

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

Linux

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 Cluster File System 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/283857>

Linux is a registered trademark of Linus Torvalds.

## Licensing and registration

Veritas Storage Foundation Cluster File System is a licensed product. See the *Veritas Storage Foundation Cluster 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
New features .....	7
Block device support in Veritas File System .....	7
System requirements .....	7
Supported Linux operating systems .....	7
Storage Foundation Cluster File System node requirements .....	8
Hardware compatibility .....	8
Memory requirements .....	8
Component product release notes .....	8
Installing for the first time .....	9
Preparing to upgrade to 5.0 MP1 .....	9
Upgrading to 5.0 MP1 .....	10
Phased upgrade stages .....	11
Full upgrade stages .....	11
Performing a phased upgrade .....	12
Performing a full upgrade .....	16
Verifying software versions .....	19
Removing 5.0 MP1 .....	19
Software limitations .....	20
Fixed issues .....	21
Known issues .....	23
Oracle-Linux .....	23
DB2-Linux .....	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/281993>

---

# New features

Veritas Storage Foundation Cluster File System provides the following new features:

## Block device support in Veritas File System

VxFS 5.0 only supported the creation of a VxFS file system on a VxVM volume. Creating a VxFS file system on raw devices, such as SCSI or IDE disks, was not supported.

VxFS 5.0 MP1 supports the creation of VxFS file systems directly on raw devices without requiring a VxVM volume. However, this feature is limited to local mount file systems only and is not supported for CFS mounts. Both partitioned and non-partitioned block devices may now be used directly with VxFS.

For example:

```
# mkfs -t vxfs /dev/sdc1
# mount -t vxfs /dev/sdc1 /mnt
# fsck -t vxfs /dev/sdc1
```

Similarly, all local mount features of VxFS, such as clones and file system snapshots, do not require the creation of a VxVM volume.

# System requirements

This section describes the system requirement for this release.

## Supported Linux operating systems

Storage Foundation Cluster File System operates on the following Linux operating systems and kernel binaries distributed by Red Hat and SUSE:

- Red Hat Enterprise Linux 4 (RHEL 4) with Update 3 (2.6.9-34 kernel) or higher on AMD Opteron or Intel Xeon EM64T (x86\_64).
- SUSE Linux Enterprise Server 9 (SLES 9) with SP3 (2.6.5-7.244, 252 kernels) or higher on AMD Opteron or Intel Xeon EM64T (x86\_64).

---

**Note:** If your system is running an older version of either Red Hat Enterprise Linux or SUSE Linux Enterprise Server, you must upgrade it before attempting to install the Veritas Storage Foundation Cluster File System software. Consult the Red Hat or SUSE documentation for more information on upgrading your system.

---

## Supported Linux Operating System Updates

Veritas products will operate on subsequent kernel and patch releases provided the operating systems maintain kernel application binary interface (ABI) compatibility.

Information about the latest supported Red Hat erratas and updates and SUSE service packs is available in the following TechNote. Read this TechNote before installing any Veritas product.

<http://entsupport.symantec.com/docs/281993>

## 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/docs/283161>

## Memory requirements

2 GB of memory is required.

## 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 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 `/boot/grub/menu.lst`, `/etc/grub.conf`, `/etc/elilo.conf`, or `/etc/lilo.conf` (as appropriate), and `/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 to 5.0 MP1
- SFCFS 5.0 RP1 to 5.0 MP1

---

**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.

### 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 12.

## 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.

### Stages of 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 16.

## 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 35](#) 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:

```
# df -T | grep vxfs
```

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

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

- 11 If there are still disk groups that are imported at this time then proceed to [step 12](#), otherwise skip to [step 15](#).

- 12 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.

- 13 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.

