be used to display throughput in exact number of bytes.

Cumulative I/O statistics

The `vxdmpadm iostat` command now has a `groupby` clause to provide cumulative I/O statistics listing per dmpnode, controller, array port id, host-array controller pair and enclosure. If the `groupby` clause is not specified then the statistics are displayed per path.

Miscellaneous improvements to DMP I/O statistics

The following improvements have been made to the DMP I/O statistics output:

The way in which average read/write time is calculated has been corrected.

By default, the average read/write time is displayed in milliseconds up to two places after the decimal point. Use the new option `-u us` to display the average read/write time in microseconds.

The average I/O size is set to 512 bytes.

The `vxdmpadm iostat show` command now shows full names for disks with target WWN. If the output exceeds 80 columns, then the output is wrapped.

Making DMP restore options persistent

The restore policy, restore interval, and restore period are now persistent across reboot. In addition to being set as options to the `vxdmpadm start restore` command, these attributes can also be set using the `vxdmpadm settune` command. The new tunables are: `dmp_restore_policy`, `dmp_restore_interval`, and `dmp_restore_cycles`.

In addition, there is a new tunable, `dmp_enable_restore`, which enables the path restoration thread to be started.

New log file location for DMP events

The log file location for DMP events is `/var/adm/vx/dmpevents.log`. For backward compatibility, `/etc/vx/dmpevents.log` is a soft link to `/var/adm/vx/dmpevents.log`.

Extended device attributes displayed in vxdisk list

The `vxdisk list` command now displays extended device attributes like hardware mirrors for certain arrays.

Display use_all_paths attribute for an enclosure

The user can obtain the value of the `use_all_paths` attribute of an enclosure. This attribute indicates whether DMP should issue I/O on both the primary and secondary paths on an Asymmetric Active/Active (A/A-A) LUN or not. For more

information, see the section on scheduling I/O on the paths of an A/A-A array in the *Veritas Volume Manager Administration Guide*.

Viewing information about the ASLs installed on the system

The `/usr/lib/vxvm/diag.d/vxcheckasl` command has been enhanced to provide all the information regarding the ASLs (all those installed in the system), the devices (all seen by OS) and all the possible ways in which these ASLs can interact with these devices.

Displaying the count of LUNs in an enclosure

The `vxdmpadm listenclosure` command now displays the count of LUNs in its default output.

Displaying LUN serial number

The `vxdmpadm getdmpnode` command now includes the option `-v` to display the LUN serial number along with other information.

Displaying HBA details

The `vxdmpadm getctlr` output has been enhanced to display HBA vendor details and the Controller ID. For iSCSI devices, the Controller ID is the IQN or IEEE-format based name. For FC devices, the Controller ID is the WWN. Because the WWN is obtained from Event Source Daemon, this field is blank if the Event Source Daemon is not running.

New exclude and include options for the vxdmpadm command

The `vxdmpadm` command now includes `exclude` and `include` commands to suppress or unsuppress devices from VxVM, respectively.

This provides a command line interface for these operations, which previously required user interaction.

New command for reporting DMP node information

The `vxdmpadm` command now includes the `list` command to display information about a DMP node, including all of the attributes that are set for that DMP node. The `vxdmpadm list` command can be used for a specified dmpnode, all dmpnodes, all dmpnodes on the path name or dmpnodename, or all dmpnodes in an enclosure.

Setting attributes for all enclosures

The `vxdmpadm setattr` command now has the `all` option for enclosure, array type and arrayname. The `all` option allows you to set the attributes (`iopolicy`, `failover_policy`, `recoveryoption`) on all the enclosures specified. Also, `vxdmpadm`

`setattr arraytype` *`array_type`* sets the attribute for all array types derived from the given *array_type*.

Support for ALUA JBOD devices

DDL has now improved the support for JBOD devices to include ALUA JBOD devices. DMP now provides immediate basic support for any ALUA compliant array.

Full support still requires an array support library (ASL) for that array. See the Hardware Compatibility List (HCL) for details about supported arrays.

PFTO is now 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 take more service time to complete an I/O successfully. In case of DMP devices, the use of PFTO is enabled by default.

Changes to site attachment

The `vxsited` daemon has been renamed to 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, Veritas Volume Manager (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. When VxVM detects that the device is online, the VxVM volume components on the involved LUN are automatically recovered, and the mirrors become usable.

The `vxattachd` daemon handles automatic reattachment and resynchronization for plexes. The `vxattachd` daemon also reattaches sites. After a plex is successfully reattached, `vxattachd` notifies root.

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 Cluster Volume Manager (CVM), upon a node crash, the mirror recovery is initiated by the CVM master. Prior to this release, the CVM master also performed all of the recovery I/O. In 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.

When distribution of volume recovery is turned on, the master distributes recovery tasks in a round-robin fashion to other nodes in the cluster. By default, all of the nodes in the cluster can participate in volume recovery. You can also exclude particular nodes from the volume recovery.

To turn on the distribution of volume recovery, add the keyword `distribute` in the file `/etc/default/vxrecover`. You can also specify the `distribute` keyword when running `vxrecover`.

To disable this feature, remove the `distribute` keyword from the file `/etc/default/vxrecover`. If the feature is not enabled, the master performs the resynchronization tasks locally, as in the previous release.

The distributed volume resynchronization functionality does not depend on the disk group version, and hence works with older version disk groups.

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

The following enhancements to `vxdisk` are related to site tagging:

- Site tagging operations on multiple disks or enclosures are now supported.
- New option to rename a site tag on a disk or enclosure.

Automatic site tagging

The `vxdg settag` command now provides an option for automatic tagging of a site. You can specify that an enclosure is automatically tagged with a particular site name. When you add a LUN belonging to that enclosure to a disk group, the LUN is tagged with the site name specified.

Site renaming

The `vxdg` command has a new `renamesite` option. The `renamesite` option renames the existing site record that is configured on the disk group. The `renamesite` option also associates all of the volume objects to the new site.

Veritas File System

Veritas File System includes the following changes in 5.0.1:

Dynamic Storage Tiering enhancements

The Dynamic Storage Tiering (DST) has the following enhancements in this release:

- Dynamic Storage Tiering APIs are enhanced to provide a new interface for managing allocation policies of Storage Checkpoints at creation time and later, and for managing named data stream allocation policies.
- The `fsppadm` command now supports UID, GID, and TAG elements in the placement policy XML file.
- Improved `fsppadm` command scan performance.
- Suppressed the processing of the chosen RULE.
- Parser support for UID, GID, and TAG elements in a DST policy.
- What-if support for `analyze` and `enforce` without requiring the policy to be assigned.
- Storage Checkpoint data placement support in a DST policy.
- Upgrade to SQLite 3.3.9
- 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.

Mount mntlock and mntunlock options

You can specify the `mntlock` option with the `mount` command, which prevents a file system from being unmounted by an application. This option is useful for an application that does not want the file systems that the application is monitoring to be improperly unmounted by other applications or administrators. Clustering applications, such as Veritas Cluster Server (VCS), are particularly expected to benefit.

The `mntunlock` option of the `vxumount` command reverses the `mntlock` option if you previously locked the file system.

Licensing changes for 5.0.1

JFS (Base VxFS) has been enhanced to include Direct I/O (DIO) without a license.

Note: Discovered direct I/O still requires a full feature license.

OnlineJFS now enables Concurrent I/O (CIO) and Extra Big File Systems (EBFS). OnlineJFS is a package that contains a license to enable some additional features over Base VxFS. In this release, OnlineJFS allows file systems greater than 32 TB to be created. Previous releases required the EBFS license for this functionality.

Veritas Storage Foundation Cluster File System

Veritas Storage Foundation Cluster File System includes the following changes in 5.0.1:

Number of parallel fsck threads to run during recovery is tunable

In prior releases the number of parallel fsck threads that could be active during recovery was set to 4. In this release the default depends on the number of CPUs in the system, but is tunable within given limits.

See the *Veritas Storage Foundation Cluster File System Administrator's Guide*.

Quick I/O and Veritas ODM

Quick I/O and Veritas Extension for Oracle Disk Manager (ODM) are now enabled by default for Storage Foundation and Storage Foundation Cluster File System. Existing licenses will enable these features after 5.0.1 is installed or upgraded on the system. This will not require the addition of a new license.

Storage Foundation for Databases supported features

The following sections list the supported features for each platform in the 5.0.1 release of Storage Foundation for Databases.

Support for new database

Veritas Storage Foundation for Oracle provides support for the Oracle™ 11g database in the 5.0.1 release.

The following features of Veritas Storage Foundation for Databases are supported in the Oracle 11g environment:

- Storage Checkpoints
- Storage mapping
- Database cloning (`clonedb`)
- Database Flashsnap
- Tiered storage for databases (DBDST)
- Quick I/O
- Veritas Extension for Oracle Disk Manager (ODM)

Supported features for HP-UX 11iv3 PA and HP-UX 11iv3 IA

For HP-UX 11iv3 PA and HP-UX 11iv3 IA, the following features of Storage Foundation for Databases are supported in this release.

Storage Foundation for Oracle features

Following is a list of supported features for Storage Foundation for Oracle in addition to the existing feature set.

Table 1-1 HP-UX 11iv3 PA and HP-UX 11iv3 IA supported features on Storage Foundation for Oracle Storage Foundation

Oracle version	Storage Foundation for Oracle features
Oracle 11g	■ Tiered storage for databases (DBDST) for RAC ■ Support for Oracle database created on multiple disk groups ■ Flashsnap commands can now be run from the CVM slave node.

System requirements

This section describes the system requirements for this release.

