

# Veritas Storage Foundation<sup>TM</sup> Cluster File System Release Notes

HP-UX

5.0 Maintenance Pack 1



# Veritas Storage Foundation Cluster File System Release Notes

Copyright © 2007 Symantec Corporation. All rights reserved.

Storage Foundation Cluster File System 5.0 Maintenance Pack 1

Symantec, the Symantec logo, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THIS DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID, SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be "commercial computer software" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation  
20330 Stevens Creek Blvd.  
Cupertino, CA 95014  
[www.symantec.com](http://www.symantec.com)

Printed in the United States of America.

## Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product 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>

HP-UX is a registered trademark of Hewlett-Packard Development Company, L.P.

## Licensing and registration

Veritas File System Installation Guide is a licensed product. See the *Veritas File System Installation Guide* for license installation instructions.

## Technical support

For technical assistance, visit

[http://www.symantec.com/enterprise/support/assistance\\_care.jsp](http://www.symantec.com/enterprise/support/assistance_care.jsp) and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.



# Contents

Introduction .....	6
System requirements .....	7
HP-UX operating system requirements .....	7
Storage Foundation Cluster File System node requirements .....	7
Hardware compatibility .....	7
Required patches .....	8
Component product release notes .....	8
Installing for the first time .....	8
Preparing to upgrade to 5.0 MP1 .....	9
Upgrading to 5.0 MP1 .....	9
Phased upgrade stages .....	10
Full upgrade stages .....	10
Performing a phased upgrade .....	11
Performing a full upgrade .....	15
Removing 5.0 MP1 .....	19
Storage Foundation Cluster File System patches .....	20
Software limitations .....	21
Fixed issues .....	22
Known issues .....	23
Oracle-HP .....	24

# Storage Foundation Cluster File System Release Notes

## Introduction

This document provides information specific to the Maintenance Pack 1 (MP1) for Storage Foundation Cluster File System (SFCFS) 5.0.

- For information on topics not covered in this document, such as SFCFS-specific software fixes and limitations, see the *Veritas Storage Foundation Cluster File System 5.0 Release Notes*.
- For information on Storage Foundation, including software fixes and limitations for Veritas File System, Veritas Volume Replicator, and Veritas Volume Manager, see the *Veritas Storage Foundation 5.0 MP1 Release Notes*.

---

**Note:** Read the following Veritas Technical Support TechNote for the latest information on updates, patches, and software issues for this release:

<http://entsupport.symantec.com/docs/281875>

---

# System requirements

This section describes the system requirement for this release.

## HP-UX operating system requirements

The system requirements for this release are:

- HP-UX 11.23

### To verify the operating system version

- ◆ Use the `swlist` command as follows:

```
# swlist | grep HPUXBaseAux
HPUXBaseAux          B.11.23.0409 HP-UX Base OS Auxiliary
```