Hardware and software requirements

For information on hardware requirements, see the hardware compatibility list. The hardware compatibility list (HCL) is available at:

http://entsupport.symantec.com/docs/283161

Supported HP-UX operating systems

This release of Veritas products can only be installed on a system running the HP-UX 11i v3 0903 OEUR release 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.11.31.0903   HP-UX Data Center Operating Environment
```

JFS must be installed on your system prior to installing any Veritas software.

To verify that JFS is installed

Use the `swlist` command as follows:

```
# swlist -l product JFS
 JFS              B.11.31         Base VxFS File System 4.1 for HP-UX
```

Database requirements

Table 1-2 identifies supported database and HP-UX combinations if you plan to use Veritas Storage Foundation for Oracle.

Table 1-2 Supported database and HP-UX combinations

Oracle Release	HP-UX 11iv3 0903 OEUR or later
9.2	Yes
10.1	Yes
10.2	Yes
11gr1	Yes

Required HP-UX patches

The 5.0.1 releases of Veritas Storage Foundation, Veritas Storage Foundation for Oracle, and Veritas Storage Foundation for Cluster File System require the following HP-UX patches.

Table 1-3 lists the required HP-UX patches.

Table 1-3 Required HP-UX patches

HP-UX Patch ID	Description
PHSS_36311	This patch fixes a security vulnerability in HP-UX IA-64 platforms. The Veritas Enterprise Administrator Service Core and VRTSobc33 depots require this OS patch on IA-64 platform.
PHKL_40022	This patch distributes vxiod threads to processors other than the moncarch CPU.

Table 1-4 lists the recommended HP-UX patches.

Table 1-4 Recommended HP-UX patches

HP-UX Patch ID	Description
PHKL_39401	This patch fixes a Virtual-Memory defect. This patch should be installed for Veritas File System (VxFS) to respond to memory pressure situations.

HP may release patches that supersede the ones in this list. To verify that you have the latest HP-UX patches, go to the Symantec support website to view the relevant TechNote.

http://www.symantec.com/techsupp

Also, you can get the patches from Hewlett-Packard's Patch Database offered under the Maintenance and Support section of the HP Services & Support - IT Resource Center. HP's Patch Database provides fast, accurate searches for the latest recommended and superseded patches available for Veritas File System or Veritas Volume Manager.

Other required HP-UX software

If you plan to install Storage Foundation from an NFS mounted directory, you must install the software `ONCplus - HP-UX 11i v3 version B.11.31.07.01`. The `ONCplus B.11.31.06` software bundled with HP-UX 11i v3 March 2009 OEUR release reports issues with long path names. This causes the installation to fail

as the installer can not copy files from the mounted directory to the systems on which you want to install Storage Foundation.

To download the software:

- Go to http://software.hp.com.
- Search for the software depot `ONCplus`.
- Download ONCplus for HP-UX 11i v3 version B.11.31.07.01.

Storage Foundation Cluster File System node requirements

All nodes in a Cluster File System must have the same operating system version and update level.

Mandatory patch required for Oracle Bug 4130116

If you are running Oracle versions 9.2.0.6 or 9.2.0.7, you must apply the Oracle patch for Oracle Bug 4130116. Contact Oracle to obtain this patch, and for details on how to apply it.

Component product release notes

In addition to reading these Release Notes, review all component product release notes before installing the product.

Software limitations

The following sections describe the Veritas Storage Foundation software limitations in this release.

Veritas Storage Foundation software limitations

Software limitations in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL.

http://entsupport.symantec.com/docs/283708

5.0.1 Veritas Storage Foundation software limitations

The following are software limitations in this release of Veritas Storage Foundation.

qiomkfile and odmmkfile commands must be run as a privileged user

The `qiomkfile` and `odmmkfile` commands must be run as a privileged user. This behavior has changed from the previous releases where these commands could be run by any user.

Veritas Volume Manager software limitations

Software limitations in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL.

http://entsupport.symantec.com/docs/283708

5.0.1 Veritas Volume Manager software limitations

The following are software limitations in this release of Veritas Volume Manager.

Rootability support for native multipathing

The following are software limitations for native rootability support:

- Migration of a root disk from DMP to native multipathing or vice versa must be followed by a system reboot.
  Online migration of native multipathing to DMP or vice versa is not supported, i.e., the changes will not be effective until the system reboots. VxVM also disables the configuration daemon (`vxconfigd`) until the system reboots. This ensures that any operation, which could lead to a configuration change, would be failed.
- Migration of disks from DMP to native multipathing or vice versa is dependent on the current boot disk. Therefore changes may not be reflected if the system is booted with a different disk.
- Target devices must be under DMP control when executing any of the rootability scripts.
- After performing the `vxrootmir` operation in native multipathing mode, the mirrored disk is not capable of booting until the mirror is broken off from the original root disk.

Rootability not supported for iSCSI devices

In this release, iSCSI devices cannot be used for VxVM rootable disks. HP-UX 11i version 3 does not support iSCSI devices as system root disks, because iSCSI depends on the network stack which is initialized after the boot.

Volume Manager commands take more time with HDS9500-ALUA arrays

Due to slow disk response time, Volume Manager commands take more time with HDS9500-ALUA arrays.

This issue only effects the HDS9500 configuration where, some `vx` commands may take longer to complete compared to 5.0 11.31. This is due to the slow disk response time and the change in default behavior for device open caching.

In 5.0, the device open caching was enabled by default. In 5.0.1, device open caching is set to `off`. You can set the default behavior to enable by using the following command:

```
vxdctl cacheenable
```

You must disable the caching before you make any dynamic configuration change.

Veritas product installer support

The Veritas product installer only supports installing on systems that have the same architecture. For example, if the installer runs on a PA system, all the target systems should be all PA systems. If the installer runs on a PA system and if any of the target systems is an IA, the installer does not proceed properly.

Disk detach policy support in a campus cluster environment

In a campus cluster environment, the disk detach `policy` value for site consistent disk groups should be set to `Global`. The value `policy=local` is not supported.

Veritas File System software limitations

Software limitations in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL.

http://entsupport.symantec.com/docs/283708

5.0.1 Veritas File system software limitations

The following are new additional software limitations in this 5.0.1 release of Veritas File System.

Support of 32 terabyte file systems

Only Veritas Storage Foundation Enterprise and Veritas Storage Foundation Enterprise HA support file systems that are greater than 32 TB.

Quick I/O, ODM, mount -o cio, and the VX_CONCURRENT advisory are mutually exclusive

The `VX_CONCURRENT` advisory cannot be set on a file that is actively open by Quick I/O or ODM. A file that has the `VX_CONCURRENT` advisory set may not be concurrently opened by Quick I/O or ODM. Quick I/O and ODM access are not allowed for any files on a file system that is mounted with the `-o cio` mount option.

Veritas Storage Foundation Cluster File System software limitations

Software limitations in the Veritas Storage Foundation Cluster File System are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes*, which is available at the following URL:

http://entsupport.symantec.com/docs/283716

Veritas Storage Foundation Cluster File System 5.0 MP3 software limitations

This section lists the software limitations in this release of Veritas Storage Foundation Cluster File System.

Compatibility with previous versions of Veritas File System

A disk layout Version 7 file system created with VxFS 5.0 software will not be accessible if the VxFS 5.0 file system software is removed and the system is reverted to VxFS 4.1.

In addition, if a disk layout Version 7 file system exists on a boot disk and VxFS 5.0 is removed, the host will not reboot successfully and will remain at the `bcheckrc` prompt. To reboot the host successfully, first edit the `fstab` file and comment out disk layout Version 7 file system.

Node and host name expansion

HP-UX 11i v3 provides expanded node and host name support. When this feature is enabled, it gives the capability to create node names of up to 64 characters and host names of up to 255 characters. You must enable a dynamic kernel tunable parameter,`expanded_node_host_names`, using the `kctune` command to allow the use of larger names on the system.

When the expanded node name setting is used, Veritas Storage Foundation 5.0.1 products support node names up to 31 characters.

The default operating system configuration is still 8 bytes for node names and 64 bytes for host names. Veritas Storage Foundation 5.0.1 products for HP-UX 11i

v3 support node names up to 7 characters for the default operating system node name configuration of 8 bytes.

Consistent distribution and kernel version for Storage Foundation Cluster File System

All the nodes in a SFCFS cluster must be at the same OS version and patch level. In addition, mixing nodes running 32-bit kernel with nodes running 64-bit kernel is not supported with SFCFS.

Veritas Storage Foundation for Oracle software limitations

Software limitations in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL.

http://entsupport.symantec.com/docs/283708

5.0.1 Veritas Storage Foundation for Oracle software limitations

There are no additional Veritas Storage Foundation for Oracle software limitations in the 5.0.1 release.

DBDST limitations with non-English filenames and placement class names (599164)

DBDST does not work on non-English database filenames or non-English placement class names, due to limitations in VxFS Dynamic Storage Tiering and VxVM volume tags. VxFS Dynamic Storage Tiering does not support placement of non-English filenames. The VxVM volume tag feature does not support non-English volume tag names.

Differing locales produces unintelligible characters in GUI (605487)

The GUI does not support Oracle users having a different locale than the superuser's locale. The GUI will display unintelligible characters if the SFDB repository server starts with a locale that is different from the Oracle user locale (client).

Some features stop working after a Global Cluster failover (563603)

Some Storage Foundation for Oracle features do not work correctly after a Global Cluster (GCO) Failover. In 5.0, the Storage Foundation for Database (SFDB) repository and tools do not manage virtual hostnames correctly in a Global Cluster environment. The SFDB repository does not correctly adjust to the secondary host after the failover.

Features such as Storage Checkpoint, Database FlashSnap, the scheduler, and Database Dynamic Storage Tiering (DBDST) will not function as normal after a failover. However, features such as Oracle Disk Manager (ODM), Quick I/O, and Concurrent I/O (CIO) will continue to work after a failover. This issue will be fixed after the next release.

Avoid using UNIX VEA via PC-based UNIX emulators

There can be problems displaying deep mapping topology in PC-based UNIX emulators such as Exceed. Use the Windows VEA client instead of running the UNIX VEA client via emulators.

No support for Intelligent Storage Provisioning

The Standard, Standard HA, Enterprise, and Enterprise HA versions of Veritas Storage Foundation for Oracle do not support Intelligent Storage Provisioning (ISP).

Disk layouts Version 5 and earlier do not display Storage Checkpoint quotas in the GUI

You can click the Quota tab for the Storage Checkpoint available via the GUI. If the file system and the Storage Checkpoint are on a version prior to Disk Layout Version 6, you will see the error 4646.

To display Storage Checkpoints Quotas via the GUI (for Disk Layout Version 5 and earlier)

1. Using Veritas File System 3.5, create a file system.
2. Upgrade to Veritas File System 5.0 but do not upgrade the file system.
3. Create a checkpoint and then click the Quota tab accessible via the GUI.

Storage Checkpoint limitations (32726, 86431)

This section describes Storage Checkpoint limitations.

The following are Storage Checkpoint limitations:

- You cannot create a clone database using a mounted Storage Checkpoint.
- You must run the dbed_update command after upgrading to Veritas Storage Foundation 4.1 for Oracle from a previous release. This will allow you to roll back to a Storage Checkpoint that was created prior to this release.
- If you create an Oracle instance using the spfile option, you must run the dbed_update command before you can successfully perform any Storage Checkpoint or Database FlashSnap functions.

VEA limitations (34446)

This section describes Veritas Enterprise Administrator (VEA) limitations.

The following are VEA limitations:

- VEA does not display tablespace information when the v$table column names are changed using the SQL*Plus profile facility. Normally this happens when SQL*Plus settings are used in `login.sql` to change column names in reports.
- VEA may display system fonts incorrectly. On a Japanese desktop, VEA may incorrectly display system fonts. Japanese characters may not be properly displayed when you select the non-default font for the VEA GUI.

Database FlashSnap limitations (34570)

This section describes Database FlashSnap limitations.

The following are Database FlashSnap limitations:

- The Database FlashSnap feature does not support RAID-5 volumes.
- When cloning a database using Database FlashSnap, the Oracle database must have at least one mandatory archive destination, otherwise dbed_vmchecksnap results in this error message.[270905]

  ```
  SFORA dbed_vmchecksnap ERROR V-81-5677 Could not find a
  mandatory, primary and valid archive destination for
  database PROD.
  Please review the LOG_ARCHIVE_DEST_n parameters and check
  v$archive_dest.
  ```

  This example shows how to establish a mandatory archive destination using SQL*Plus.

  ```
  alter system set log_archive_dest_1 =
  'LOCATION=/ora_mnt/oracle/oradata/PROD/archivelogs MANDATORY
  [REOPEN]' [scope=both];
  ```

  See your Oracle documentation.
- Existing snapshot plexes created by the vxassist command are not supported. A combination of snapshot plexes created by vxassist and vxsnap is also not supported.

Oracle Disk Manager limitations (34281)

This section describes Oracle Disk Manager limitations.

The following are Oracle Disk Manager limitations:

- Because Oracle Disk Manager uses the Quick I/O driver to perform asynchronous I/O, do not turn off the Quick I/O mount option. The default option is the correct option to use.
- Using Oracle Disk Manager with Cached Quick I/O enabled is not supported and could cause your system to panic

To avoid a system panic, ensure the following:

- If you are using Oracle Disk Manager, do not enable Cached Quick I/O on your file system.
- If you are converting from Quick I/O to Oracle Disk Manager, make sure you disable Cached Quick I/O.

Clone database ORACLE_SID character limit

When cloning an Oracle instance using the dbed_clonedb or dbed_vmclonedb command, the clone database's ORACLE_SID can be only eight characters or less. You will receive an error (ERROR V-81-5713) if the ORACLE_SID is more than eight characters.

Renaming columns in login.sql

Renaming columns in login.sql can sometimes cause scripts to fail or produce incorrect results.

To prevent this, make the following changes in the user environment to generally avoid loading login.sql

1. Move login.sql to another directory, for example, to ~oracle/login.sql.
2. Make sure this new directory is included in SQLPATH, for example:

   ```
   export SQLPATH=~oracle/sql:$SQLPATH
   ```

3. Do not make SQLPATH read-only, so that Storage Foundation for Oracle scripts can unset it at runtime.
4. Avoid starting Storage Foundation for Oracle scripts from the directory where login.sql resides, unless you are sure that login.sql does not contain any settings or commands that change the default output for queries against the data dictionary or increase the startup time for SQL*Plus.

   Also avoid using any settings or commands in the glogin.sql file that change the default output for queries against the data dictionary, or that may increase the startup time for SQL*Plus.

DBDST class names limited to 29 characters (601746)

The `dbdst_admin -o rmclass` command fails when attempting to remove a class name of 30 characters or more. The maximum class name length is 29 characters.

Selected utilities require setuid (643964)

Some Veritas Storage Foundation for Oracle programs are setuid binaries because they are meant to be run as a database administrator and the APIs used are root access-only Symantec internal APIs. The affected binaries are used mainly for information query purposes.

For these reasons, the following programs are setuid-enabled:

- `/opt/VRTSdbed/.dba/dbed_analyzer`
- `/opt/VRTSdbed/.dba/vxckptplan`
- `/opt/VRTSdbcom/bin/vxstorage_stats`
- `/opt/VRTSdbcom/.dba/vxdbd_start`
- `/opt/VRTSdbcom/.dba/vxckpt_ismounted`

Multiple archive log destinations with RAC (795617)

Multiple archive log locations are not supported in RAC configurations.

Repository hostnames are case insensitive (859863)

Because DNS host name lookup queries are by definition case insensitive, make sure the SFDB repository is running on a host with a name that is truly unique -- regardless of case -- within the local subnet. Errors may occur if the repository host name differs from another host name only by case.

One-time scheduled tasks need specific date (861274)

When scheduling a one-time task from the GUI, the task may not be executed if a Specific Date (Include Date) is not set for it.

Veritas Volume Replicator software limitations

Software limitations in the 5.0 release are listed in the *Veritas Volume Replicator 5.0 Release Notes*, which is available at the following URL.

5.0.1 Veritas Volume Replicator software limitations

The following are additional software limitations in the 5.0.1 release of Veritas Volume Replicator.

RAID-5 volume

VVR does not support Volume Manager RAID-5 volumes as part of RVGs. Hardware RAID-5 is supported.

Disk Group Split and Join

VVR does not support Disk Group Split and Join. This means that you cannot use Disk Group Split and Join on data volumes in an RVG. However, you can take snapshots of data volumes and use DGSJ on the snapshots.

Importing a disk group with VVR objects

If a private disk group containing VVR objects is imported on two nodes, as well as the problems documented for VxVM objects, the SRL is corrupted and a full resynchronization of the Secondary is required.

Volumes in boot disk group

Symantec does not recommended having replicated volumes in the boot disk group (`bootdg`).

Selecting records using search expressions

Selecting RVG and RLINK records using search expressions with the `vxprint -e pattern` command is not supported.

Adding a Secondary

When adding a Secondary to the RDS, the `vradmin addsec` command cannot be entered from the host being added. If the RDS contains the Primary only, the command must be entered on the Primary.

Issues related to replication in a shared environment

The following issues relate to replication in a shared environment:

Creating a Primary RVG when the data volumes and SRL are on a shared disk

When creating a Primary whose data volumes and SRL are located on a shared disk group, the `vradmin createpri` command must be issued on the master node of the cluster.

Creating a Secondary RVG when the data volumes and SRL are on a shared disk

When adding a Secondary whose data volumes and SRL are located on a shared disk group to a RDS, the `vradmin addsec` command requires the Secondary host name must be resolvable and up on the master node of the Secondary cluster.

Replication not supported between Solaris and HP-UX for shared disk groups (592349)

Replication in a shared disk group is not supported between Solaris and HP-UX.

VCS Agents for VVR limitations

The following VCS Agents for VVR limitations exist.

Onlining and offlining the RVG resources

Currently a problem with Veritas Volume Manager affects the RVG resources. When many RVGs are defined, the number of requests to the Volume Manager to online or offline can overload the `vxconfigd` process of VxVM, preventing some RVGs from going online or offline.

Workaround

The RVG resources have been set to allow only a single instance of an RVG resource to be brought online or offline. This slows down the onlining and offlining of service groups containing RVGs.

Limitations in the RVGSnapshot agent

Fire drill setup (`fdsetup`) supports applications using one disk group only. The RVGSnapshot agent does not support volume sets.

Fixed issues

The following sections describe Veritas Storage Foundation issues that were fixed in this release.

Veritas Volume Manager fixed issues

The following table describes release 5.0v3 known issues fixed in Veritas Volume Manager release 5.0.1.

Table 1-5 Veritas Volume Manager issues fixed in release 5.0.1

Incident number	Description
592180	Disk group is disabled if private region sizes differ.
1190287	Site failure in an EDC setup with CVM causes application I/O and VxVM commands to hang.
528677	Volume relayout for site-confined or site-consistent volumes.

Table 1-5 Veritas Volume Manager issues fixed in release 5.0.1 *(continued)*

Incident number	Description
533208	Adding a remote mirror to a new site for a site-consistent volume.
536853, 536881	Replacing a failed disk while a site is detached.
540351	Reattaching a site when disks are in the serial split brain condition.
563524	Site records are not propagated during disk group split, move, or join.
1229651	VxVM FlashSnap operation on volume sets with mounted VxFS file systems can cause the loss of changes.
614061, 614787	Cache volumes in volume sets. Do not add cache volumes (used by space-optimized instant snapshots) to volume sets. This causes data corruption and system panics.

Veritas Storage Foundation Cluster File System fixed issues

The following sections list fixed issues for Storage Foundation Cluster File System 5.0 Maintenance Packs. Fixed issues in the Veritas Storage Foundation Cluster File System 5.0 release are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes*, which is available at the following URL:

http://entsupport.symantec.com/docs/283716

Veritas Storage Foundation Cluster File System 5.0.1 fixed issues

Table 1-6 describes fixed issues in the Veritas Storage Foundation Cluster File System 5.0.1 release.

Table 1-6 Veritas Storage Foundation Cluster File System fixed issues

Incident	Description

1375199	Previously, ODM tried to mount in cluster mode by default. ODM tried to register its port with gab. If gab was not configured, the following warning appeared in the syslog:
	
	`GAB WARNING V-15-1-20115 Port d registration failed, GAB not configured vxgms: GAB_API_REGISTER error=261 ODM WARNING V-41-6-5 odm_gms_api_start_msgs fails`
	
	The above message does not appear now. ODM now starts in cluster mode only if GAB is configured; otherwise, ODM starts in local mount mode. If required , enable cluster mode by using the following command after the gab is configured:
	
	`$echo "on" >> /dev/odm/cluster`

Known issues

The following are additional known issues for Veritas Storage Foundation.

Installation known issues

The following are known issues related to the installation process.

Software disk cannot be ejected during installation

During installation, if any of the products that contain VxVM are configured and started, the software disk cannot be ejected. This prevents further use of the disk drive. If a product was installed or upgraded that required a system reboot to complete the installation, this problem is not an issue.

Workaround

If you cannot eject the software disk, do the following.

To eject the software disk

1. Use the `fuser`(1M) command to identify processes using the disk device. Note the process IDs. If `fuser` cannot identify the processes, use the `ps`(1M) command.
2. Identify the process ID for the `swagentd` process. Enter the following command:

   ```
   # ps -aef | grep swagentd
   ```

3. Kill the `swagentd` process. Use the `kill -9` command with the process ID you identified in step 2. Enter the following command:

   ```
   # kill -9 swagentd_process_ID
   ```

4. Eject the software disk.
5. Restart the `swagentd` process. Enter the following command:

   ```
   # /usr/sbin/swagentd -r
   ```

If only the required packages are installed, installing the `VRTSvxfs`, `VRTSdbed`, and `VRTSodm` packages fail on the second node of the cluster (1852746)

This issue is caused by the dependencies of corequisite packages. The `VRTSfsman` package is a corequisite package of the `VRTSvxfs` package. However, `VRTSfsman` is an optional package. For example, if you install Veritas Storage Foundation for Oracle RAC (SFRAC) and choose to only install the required packages, `VRTSfsman` is not installed on the remote system before `VRTSvxfs`, and errors occur during `VRTSvxfs` installation. The installer does not copy the corequisite packages to the remote systems.

Further, the `VRTSdbed` and `VRTSodm` packages depend on `VRTSvxfs` being installed properly, so errors occur during those installations as well.

However, if you choose to install all product packages, including the optional ones, `VRTSfsman` is installed before `VRTSvxfs` and no errors occur.

Note: This issue was first noticed during SFRAC installation; however, it applies to all Storage Foundation products.

Workaround

To avoid this issue with any of the Storage Foundation products, install all product packages, both required and optional.

Configuration fails if the `VAILAgent` process fails to start (1853906)

The product installer checks to see if the `VAILAgent` process is running. If the process is not running, the configuration cannot continue and it terminates with an error. This is a timing issue between when the `VAILAgent` process starts and the product installer makes its check.

Workaround

To workaround this issue, do the following.

To start the `VAILAgent` process and configure the CFS cluster

1. Add the missing registry keys to the registry. Enter the following command. If you enter the command on a single line, omit the `\` continuation character at the end of the first line.

   ```
   # /opt/VRTSvail/bin/vaildiag \
   /JOINDOMAIN -agentname VAILAgent -standalonemode
   ```

   The `vaildiag` command displays the following error, which can be ignored:

   ```
   ERROR: Authentication Broker value not available in the
   /etc/default/csf_resolv.conf file
   Please ignore this message if the installation was done in
   Host controller mode configuration.
   ```

2. Start the `VAILAgent` process. Enter the following command:

   ```
   # /opt/VRTSobc/pal33/bin/vxpal -a VAILAgent
   ```

3. Configure the CFS cluster. From the first node of the cluster, enter the following command:

   ```
   # /opt/VRTS/bin/cfscluster config -t 200 -s
   ```

Veritas Storage Foundation known issues

Known issues in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL:

http://entsupport.symantec.com/docs/283708

Veritas Storage Foundation 5.0.1 known issues

The following are additional known issues in the 5.0.1 release of Veritas Storage Foundation.

AVXTOOL error messages reported during installation or upgrade of 5.0.1 (1741126)

During installation and upgrade of 5.0.1 Storage Foundation Products, following warning messages appear in the /var/adm/sw/swremove.log file for AVXTOOL fileset.