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.11.23          The Base VxFS File System
```

## Storage Foundation Cluster File System node requirements

All nodes in a cluster file system must have the same operating system version and update level.

## Hardware compatibility

For information on servers, network controllers, and disk subsystems that support SFCFS, see the hardware compatibility list available at <http://entsupport.symantec.com>.

## Required patches

Patch requirements and versions are determined at the time of product release. Contact your vendor for the most current patch version and information. For a list of current patches for component products, see the *Veritas Storage Foundation Cluster File System 5.0 Installation Guide*.

## Component product release notes

In addition to reading these Release Notes, review all component product release notes before installing the product. The following component product release notes are included as PDF files on the Veritas software disc:

*Veritas Storage Foundation Release Notes (sf\_notes.pdf)*

*Veritas Cluster Server Release Notes (vcs\_notes.pdf)*

Because release notes are not installed by any packages and are not on the documentation disc, it is recommended that you copy them from the software disc to the `/opt/VRTS/docs` directory so they are available on your system for reference.

## Installing for the first time

For new installations, first install Veritas Storage Foundation Cluster File System 5.0, then upgrade to the MP1 level. Review and follow the installation procedure for Storage Foundation Cluster File System 5.0 described in the *Veritas Storage Foundation Cluster File System 5.0 Installation Guide*. Follow the procedures in this document to upgrade Storage Foundation Cluster File System to 5.0 MP1.



## Preparing to upgrade to 5.0 MP1

If you are upgrading an installed Veritas Storage Foundation 5.0 or 5.0 RP1 system, read the *Veritas Storage Foundation Installation Guide* or *Veritas Storage Foundation Cluster File System 5.0 Installation Guide* for instructions on how to preserve the existing configuration information.

In particular, perform the following actions:

- Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/fstab` file. You will need to recreate these entries in the `/etc/fstab` file on the freshly installed system.
- Before upgrading, ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in files such as `/etc/fstab`. You should also run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands, and record the output from these. You may need this information to reconfigure your system after the upgrade.
- Use the `vxlicrep` command to make a record of the currently installed Veritas licenses.

You should also review the *Veritas Storage Foundation Release Notes* or *Veritas Storage Foundation Cluster File System Release Notes*, and all documents in the `release_notes` directory for important release information.

## Upgrading to 5.0 MP1

An upgrade requires stopping cluster failover functionality during the entire procedure. The upgrade is performed in a number of stages depending on the type of upgrade you are performing. The supported upgrade paths are:

- SFCFS 5.0
- SFCFS 5.0 RP1

---

**Caution:** Phased upgrade procedure results in system PANIC on configurations where LLT is configured over UDP and this known issue is fixed in 5.0 MP1. This issue is specific to configurations where LLT is configured over UDP and not present in usual LLT Ethernet configurations.

Full upgrade procedure should be used for upgrading from SFCFS 5.0 or SFCFS 5.0 RP1 to SFCFS 5.0 MP1 on configurations where LLT is configured over UDP.

---

## Phased upgrade stages

A phased upgrade minimizes downtime by upgrading portions of the cluster, one at a time. Although the entire cluster is offline for a shorter period than a full upgrade, this method requires command-line interaction and some manual configuration.

### Phased upgrade

If your system is running SFCFS 5.0 or SFCFS 5.0 RP1 the following are the stages of a phased upgrade.

#### Stages of a phased upgrade

- 1 Freeze service group operations and stop VCS on the cluster.
- 2 Select a group of one or more cluster nodes to upgrade, and leave a group of one or more nodes running.
- 3 Take the first group offline and install the software patches.
- 4 Take the second group offline **before** bringing the first group online.
- 5 Bring the first group (with the newly installed patches) online to restart cluster failover services.
- 6 Upgrade the remaining nodes in the second group and bring them online. The cluster is fully restored.

Proceed to [“Performing a phased upgrade”](#) on page 11.

## Full upgrade stages

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

#### To perform a full upgrade

- 1 Freeze service group operations and stop VCS on the cluster.
- 2 Take all nodes in the cluster offline and install the software patches.
- 3 Bring all the nodes (with the newly installed patches) online to restart cluster failover services. The cluster is fully restored.

Proceed to [“Performing a full upgrade”](#) on page 15.

## Performing a phased upgrade

This section describes how to upgrade to 5.0 MP1.

### To upgrade to 5.0 MP1

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
- 3 From any node in the cluster, make the VCS configuration writable:  

```
# haconf -makerw
```
- 4 Enter the following command to freeze HA service group operations on each node:  

```
# hasys -freeze -persistent node_name
```
- 5 Make the configuration read-only:  

```
# haconf -dump -makero
```
- 6 Select the group of nodes that are to be upgraded first, and follow [step 7](#) through [step 34](#) for these nodes.
- 7 Stop VCS by entering the following command on each node in the group being upgraded:  

```
# hastop -local
```
- 8 Check if each node's root disk is under VxVM control by running this command:  

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- a Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Do not remove the plexes on the root disk that correspond to the original disk partitions.

- b** Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.

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

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

- 9** On each node, use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# bdf -t vxfs
```

- 10** On each node in the cluster, unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
```

```
# umount /filesystem
```

- 11** If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- a** Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- b** Use the `vxrvg stop` command to stop each RVG individually:

```
# vxrvg -g diskgroup stop rvg_name
```

- c** On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

To avoid data corruption, do not proceed until all RLINKs are up-to-date.

- 12** Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

- 13** On each node, stop all VxVM volumes by entering the following command for each disk group:

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

- 14** Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```

- 15** Insert the Veritas software disc on any node and determine the block device file for the CD drive:

```
# ioscan -fnC disk
```

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

- 16 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /cdrom
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /cdrom
```
- 17 Change the directory to /cdrom:

```
# cd /cdrom
```
- 18 To upgrade the Storage Foundation Cluster File System, you must invoke the `installmp` command from one of your cluster nodes using the option that corresponds to your configuration:
  - To install on the local system only, enter the following command:

```
# ./installmp
```
  - To install on more than one system using secure shell (SSH) utilities, enter the following command:

```
# ./installmp node_name1 node_name2 ...
```
  - To install on more than one system using remote shell (RSH) utilities, enter the following command:

```
# ./installmp node_name1 node_name2 ... -rsh
```
- 19 After the initial system checks are complete, press **Return** to start the requirements checks.
- 20 After the requirements checks are complete, press **Return** to start upgrading the packages. If you are upgrading multiple nodes, you have the option of upgrading them simultaneously. You will be prompted after the upgrade is complete.
- 21 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.
- 22 Stop VCS on each of the second group of (non-upgraded) nodes:

```
# hastop -local
```
- 23 Reboot the newly upgraded nodes. Assuming the reboot was successful, application failover is now available for the first group of nodes.
- 24 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.
- 25 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.
- 26 If any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.

- 27 Make the VCS configuration writable again from any node in the upgraded group:  

```
# haconf -makerw
```
- 28 Enter the following command on each node in the upgraded group to unfreeze HA service group operations:  

```
# hasys -unfreeze -persistent node_name
```
- 29 Make the configuration read-only:  

```
# haconf -dump -makero
```
- 30 Bring the CVM service group online on each node in the upgraded group:  

```
# hagrps -online cvm -sys node_name
```
- 31 Restart all the volumes by entering the following command for each disk group:  

```
# vxvol -g diskgroup startall
```
- 32 If you stopped any RVGs in [step 11](#), restart each RVG:  

```
# vxrvrg -g diskgroup start rvg_name
```
- 33 Remount all VxFS file systems and Storage Checkpoints on all nodes:  

```
# mount /filesystem  
# mount /checkpoint_name
```
- 34 Check if the VEA service was restarted:  

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```
- 35 Repeat [step 9](#) through [step 34](#), excluding [step 22](#), for the second group of nodes.

## Performing a full upgrade

The following procedure describes upgrading a standalone system.

### To upgrade the Storage Foundation software on a standalone system

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.

- 3 On any node in the cluster, run the following command:

```
# hastop -all
```

This stops VCS on all nodes in the cluster.

- 4 Check if each node's root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- a Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Do not remove the plexes on the root disk that correspond to the original disk partitions.

- b Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.

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

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

- 5 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# bdf -t vxfs
```

- 6 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

## 7 Verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | fsdb -F vxfs filesystem | grep clean
flags 0 mod 0 clean clean_value
```

A *clean\_value* value of 0x5a indicates the file system is clean, 0x3c indicates the file system is dirty, and 0x69 indicates the file system is dusty. A dusty file system has pending extended operations.

- a If a file system is not clean, enter the following commands for that file system:

```
# fsck -F vxfs filesystem
# mount -F vxfs filesystem mountpoint
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.

There may be a pending large fileset clone removal extended operation if the `umount` command fails with the following error:

```
file system device busy
```

You know for certain that an extended operation is pending if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system
file system still in progress.
```

- b If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
- c Repeat [step 7](#) to verify that the unclean file system is now clean.
- 8 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:
- a Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- b Use the `vxrvrg stop` command to stop each RVG individually:
- ```
# vxrvrg -g diskgroup stop rvg_name
```
- c On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:
- ```
# vxrlink -g diskgroup status rlink_name
```
- To avoid data corruption, do not proceed until all RLINKs are up-to-date.
- 9 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.



- 10 Stop all VxVM volumes by entering the following command for each disk group:  

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

To verify that no volumes remain open, use the following command:  

```
# vxprint -Aht -e v_open
```
- 11 Check if the VEA service is running:  

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is running, stop it:  

```
# /opt/VRTS/bin/vxsvcctl stop
```
- 12 Insert the Veritas software disc on any node and determine the block device file for the CD drive:  