```
The fileset "AVXVM.AVXVMPKG,l=/,r=B.05.00.01" requires the
         selected fileset "AVXTOOL.AVXTOOLPKG,l=/,r=B.05.00.01" as a
         corequisite.
    WARNING: The dependencies for fileset
         "AVXTOOL.AVXTOOLPKG,l=/,r=B.05.00.01" cannot be resolved (see
         previous lines).
```

Cause: Both AVXVM and AVXTOOL filesets are uninstalled by the installer. The two filesets have dependencies on each other. Therefore, there is no way to avoid the warning message, when either of these two filesets are uninstalled.

Workaround:

This message can be safely ignored.

`VAILAgent` process does not start after configuring Storage Foundation (1597587)

When the Storage Foundation product is configured, the `VAILAgent` process does not start.

Workaround

Use the following procedure to start the `VAILAgent` process.

To start the `VAILAgent` process

1 Add the missing registry keys to the registry. If you enter the first command on a single line, omit the `\` continuation character at the end of the first line.

```
# /opt/VRTSvail/bin/vaildiag \
/JOINDOMAIN -agentname VAILAgent -standalonemode
```

The `vaildiag` command displays the following error, which can be ignored:

```
ERROR: Authentication Broker value not available in the
/etc/default/csf_resolv.conf file
Please ignore this message if the installation was done in
Host controller mode configuration.
```

2 Start the `VAILAgent` process. Enter the following command:

```
# /opt/VRTSobc/pal33/bin/vxpal -a VAILAgent
```

Installer error during installation on IVM host system with local IVM guest OS (1765419)

The following installer error may occur during installation on IVM host system with local IVM guest OS:

```
CPI ERROR V-9-0-0 thread 2 threw error: Can't call
method "pid" on an undefined value at
authentication_service/scripts/CPI/common/misc.pl line 494.
```

This message indicates that the installer failed because too many processes are running on the system. The installer failed to create additional multi-threaded installation processes.

Workaround:

Rerun the installer and specify the `-serial` option:

```
# ./installer -serial
```

Startup or shutdown failure messages reported for LLT, GAB, and VXFEN

If you need to reboot the system when you install Storage Foundation, the init scripts for LLT, GAB, and VXFEN report start or stop failure messages. This is because Storage Foundation is not yet configured and the required configuration files are not yet generated for these components. These messages may be ignored. [1666327]

This issue also applies when you install Storage Foundation Cluster File System.

VRTSgab error messages reported during upgrade from Storage Foundation 4.1 MP2

When you upgrade Storage Foundation from version 4.1 MP2 (HP-UX 11i v2) to version 5.0.1 (HP-UX 11i v3) using the Veritas installation programs, the installer reports an error message indicating failed uninstallation of GAB (Global Atomic Broadcast). The error message is logged in the installation log files and the corresponding swremove log files. [1719136]

The error message reported in the installation log file is as follows:

```
The following depots failed to uninstall on galaxy:
VRTSgab
```

The message may be ignored as the VRTSgab package for version 5.0.1 is successfully installed.

VRTSpbx error messages reported during upgrade from Storage Foundation 4.1 MP2 (1738121)

When you upgrade Storage Foundation from version 4.1 MP2 on HP-UX 11i v2 to version 5.0.1 HP-UX 11i v3 using the Veritas installation programs, a warning message is logged in the installation log files and the corresponding swremove log files. The message may be ignored as the VRTSpbx package for version 5.0.1 is successfully installed.

VRTSat error messages reported during upgrade from Storage Foundation 4.1 MP2 (1738097)

When you upgrade Storage Foundation from version 4.1 MP2 (HP-UX 11i v2) to version 5.0.1 (HP-UX 11i v3) using the Veritas installation programs, The warning message is logged in the installation log files and the corresponding swremove log files. The message may be ignored as the VRTSat package for version 5.0.1 is successfully installed.

Storage agent gives warning and fails to discover objects in VEA GUI (1588508)

If you are going to use the VEA GUI, you must restart the Storage agents after you complete the configuration of Veritas Storage Foundation products.

Database repository daemon fails to start after upgrade (1706956)

In some cases, the `vxdbd` daemon fails to start after you upgrade Storage Foundation products from a previous version. This issue may apply to Storage Foundation, Storage Foundation High Availability, Storage Foundation for Oracle, or Storage Foundation for Oracle High Availability.

Work around

Restart the `vxdbd` daemon using the following command:

```
# /opt/VRTSdbcom/bin/vxdbdctrl start
```

Warning messages observed during upgrade (1725491)

The Veritas product installer displays warning messages during the Veritas Volume Manager (VxVM) upgrade to 5.0.1. These messages are displayed during upgrade of any Storage Foundation product that includes VxVM.

The HP-UX 11iv3 includes the AVXVM package that HP OEM bundle, Base-VxVM-50 has a wrapper pkg named AVXVM. The AVXVM pkg is an empty pkg but it has a dependency on VRTSvxvm pkg. These warning messages can be safely ignored.

While configuring Storage Foundation Management Server and the Cluster Management Console through the CPI for authentication passwords, some special characters are not passed correctly through the CPI (1245237)

While configuring Storage Foundation Management Server and the Cluster Management Console through the Common Package Interface (CPI) for authentication passwords, some special characters are not correctly passed through the CPI to the nodes, even though these special characters are accepted by authentication.

The following special characters are not correctly passed through the CPI to the nodes:

- \' (single quote)
- \" (double quote)
- \@ (at sign)
- \$ (dollar)
- \\ (slash)
- * (star)

Workaround

There is no workaround for this issue. When entering authentication passwords, do not use any of the special characters listed above.

Product installer does not support upgrading to a different Storage Foundation product and upgrading the product version in a single process

The Veritas product installer does not support upgrading to a different product in the Storage Foundation product line and upgrading the product version in a

single process. For example, you cannot upgrade directly from Storage Foundation 4.0 to Storage Foundation Cluster File System 5.0.1. First, you must upgrade the Storage Foundation version from 4.0 to 5.0.1. Then, upgrade the product from Storage Foundation 5.0.1 to Storage Foundation Cluster File System 5.0.1.

Uninstalling the VRTSmapro package (617740)

Uninstallating the `VRTSmapro` (mapping provider) package does not remove the corresponding entry from the VEA registry.

Workaround

Run the following command before uninstalling the `VRTSmapro` package:

```
# /opt/VRTSmapro/bin/vxmapping_prov.config -r
```

DBMS security issue

The following are DBMS security issues.

The Symantec Shared DBMS feature creates the following configuration files:

- `/etc/vxdbms/VERITAS_DBMS3_hostname/conf/databases.conf`
- `/etc/vxdbms/VERITAS_DBMS3_hostname/conf/databases1.conf`
- `/etc/vxdbms/VERITAS_DBMS3_hostname/conf/registration.dat`

These configuration files are created or modified by `vxdbms_start_db.pl`, `vxdbms_start-server.pl`, and `vxdbms_register.pl` respectively.

The files are writable by everyone if the file mode creation mask is insufficiently restrictive. Symantec strongly recommends that you restrict the default file mode creation permissions (using the `umask` command) for `root` and `administrator` accounts to avoid a potential security issue. Specifically, change the `group|world` write and execute permissions in the default umask. At its least restrictive, the default umask for `root` should be `022`. Symantec recommends setting to `077`.

Host name may need to be entered manually on clustered host

When installing SF Management Server on a clustered host on which the Domain Name Service (DNS) is not configured, you will be prompted to enter the fully-qualified host name manually to proceed with the installation.

When installing on a clustered host on which the DNS is running slowly, you may occasionally be prompted to enter the fully-qualified host name manually to proceed with the installation.

Unconfigure VEA Action Agent after a MANAGED host installation (616057)

When installing in MANAGED mode you must execute the following script at any time after the Veritas packages have been installed:

```
/opt/VRTSaa/config/remove_vxaa.sh
```

This script unconfigures the VEA Action Agent and prevents it from starting on a MANAGED host installation.

Warning and error messages appear during the Storage Foundation 5.0.1 installation or upgrade (1836854)

During Storage Foundation 5.0.1 stack installation, warning messages are displayed for the `VRTSvxfs` and `VRTSfsman` packages. swverify errors are displayed for `AONLINEJFS.AOLINEJFSPKG`.

Note: These messages only occur with HP-UX 11i v3 0909 OEUR. You can ignore them.

Veritas Volume Manager known issues

Known issues in the 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL:

http://entsupport.symantec.com/docs/283708

This release, 5.0.1, fixes some of the 5.0v3 known issues that are documented in the link above.

See "Veritas Volume Manager fixed issues" on page 33.

Veritas Volume Manager 5.0.1 known issues

The following are new additional known issues in this 5.0.1 release of Veritas Storage Foundation.

Incorrect paths for vxvm-iscsi script in HP /sbin/init.d/iscsi startup script (1656156)

The HP-UX operating system uses the `/sbin/init.d/iscsi` startup script to enable iSCSI services. During system startup with iSCSI configuration, the `/sbin/init.d/iscsi` script invokes the `vxvm-iscsi` script. The `vxvm-iscsci` script is required to perform appropriate VxVM discovery of iSCSI devices, import diskgroups on those devices, and start the volumes. In HP-UX 11iv3 0903 OEUR, the `/sbin/init.d/iscsi` script has the incorrect path for the `vxvm-iscsi` script.

This issue leads to iSCSI diskgroups not getting automatically discovered and imported.

Workaround:

Correct the `/sbin/init.d/iscsi` script by changing all occurences of `/sbin/init.d/vxvm-iscsi` to `/sbin/vxvm-iscsi`.

Issues with persistent disk access records (1725501)

If you initialize a disk using the following commands, Veritas Volume Manager creates persistent disk access records.

```
# vxdisk rm disk_access_name
# vxdisk -f init disk_access_name

# vxdisk rm disk_access_name
# vxdisk define disk_access_name
```

The recommended method to initialize a disk for use by VxVM is to use the `vxdisksetup` command. This method does not create persistent disk access records.

To initialize a disk for regular use:

```
# /etc/vx/bin/vxdisksetup -I disk_access_name
```

In some scenarios, issues can occur with persistent disk access records. If a disk group is created using the persistent disk access records, these persistent disks may enter an error state when the naming scheme changes. Some hardware reconfigurations may also cause the persistent disks to be in an error state.

Workaround

If the persistent disks are in an error state due to a change in the naming scheme or hardware reconfiguration, you can remove the error state using one of the following procedures. The procedure depends on whether the host is in a cluster or is a stand-alone host.

To remove the error state for persistent disks in a disk group for which host is not part of a cluster

1. If some volumes are in use, unmount them.
2. Enter the following command to restart the VxVM configuration daemon:

   ```
   # vxconfigd -kr reset
   ```

To remove the error state for persistent disks in a disk group for which the host is part of a cluster

1. If some volumes are in use, unmount them.
2. Deport the disk group.

```
# vxdg deport diskgroup
```

3. Rescan the devices.

```
# vxdisk scandisks
```

4. Import the disk group.

```
# vxdg import diskgroup
```

I/O failures result in the disk failing flag (1205380)

In some DMP failover scenarios, I/O retry causes the disk failing flag to be set, although there is nothing wrong with the disks except for the failing flag.

Workaround

Clear the failing flag using the `vxedit` command.

Upgrading external ASLs and APMs

If your system uses any existing ASLs and APMs from older releases, you may need to download the latest version from Symantec. Check the latest array support information to determine whether an updated ASL or APM package is available for your arrays.

See the 5.0.1 Hardware Compatibility List for information about supported arrays.

Before upgrading a Storage Foundation product, you must manually remove any external ASL or APM packages from older releases which are not supported. After completing the upgrade, obtain the required updated ASLs or APMs to ensure the array is claimed correctly.

Disabling HP_UX native ALUA for arrays operating in ALUA mode (1822365)

For Storage Foundation 5.0, HP-UX 11iv3, and for the storage arrays operating in Asymmetric Logical Unit Access (ALUA) mode, you must disable Native Multi-Pathing ALUA (nMP) mode to ensure that DMP functions properly in the Storage Foundation stack.

This can be done in the following ways:

- By disabling nMP mode at the LUN Level
- By disabling nMP mode at the `esdisk` driver level

Procedures for each method are shown below.

Disabling Native Multi-Pathing ALUA (nMP) mode at the LUN level

This method sets the `alua_enabled` attribute for each LUN and its multiple paths to `false`. Additionally, the `alua_enabled` attribute must be made persistent by setting `save_attr`.

The `scsimgr`(1M) command is used to enable or disable ALUA mode for a LUN.

To disable Native Multi-Pathing ALUA (nMP) at the LUN level

1 Determine if the attribute has been already set to disabled; that is, `alua_enabled = false`. The following command shows the `scsimgr` options to display the `alua_enabled`attribute and its persistence.

```
# scsimgr get_attr -D  /dev/rdisk/disk460 -a alua_enabled

  SCSI ATTRIBUTES FOR LUN : /dev/rdisk/disk460

  name = alua_enabled
  current = true
  default = true
  saved =
```

2 Disable nMP ALUA for a LUN. Use the `scsimgr` command to set the attribute `alua_enabled = 0`. Enter the following command:

```
# scsimgr set_attr -D  /dev/rdisk/disk460 -a alua_enabled=0

  Value of attribute alua_enabled set successfully
```

3 Verify that the attribute is set correctly. Re-enter the `scsimgr`command.

```
# scsimgr get_attr -D  /dev/rdisk/disk460 -a alua_enabled

  SCSI ATTRIBUTES FOR LUN : /dev/rdisk/disk460

  name = alua_enabled
  current = false
  default = true
  saved =
```

Note: In the output above, the `current` state has been changed to `false`, indicating that the native ALUA attribute for this particular LUN path has been disabled.

4 Make this attribute persistent across host reboots. Enter the `scsimgr` command with the `save_attr` option.