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.
- 13 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:  

```
# mkdir -p /cdrom  
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /cdrom
```
- 14 Change the directory to /cdrom:  

```
# cd /cdrom
```
- 15 To upgrade the Storage Foundation Cluster File System, you must invoke the `installmp` command from one of your cluster nodes using the option that corresponds to your configuration:
  - To install on the local system only, enter the following command:  

```
# ./installmp
```
  - To install on more than one system using secure shell (SSH) utilities, enter the following command:  

```
# ./installmp node_name1 node_name2 ...
```
  - To install on more than one system using remote shell (RSH) utilities, enter the following command:  

```
# ./installmp node_name1 node_name2 ... -rsh
```
- 16 After the initial system checks are complete, press **Return** to start the requirements checks.
- 17 After the requirements checks are complete, press **Return** to start upgrading the packages. If you are upgrading multiple nodes, you have the option of upgrading them simultaneously. You will be prompted after the upgrade is complete.
- 18 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.
- 19 Shut down and reboot the system.

- 20 If necessary, reinstate any missing mount points in the `/etc/fstab` file.
- 21 If any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.
- 22 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```
- 23 If you stopped any RVGs in [step 8](#), restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```
- 24 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem
# mount /checkpoint_name
```
- 25 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

# Removing 5.0 MP1

The following procedure removes the patches if you want to uninstall 5.0 MP1.

## To remove the patches

- 1 Log in as superuser (`root`).
- 2 Stop the cluster:

```
# hstop -all
```
- 3 On all the nodes, stop VxFEN processes:

```
# /sbin/init.d/vxfen stop
```
- 4 On all the nodes, stop the currently running VxPAL agents (see `vxpalctrl` (1M)):
  - a Stop the storage agent:

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a StorageAgent -c stop
```
  - b Check the status of the storage agent:

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a StorageAgent -c status
```
  - c Stop the action agent:

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a actionagent -c stop
```
  - d Check the status of the action agent:

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a actionagent -c status
```
  - e Stop the gridnode agent:

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a gridnode -c stop
```
  - f Check the status of the gridnode agent:

```
# /opt/VRTSobc/pal33/bin/vxpalctrl -a gridnode -c status
```
- 5 On all the nodes, remove all the patches using the `swremove` command:

```
# swremove -x autoreboot=true patch_name1, patch_name2 ...
```

See “[Storage Foundation Cluster File System patches](#)” on page 20.

Symantec recommends that all the patches installed during 5.0 MP1 installation be removed through a single command line.

The system automatically reboots after removing the patches.

---

**Caution:** The patches PHKL\_35178, PHCO\_35179, and PVCO\_03696 should always be installed and removed in a single command line. Not installing or removing them together can result in data loss.

---

## Storage Foundation Cluster File System patches

The following are the patches that are installed as a part of the SFCFS 5.0 MP1 installation:

Patch ID	Patch Name
PHCO_35124	VRTSvmpro Command Patch
PHCO_35125	VRTSddlpr Command Patch
PHCO_35126	VRTSalloc Command Patch
PHCO_35127	VRTSdcli Command Patch
PHCO_35179	VRTSvxvm Command Patch
PHCO_35180	VRTSfspro Command Patch
PHCO_35212	VRTSob Command Patch
PHCO_35213	VRTSobc33 Command Patch
PHCO_35214	VRTSobgui Command Patch
PHCO_35215	VRTSaa Command Patch
PHCO_35216	VRTSccg Command Patch
PHCO_35217	VRTSmh Command Patch
PHCO_35301	VRTSat Command Patch
PHCO_35332	VRTSvxfS Command Patch
PHCO_35375	VRTSfsman Command Patch
PHCO_35425	VRTSvxfen Command Patch
PHKL_35178	VRTSvxvm Kernel Patch
PHKL_35305	VRTSvxfS Kernel Patch
PHKL_35312	VRTSglm Kernel Patch
PHNE_35413	VRTSilt Kernel Patch
PHKL_35424	VRTSvxfen Kernel Patch
PVCO_03671	VRTSvrpro Command Patch
PVCO_03672	VRTSvrw GUI Patch

Patch ID	Patch Name
PVCO_03676	VRTSvcs/VRTSvcsag Command Patch
PVCO_03678	VRTScscm Command Patch
PVCO_03680	VRTScssim Command Patch
PVCO_03686	VRTScmccc Command Patch
PVCO_03689	VRTScmcs Command Patch
PVCO_03690	VRTScscw Command Patch
PVCO_03694	VRTScavf Command Patch
PVCO_03696	VRTSsvsc Command Patch
PVCO_03697	SYMClma Command Patch
PVCO_03698	VRTSsmf Command Patch
PHNE_35783	VRTSgab Kernel Patch

## Software limitations

Veritas Storage Foundation Cluster File System software limitations 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>

The Veritas Storage Foundation, Veritas Volume Manager, Veritas File System, Veritas Volume Replicator, and Veritas Storage Foundation for Databases software limitations are listed in the *Veritas Storage Foundation 5.0 MP1 Release Notes*.