```
# scsimgr save_attr -D /dev/rdisk/disk460 -a alua_enabled=0

  Value of attribute alua_enabled set successfully
```

Note: You do not have to reboot the host to make these changes effective. You should not run I/O when making these changes.

5 Validate that both the `name` and `saved` attributes are set correctly.

```
# scsimgr get_attr -D  /dev/rdisk/disk460 -a alua_enabled

  SCSI ATTRIBUTES FOR LUN : /dev/rdisk/disk460

  name = alua_enabled
  current = false
  default = true
  saved = false
```

6 Display the `alua_enabled` attribute for all LUNs in a server.

```
# scsimgr -p get_attr all_lun -a device_file -a alua_enabled
```

Note: You should execute the steps above for all agile DSFs in the `/dev/rdisk/` directory. If new LUN(s) or DSF(s) are added to the host at runtime on the HP-UX host (SF is already configured and then new LUN(s) are added), then the steps above must be executed separately for the newly added LUN(s).

Disabling Native Multi-Pathing ALUA (nMP) mode at the `esdisk` driver level

This method sets the `alua_enabled` attribute at the `esdisk` driver level, thereby disabling or enabling ALUA based on attribute value for all LUNs bound to the `esdisk` driver.

The following procedure shows how to set, save, and display the current and default settings for the `alua_enabled` attribute and also how to disable the ALUA persistently at the driver level.

To disable Native Multi-Pathing ALUA (nMP) mode at the `esdisk` driver level

1 Disable ALUA at the `esdisk` driver. Enter the following command:

```
# scsimgr  set_attr -N  "/escsi/esdisk" -a alua_enabled = 0

  Value of attribute alua_enabled set successfully
```

2 Make this attribute persistent across host reboots. Save it by using the following command:

```
# scsimgr save_attr -N "/escsi/esdisk" -a alua_enabled=0

  Value of attribute alua_enabled saved successfully
```

Note: You do not have to reboot the host to make these changes effective. You should not run I/O when making these changes.

3 Use `get_attr` to check or display the attribute changes. Enter the following command:

```
# scsimgr get_attr -N "/escsi/esdisk" -a alua_enabled

  SCSI ATTRIBUTES FOR SETTABLE ATTRIBUTE SCOPE : "/escsi/esdisk"

  name = alua_enabled
  current = false
  default = true
  saved = false
```

4 Check the `alua_enable` attributes for `disk460`. Enter the following command:

```
# scsimgr get_attr -D /dev/rdisk/disk460 -a alua_enabled

  SCSI ATTRIBUTES FOR LUN : /dev/rdisk/disk460

  name = alua_enabled
  current = false
  default = false
  saved =
```

For more information on `scsimgr`, see the `scsimgr`(1M) man pages.

Veritas Volume Manager cluster issues

The following are cluster issues in this release of Veritas Volume Manager.

Cluster Volume Manager (CVM) behavior when the disk group failure policy is set to leave (1510854)

If the master node loses access to all copies of the logs, the behavior depends on the disk group failure policy.. If the disk group failure policy is set to leave, the master node panics so that another node that has access to the disk group can become the master. If the detach policy is set to global, the master node panics immediately. However, if the detach policy is set to local, the master node does not panic immediately. The panic occurs later, when an event occurs that requires an update to the kernel log. For example, after all slave I/O stops.

Handling intermittently failing paths in a Campus Cluster

In remote mirror configurations, a site is reattached when its disks come back online. Recovery is then initiated for the plexes of a volume that are configured at that site. Depending on the configuration, recovery of the plexes can take a considerable time and consume considerable resources. To minimize the frequency

of having to perform a site reattachment operation, it is recommended that you use the `vxdmpadm settune` command to configure a value smaller than 60 seconds for `dmp_health_time`, and a value larger than 300 seconds for `dmp_path_age`.

Domain controller mode in CVM clusters (603213)

The slave nodes in a CVM cluster only have access to I/O objects. If non-I/O related information (for example, volume tags) are to be made available on a slave node, a command must to be shipped to the Storage Agent on the master node for execution. The results are then communicated back to the slave node.

The domain controller mode of VEA allows all nodes of a CVM cluster to be placed in the same domain with a central authentication server. This allows commands to be executed on any node within the domain if the executing process has sufficient rights.

Provided domain controller mode is configured, non-I/O related information is accessible via VEA on any node in a CVM cluster.

However, even if domain controller mode is enabled in a CVM cluster, ISP commands must be run on the master node. ISP commands that are run on a slave node are not redirected to the Storage Agent on the master node. Such commands fail if they require access to non-I/O related information that is unavailable on a slave node.

Volume recovery can be aborted due to an active joining node, leaving the volume in the NEEDSYNC/SYNC state (1824240)

Volumes needing recovery go into the NEEDSYNC state. Volume recovery moves the volume from the NEEDSYNC to ACTIVE state. Sometimes this operation fails because a node joins the reconfiguration being processed at the same time, leaving the volume state in NEEDSYNC/SYNC state.

Symptom

The volume under recovery is in the NEEDSYNC/SYNC state and `vxtask list` does not show any active tasks.

Workaround

Manually recover volumes in NEEDSYNC/SYNC state. On the Master Node, use the following command to move volumes from NEEDSYNC/SYNC state to ACTIVE state:

vxvol -g *dg* -f resync *vol*

In certain cases, when the master node is rebooted, mirrored volumes detach and I/O failures occur (1828706)

If the VxVM disk group is configured on iSCSI devices, and if the CVM master node is rebooted (ungraceful shutdown) with I/Os in progress on the VxVM mirrored volumes from the same node, the mirrored volumes are detached on all the slave nodes in the cluster and I/O failures occur. This behavior occurs when the detach `policy` value is configured as `Global`.

Node join issues when the disk group is in a disabled state on the master node (1835578)

In a CVM environment, a shared disk group goes into a disabled state on the master node when it loses connectivity to all the configuration and logs copies of the disk group. Subsequent node joins fail.

Workarounds

You can correct this issue with one of the following workarounds:

- Resolve the storage connectivity issue, stop I/Os on volumes in this disk group from all slave nodes, and deport and re-import the shared disk group. The subsequent slave join will succeed.
- If the node join is needed before the connectivity problem is resolved, the master role will need to be failed over to a node that has connectivity to all shared disk groups. The subsequent slave join (with the new master) will succeed.

Deport operation on a shared disk group fails (1368377)

With all primary paths inaccessible, the deport operation on a shared disk group fails to clear the PGR keys as the DMP database is not up-to-date. The deport operation succeeds but the PGR keys are not cleared as the DMP database is not updated to reflect the inaccessibility of failed primary paths.

Workaround

Running `vxdisk scandisks` before the DG deport operation triggers DMP reconfiguration which updates the DMP database such that a disk is accessible through active paths.

Veritas File System known issues

The Veritas File System 5.0.1 known issues are listed by release in this section. Known issues in the Veritas File System 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes* which is available at the following URL:

http://entsupport.symantec.com/docs/283708

Veritas File System 5.0.1 known issues

The following are additional known issues in this 5.0.1 release of Veritas Storage Foundation.

File Change Log tunable setting for proper functioning of Dynamic Storage Tiering applications

If the active placement policy of a given file system uses I/O or access temperatures, after the policy becomes active by being assigned, you must tune the file system's *fcl_malloc* tunable with the following command:

```
# vxtunefs -o fcl_maxalloc=0 mount_point
```

However, if any applications other than DST use FCL, this setting may conflict with those applications.

Veritas Storage Foundation Cluster File System known issues

Veritas Storage Foundation Cluster File System known issues in the 5.0 release are listed in the *Veritas Storage Foundation Cluster File System 5.0 Release Notes*.

The *Veritas Storage Foundation Cluster File System 5.0 Release Notes* can be viewed at the following URL:

http://entsupport.symantec.com/docs/283716

Veritas Storage Foundation Cluster File System 5.0.1 known issues

The following sections include new additional Veritas Storage Foundation Cluster File System known issues in this release.

The following issue also applies to Veritas Storage Foundation Cluster File System:

See "Startup or shutdown failure messages reported for LLT, GAB, and VXFEN" on page 39.

Accessing the files through ODM may not work (1842547)

Accessing the files through ODM in serially-exclusive mode does not work, ODM access is only available for files that are not simultaneously accessed from multiple nodes.

Workaround

Currently no workaround exists for this issue. This will be fixed in a subsequent release of the SFCFS product.

Oracle-HP (840486)

Problems uninstalling or upgrading Veritas Storage Foundation for Oracle when Veritas Storage Foundation Cluster File System is installed on the same system.

If Veritas Storage Foundation for Oracle and Veritas Storage Foundation Cluster File System are installed on the same machine, do not use the installer to uninstall if you are planning to uninstall only one product.

If you want to uninstall the product, you must uninstall the Veritas Storage Foundation for Oracle packages manually.

To uninstall the Veritas Storage Foundation for Oracle packages

1. Review the uninstallation requirements in the *Veritas Storage Foundation Installation Guide*.
2. Stop the repository database and unmount the repository volume.

In a stand-alone configuration:	Stop the database repository:
	`# /opt/VRTSdbcom/bin/sfua_db_config -o stopdb`
	Unmount the database repository:
	`# /sbin/init.d/sfua_rep_mount stop`

In an HA configuration:	Stop VCS processes on either the local system or all systems.
	To stop VCS processes on the local system:
	`# hastop -local`
	To stop VCS processes on all systems:
	`# hastop -all`

3. Remove the Veritas Storage Foundation for Oracle packages using the `swremove` command.

```
# swremove VRTSorgui VRTSdbed VRTSdbcom VRTSdbdoc
```

If Veritas Storage Foundation for Oracle and Veritas Storage Foundation Cluster File System are installed on the same machine and you are upgrading both products, use the installer to upgrade Veritas Storage Foundation Cluster File System first. Then, use the installer to upgrade Veritas Storage Foundation for Oracle.

If the second upgrade fails, remove the Veritas Storage Foundation for Oracle packages as described above, then run the installer to upgrade Veritas Storage Foundation for Oracle.

CFSMount resource may fault

During cluster startup in a SFCFS for Oracle RAC or SFCFS environment, a CFSMount resource may fault even though the underlying CVMVolDg resource becomes online successfully. If the CVMVolDg resource contains layered VxVM volumes, the reason for the fault could be that the CVMVolDg resource went online before all the subvolumes of the layered volume could be started.

Workaround

In order to ensure that a CVMVolDg resource containing layered volumes becomes online only after all the subvolumes are enabled, the CVMVolume attribute in the `main.cf` file should be populated with the names of the layered volumes under that CVMVolDg resource.

See the *Veritas Cluster Server User's Guide*.

Incorrect directory permissions after an upgrade from SFCFS 4.1 or 5.0 releases on HP-UX 11i v2 (1733255)

In Storage Foundation Cluster File System (SFCFS) 4.1 or 5.0 releases on HP-UX 11i v2, the permissions for the `/var/adm/cfs` directory were set to mode 766. In this release, the permissions are set to 755. Therefore, after an upgrade to SFCFS 5.0.1, the following warning displays:

```
WARNING: Directory "/var/adm/cfs" should have mode "755" but the actual
         mode is "766".
WARNING: Fileset "VRTScavf.VXFS-CFS-ADM-CMDS,l=/,r=5.0.31.5.%20090531"
         had file warnings.
```

You can safely ignore this warning, or you can manually change the permissions to correct the issue.

To change the permissions, use the following command:

```
# chmod 755 /var/adm/cfs
```

Product installer does not support configuring I/O fencing in enabled mode (1593808)

For this release, the product installer does not support configuring I/O fencing in enabled mode. If you use the product installer to configure SFCFS, you must answer "n" to the following prompt:

```
Will you be configuring I/O fencing in
enabled mode? (y,n, q) n
```

If you select "y", the VRTSvxfen process fails. This failure can be ignored. At the end of the configuration, you can enable the I/O fencing manually.

For details about enabling I/O fencing, see the *Storage Foundation Cluster File System Administrator's Guide*.

Messages reported during 5.0.1 upgrade (1837574)

During the upgrade of Storage Foundation 5.0.1 products, if you are using HP-UX 11i v3 0909 OEUR, the following messages are displayed.

Note: In each of these cases, you may ignore the messages.

The following warning messages appear for VRTSvlic.

```
      * The fileset "OnlineJFS01.VXFS50-AD-RN,l=/,r=5.0" requires the
        selected fileset "VRTSvlic.VLIC-RUN,l=/,r=3.02.24.0" as a
        corequisite.
WARNING: The dependencies for fileset
        "VRTSvlic.VLIC-RUN,l=/,r=3.02.24.0" cannot be resolved (see
        previous lines).
```

The following messages appear while removing the VRTScutil.VRTScutil fileset.

```
rmdir: /opt/VRTSvcs/hasnap/fw/lib/service/repositoryclient/cvsclient:
No such file or directory
rmdir: /opt/VRTSvcs/hasnap/fw/lib/service/repositoryclient:
No such file or directory
rmdir: /opt/VRTSvcs/hasnap/fw/lib/service: No such file or directory
rmdir: /opt/VRTSvcs/hasnap/fw/lib/project/ant: No such file or directory
rmdir: /opt/VRTSvcs/hasnap/fw/lib/core: No such file or directory
rmdir: /opt/VRTSvcs/bin: Directory not empty
rmdir: /opt/VRTSvcs: Directory not empty
```

The following messages appear while removing the VRTSvxfen.VRTSVXFEN-KRN fileset.

```
rmdir: /opt/VRTSvcs/vxfen/newconfig/etc/rc.config.d: Directory not empty
rmdir: /opt/VRTSvcs/vxfen/newconfig/etc: Directory not empty
rmdir: /opt/VRTSvcs/vxfen/newconfig: Directory not empty
rmdir: /opt/VRTS/messages/en: Directory not empty
rmdir: /opt/VRTS/messages: Directory not empty
rmdir: /opt/VRTSvcs/vxfen/bin/customized/script: No such file or directory
rmdir: /opt/VRTSvcs/vxfen/bin/customized/sanvm: No such file or directory
rmdir: /opt/VRTSvcs/vxfen/bin/customized/netapp: No such file or directory
rmdir: /opt/VRTSvcs/vxfen/bin: Directory not empty
rmdir: /opt/VRTSvcs/vxfen: Directory not empty
```

Veritas Volume Replicator known issues

Known issues in the Veritas Volume Replicator 5.0 release are listed in the *Veritas Volume Replicator 5.0 Release Notes*, which is available at the following URL:

http://entsupport.symantec.com/docs/283745

Veritas Volume Replicator 5.0.1 known issues

The following are new additional known issues in this 5.0.1 release of Veritas Volume Replicator.

Mirrors are not synchronized when volumes created using init=active option (1287111)

For volumes created using `init=active` option, the mirrors (plexes) are not synchronized. The `vradmin verifydata` command could incorrectly report differences for such volumes. To rectify this situation, synchronize mirrors (plexes) and resynchronize the secondary by doing Automatic Synchronization, Full Synchronization, or Difference-based Synchronization.

Issue with VVR VEA in the Japanese locale (616709)

In the Japanese locale, the Add Bunker wizard page has truncated text. When you add a bunker using VVR VEA, the description text for the Bunker DG and Protocol fields is truncated.

The incomplete text should read as follows:

- Bunker DG: If protocol is Storage the Bunker DG is expected to have been imported on the Primary host.
- Protocol: Protocol should be set to Storage when Bunker storage is directly accessible from the Primary host.

Veritas Storage Foundation for Oracle known issues

Known issues in the Veritas Storage Foundation for Oracle 5.0 release are listed in the *Veritas Storage Foundation 5.0 Release Notes*, which is available at the following URL.

http://entsupport.symantec.com/docs/283708

Veritas Storage Foundation for Oracle 5.0.1 known issues

The following new known issues exist in this release of Veritas Storage Foundation for Oracle.

Checkpoint Clonedb gives "unexpected error" message on the GUI (1703134)

When creating a checkpoint clone, you may see the following error:

```
Error V-40-49408-1
  An unexpected error occurred
```

In spite of the above error, the clone operation has completed in the background. A rescan in the GUI displays the appropriate objects.

Workaround:

Edit the `/etc/vx/isis/Registry` as follows:

In the section`SOFTWARE\VERITAS\VRTSobc\pal33\Agents\__defaults`, add the following key:

```
[REG_INT] "USE_RT_TIMEOUT" = 0;
```

dbed_clonedb -o restartdb fails after database group switches to a second node with spfile (1721965)

In Storage Foundation for Oracle High Availability, the `restartdb` command for a checkpoint clone database can be run only on the node on which you have created the clone.

Reverse Resync not supported if database is created using Oracle Managed Files (1192729)

If an Oracle database is created with Oracle Managed Files (OMF), then `reverse_resync` operations would fail.

The following errors are displayed:

```
oracle@swlx07:~> /opt/VRTSdbed/bin/dbed_vmsnap -S $ORACLE_SID \
-f sp4 -o

reverse_resync_begin
dbed_vmsnap started at 2007-12-28 12:02:42

SFORA dbed_vmsnap WARNING V-81-5725 After reverse_resync_commit
is performed, you need to recreate the Authentication Password
File using the ORAPWD utility.
SFORA dbed_vmsnap ERROR V-81-4882 An error occurred while
reconfiguring Oracle instance 'sfora'.
SFORA dbed_vmsnap ERROR V-81-4881 Log file is at
/tmp/dbed_vmclonedb.12313/nomount.log.
SFORA dbed_vmsnap ERROR V-81-4918 Database sfora has not been
correctly recovered.
SFORA dbed_vmsnap ERROR V-81-4881 Log file is at
/tmp/dbed_vmclonedb.12313/recover.log.
```

Workaround

The `reverse_resync` operation for a database created with OMF is not supported in the 5.0.1 release.

There is no workaround for this issue.

Checkpoint Rollback corrupts Oracle database when upgrading from 9i to 10g(1121064)

If an Oracle 9i database is upgraded to 10g, and at a later point in time the database is rolled back to the 9i environment, then the recovery procedure for 9i will give the following error in the alert log file:

```
ORA-00600: internal error code, arguments: [2662]..."
```

This bug has been documented by Oracle as bug number 2216823.

Workaround

Do not use the pre-existing `tempfile`. Instead you may either back up the `tempfile` with rest of the database or remove the `tempfile` and then re-create a new `tempfile` once the database is open.

Veritas DBMS cannot distinguish between hosts with long hostnames (1802583)

If two hostnames are more than 8 characters, but the first 8 characters are identical, the VxDBMS cannot distinguish between them.

DBMS server cannot be started if the `base_pagesize` is 64 (1832119)

The DBMS server cannot be started if the `base_pagesize` is set to 64. This can happen while installing or upgrading to 5.0.1 in the IVM host, where the base_pagesize is set to a value of 64 by default. The following error messages are printed in the swm and swagent logs:

```
------------------
Allocating a communication buffer failed

Ping server failed -- Connection error: Found server but communication erro
occurred.