See the *Veritas Storage Foundation 5.0 MP1 Release Notes*.

There are no new Veritas Storage Foundation Cluster File System software limitations in this MP1 release.

## Fixed issues

Veritas Storage Foundation Cluster File System fixed 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>

The Veritas Storage Foundation, Veritas Volume Manager, Veritas File System, Veritas Volume Replicator, and Veritas Storage Foundation for Databases fixed issues are listed in the *Veritas Storage Foundation 5.0 MP1 Release Notes*.

See the *Veritas Storage Foundation 5.0 MP1 Release Notes*.

The following are new Veritas Storage Foundation Cluster File System fixed issues in this MP1 release:

Incident	Description
156186	In environments using LDAP or NIS+, the CFSfsckd agent scripts could hang or time out because of unresponsive network-based name services occurring due to incorrect options included in the ps command. The problem has been resolved by setting the correct options for the ps command to avoid name lookups over the network.
647359	Under certain conditions, the <code>cmhaltcl (1M)</code> command was failing to halt the cluster nodes in a ServiceGuard cluster when invoked with the <code>-f</code> option. This was occurring due to the incorrect freeing of CFS response messages during the retry of a previously sent message in case of an error. The problem has been resolved by freeing the cluster node's response message in case of an error so that the failed message does not get allocated again.
704284	The <code>fsclustadm (1M)</code> was continuously returning a "device busy" message when using the <code>cfsdeinit</code> command line option, in a ServiceGuard cluster. During a forced unmount of a CFS file system the hold count was not being correctly updated causing an EBUSY error status to be returned to <code>fsclustadm (1M)</code> . The problem has been resolved by correctly updating the hold count during inode inactivation.
771108	Oracle 10gr1 database creation on raw volumes with ODM was failing. The problem has been resolved by adding a new permission bit to the ODM interface.

Incident	Description
797453	Performance enhancements were required in the handling of inherited extended attributes. The problem has been resolved by ensuring that LCT buffers are written to disk only when necessary instead of every time they're modified.
799189	The Oracle instance was crashing because ODM was returning ENOMEM. The problem has been resolved by changing the arena allocation call to use the WAITOK flag.
802327	A Data Page Fault was being encountered in a fast lookup path for NFS access when a node joins the cluster. The problem has been resolved by ensuring that the code paths that provide fast access for NFS requests, initialize the vfs pointer whenever it has not yet been initialized.
816497	The direct access of a 64-bit data item in a CFS message was causing an alignment trap. After a cluster node was rebooted, a remaining cluster node was crashing in <code>vx_do_rbdele_resp()</code> due to the alignment trap. The problem has been resolved by avoiding direct access of the 64-bit data item in the CFS message by using <code>bzero()</code> and <code>bcopy()</code> to move data in and out of the CFS message buffer.
817879	A blocking file lock request over NFS was resulting in a hung cluster, for a file that is residing on a cluster mounted file system. The problem has been resolved by ensuring that the file locking code in CFS handles the blocking file lock requests over NFS appropriately.
847819	A Data Page Fault was being encountered in <code>vx_find_fsparams()</code> due to a missing check for an unmounted file system. The problem has been resolved by ensuring that the file system has not been unmounted before accessing the in-core structures of the file system in <code>vx_async_getdele()</code> .

For a list of additional issues fixed in this release, see the following TechNote:  
<http://entsupport.symantec.com/docs/285869>

## 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>

The Veritas Storage Foundation, Veritas Volume Manager, Veritas File System, Veritas Volume Replicator, and Veritas Storage Foundation for Databases known issues are listed in the *Veritas Storage Foundation 5.0 MP1 Release Notes*. See the *Veritas Storage Foundation 5.0 MP1 Release Notes*.

## Oracle-HP

### 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:

```
# /opt/VRTSdbcom/config/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. [840486]