ERROR:   *** Could not verify DBMS server instance. Please check error.
ERROR:   *** Return code: 1
ERROR:   *** Initialization has failed!
ERROR:   The "configure" script for "VRTSdbms3" failed (exit code "1").
         The script location was
```

```
        "/var/adm/sw/products/VRTSdbms3/pfiles/configure".
-----------------
```

Workaround

To start the DBMS server successfully, use `kctune`(1M) to set `base_pagesize` to a value less than 64.

Offline checkpoint clonedb shows checksum error in alert log (1545717)

If a clone database operation is performed for an offline checkpoint, then the `dbed_clonedb` command executes successfully and the database is cloned properly. However, the alert log for the clone database shows the following errors:

```
Errors in file /oracle/11g/diag/rdbms/cl6/cl6/trace/cl6_ora_21133.trc:
ORA-00367: checksum error in log file header
ORA-00318: log 1 of thread 1, expected file size 1024 doesn't match 1024
ORA-00312: online log 1 thread 1: '/tmp/cl6/datavol/FLAS11r1/redo01.log'
```

These errors are harmless and can be ignored.

Veritas Storage Foundation Graphical User Interface known issues

Cannot convert mountpoint to volume set using the Java Graphical User Interface (1176531)

When you use the Java Graphical User interface to convert a mountpoint into a volume set, the operation is not successful.

Workaround

You must use the `dbdst_convert` command line interface to convert a mountpoint volume to a volume set. You must not use the Java Graphical User Interface to convert volumes.

Adding two storage classes consecutively through VEA Java Graphical User Interface fails (1231856)

When you use the VEA Java Graphical User Interface to add the first class, the operation is successful. However, when you use the VEA Java Graphical User Interface to add a second class, you may see the following error:

```
SFORA dbdst_admin V-81-6212
Do not add or remove class in a single command.
```

This issue is encountered only with Graphical User Interface and not with the command line interface.

Workaround

You must logout and disconnect from VEA. Then you must start a new session to add a second class through the Graphical User Interface.

Alternatively, you may use the `dbdst_admin` command in the command line interface to add a new class:

```
# /opt/VRTS/bin/dbdst_admin -D DB4 -o addclass=NEWCLASS:"newclass"
```

Some disks may not appear in the VEA GUI (1826607)

Whenever the VxVM disk naming scheme is modified, the VEA GUI needs to recognize the new naming scheme.

Disks or volumes may not appear in the VEA GUI (1727003)

After you configure any Storage Foundation product, some disks or volumes may not appear in the VEA GUI.

Workaround

Restart the Storage Agent using the following command:

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a StorageAgent -c restart
```

No longer supported

This section describes Veritas Storage Foundation features that will not be supported in future releases. Symantec advises customers to minimize the use of these features.

The following features will not be supported in the next release of Veritas Storage Foundation products:

- The `seconly` feature in SFCFS
- Storage Expert
- The use of the `vxvoladm` command line utility
- Intelligent Storage Provisioning (ISP)

Documentation

Relevant component product release notes

Read the relevant component product release notes before installing any version of Veritas Storage Foundation.

The following documents are relevant component product release notes:

- *Veritas Cluster Server Release Notes* (`vcs_notes.pdf`)
- *Veritas Storage Foundation for Oracle RAC Release Notes* (`sfrac_notes.pdf`)

Storage Foundation guides

The following manuals, along with the online help, comprise the Veritas Storage Foundation documentation set:

Table 1-7 describes the guides in the Veritas Storage Foundation documentation set.

Table 1-7 Guides in Veritas Storage Foundation documentation set

Guide Title	Filename
Third-party Legal Notices	`3rdpartyattributions.pdf`
Veritas FlashSnap Point-In-Time Copy Solutions Administrator's Guide	`flashsnap_admin.pdf`
Veritas Storage Foundation and High Availability Getting Started Guide	`getting_started.pdf`
Read me end user license agreement documentation	`README_EULA`
Veritas Storage Foundation read me first	`readme_first.txt`
Veritas Storage Foundation Release Notes (this document)	`sf_notes.pdf`
Veritas Storage Foundation Installation Guide	`sf_install.pdf`
Veritas Storage Foundation for Oracle Administrator's Guide	`sf_ora_admin.pdf`
Veritas Storage Foundation for Oracle Graphical User Interface Guide	`sf_ora_gui.pdf`

Table 1-7 Guides in Veritas Storage Foundation documentation set *(continued)*

Guide Title	Filename
Veritas Storage Foundation Intelligent Storage Provisioning Administrator's Guide	`sf_isp_admin.pdf`
Veritas Storage Foundation Intelligent Storage Provisioning Solutions Guide	`sf_isp_solutions.pdf`
Veritas Storage Foundation Cross-Platform Data Sharing Administrator's Guide	`sf_cds_admin.pdf`
Veritas Enterprise Administrator User's Guide	`vea_users.pdf`
Veritas File System Administrator's Guide	`vxfs_admin.pdf`
Veritas File System Programmer's Reference Guide	`vxfs_ref.pdf`
Veritas Volume Manager Administrator's Guide	`vxvm_admin.pdf`
Veritas Volume Manager Migration Guide	`vxvm_migration.pdf`
Veritas Volume Manager Troubleshooting Guide	`vxvm_tshoot.pdf`

Veritas Storage Foundation Cluster File System documentation

The following Veritas Storage Foundation Cluster File System documentation is available with all Veritas Storage Foundation HA product suites:

Table 1-8 describes the Veritas Storage Foundation Cluster File System (CFS) documentation set.

Table 1-8 Guides in Veritas Storage Foundation Cluster File System documentation set

Guide Title	Filename
Veritas Storage Foundation Cluster File System Installation Guide	`sfcfs_install.pdf`
Veritas Storage Foundation Cluster File System Administrator's Guide	`sfcfs_admin.pdf`

Veritas Cluster Server documentation

The following Veritas Cluster Server documentation is available with all Veritas Storage Foundation HA product suites:

Table 1-9 describes the Veritas Cluster Server documentation set.

Table 1-9 Guides in Veritas Cluster Server documentation set

Guide Title	Filename
Veritas Cluster Server Release Notes	vcs_notes.pdf
Veritas Cluster Server Installation Guide	vcs_install.pdf
Veritas Cluster Server User's Guide	vcs_users.pdf
Veritas Cluster Server Agent Developer's Guide	vcs_agent_dev.pdf
Veritas Cluster Server Bundled Agents Reference Guide	vcs_bundled_agents.pdf
VCS Enterprise Agent for Oracle Installation and Configuration Guide	vcs_oracle_install.pdf
VCS Enterprise Agent for DB2 Installation and Configuration Guide	vcs_db2_install.pdf
VCS Enterprise Agent for Sybase Installation and Configuration Guide	vcs_sybase_install.pdf

Veritas Volume Replicator documentation

The following Veritas Volume Replicator documentation is available with the Veritas Volume Replicator option:

Table 1-10 describes the Veritas Volume Replicator documentation set.

Table 1-10 Guides in Veritas Volume Replicator documentation set

Guide Title	Filename
Veritas Volume Replicator Administrator's Guide	vvr_admin.pdf
Veritas Volume Replicator Planning and Tuning Guide	vvr_planning.pdf
Veritas Volume Replicator Web Console Administrator's Guide	vvr_web_admin.pdf
Veritas Volume Replicator Advisor User's Guide	vvr_advisor_users.pdf
Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide	vvr_agents_config.pdf

Veritas Storage Foundation for Oracle RAC documentation

The following Storage Foundation for Oracle RAC documentation is available with all Veritas Storage Foundation HA product suites:

Table 1-11 describes the Storage Foundation for Oracle RAC documentation set.

Table 1-11 Guides in Storage Foundation for Oracle RAC documentation set

Guide Title	Filename
Veritas Storage Foundation™ for Oracle RAC Installation and Configuration Guide	`sfrac_install.pdf`
Veritas Storage Foundation™ for Oracle RAC Release Notes	`sfrac_notes.pdf`
Veritas Storage Foundation™ for Oracle RAC Administrator's Guide	`sfrac_admin.pdf`

Manual Pages

The Veritas online manual pages are installed in the `/opt/VRTS/man` directory. This directory can be added to the `MANPATH` environment variable.

If the `MANPATH` environment variable does not include `/opt/VRTS/man`, you can view the desired manual page by entering the following command:

```
# man -M /opt/VRTS/man manual_page_name
```

Note: Installing documentation and online manual pages is optional.

Documentation errata

The following describes documentation errata.

Manual pages errata

Several manual pages have been updated in the 5.0.1 release to include corrections for serveral errors or omissions.

Dbdst_report manual page errata (1361604)

For the `dbdst_report` manual page, one of the usage examples is incorrectly displayed as follows:

```
$ /opt/VRTS/bin/dbdst_show_fs -S $ORACLE_SID -o audit
```

The correct usage example is:

```
$ /opt/VRTS/bin/dbdst_report -S $ORACLE_SID -o audit
```

Workaround

There is no workaround for this issue.

Dbdst manual pages errata (1121091)

You cannot access the correct manual pages for the `dbdst` commands if you type in the command name as follows:

```
# man dsdst_command_name
```

Workaround

You must append the correct database suffix for each `dbdst` command in order to access the correct man page for that `dbdst` command.

To obtain correct usage notes for the `dbdst` command in an Oracle environment, you must use the following command format:

```
# man dsdst_command_name_ora
```