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

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

- 15 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

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

- 16 Insert the Veritas software disc on any node and mount the CD-ROM. For example:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```
- 17 Move to the top-level directory on the CD-ROM:

```
# cd /mnt/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 (Optional) If you are going to upgrade your system, perform the following:
  - a Upgrade your system, and patch it to a kernel version. For more information on supported operating systems and kernel versions, see [“Supported Linux operating systems”](#) on page 7.
- 23 Stop VCS on each of the second group of nodes:

```
# hastop -local
```
- 24 Reboot the upgraded nodes. Assuming the reboot was successful, application failover is now available for the first group of nodes.
- 25 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*.
- 26 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.

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

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

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

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

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

```
# vxvol -g diskgroup startall
```
- 33 If you stopped any RVGs in [step 12](#), restart each RVG:

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

```
# mount /filesystem
# mount /checkpoint_name
```
- 35 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
```
- 36 Repeat [step 7](#) through [step 35](#) 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 `rootdsk` 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:

```
# df -T | grep vxfs
```

- 6 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
```

```
# umount /filesystem
```



- 7 If there are still disk groups that are imported at this time then proceed to [step 8](#), otherwise skip to [step 11](#).
- 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/vxsvcctrl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```
- 12 Insert the Veritas software disc on any node and mount the CD-ROM. For example:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```
- 13 Change to the top-level directory on the CD-ROM:

```
# cd /mnt/cdrom
```
- 14 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
```

- 15 After the initial system checks are complete, press **Return** to start the requirements checks.
- 16 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.
- 17 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.
- 18 (Optional) If you are going to upgrade your system, perform the following:
  - a Upgrade your system, and patch it to a kernel version. For more information on supported operating systems and kernel versions, see [“Supported Linux operating systems”](#) on page 7.
- 19 Shut down and reboot the system.
- 20 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*.
- 21 If necessary, reinstate any missing mount points in the `/etc/fstab` file.
- 22 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.
- 23 Restart all the volumes by entering the following command for each disk group:

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

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

```
# mount /filesystem
# mount /checkpoint_name
```
- 26 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
```

## Verifying software versions

To list the Veritas packages installed on your system, enter the following command:

```
# rpm -qa | grep VRTS
```

## Removing 5.0 MP1

Roll back of the 5.0 MP1 packages to the release 5.0 version of the packages is not supported. To restore release 5.0 on your system, you must completely remove 5.0 MP1, then reinstall Veritas Storage Foundation Cluster File System 5.0.

### To uninstall the Veritas software

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 3 Run the `hastop -all` command to take all service groups offline, and shut down VCS:

```
# /opt/VRTSvcs/bin/hastop -all
```

---

**Note:** Do not use the `-force` option when executing `hastop`. This will leave all service groups online and shut down VCS, causing undesired results during uninstallation of the packages.

---

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

```
# df -T | grep vxfs
```

- 5 Unmount all Storage Checkpoints and file systems:

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

- 6 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.

- 7 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
```

- 8 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
```
- 9 To shut down and remove the installed Veritas packages, use the appropriate command in the `/opt/VRTS/install` directory. For example, to uninstall the Storage Foundation Cluster File System, use the following commands:  

```
# cd /opt/VRTS/install  
# ./uninstallsfcfs [-rsh]
```

You can use this command to remove the packages from one or more systems. The `-rsh` option is required if you are using the remote shell (RSH) rather than the secure shell (SSH) to uninstall the software simultaneously on several systems.

---

**Note:** Provided that the remote shell (RSH) or secure shell (SSH) has been configured correctly, this command can be run on a single node of the cluster to install the software on all the cluster nodes.

---

After uninstalling the Veritas software, reinstall the release 5.0 software as described in the *Veritas Storage Foundation Installation Guide*, *Veritas Storage Foundation Cluster File System Installation Guide*, or *Veritas Cluster Server (VCS) Installation Guide*.

## 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/283857>

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/283857>

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
612406	Push the entire extent for an ilist hole. This will prevent multi-entry holes.
645227	The file system was hanging due to a rare race between the merge and activation of File Change Log (FCL) which lead to a deadlock. The problem has been resolved by ensuring that the merge threads try for the FCL global lock and proceed only if the lock is available.
770917	In some code paths related to large directory support, the cluster member (node) was not taking and holding the ownership of the inode for which the extents are to be re-organized. So, an assertion was being triggered when another member node of the cluster tried to revoke ownership of the inode. The problem has been resolved by correcting the ownership issues in large directory related code paths.
771892	A deadlock was occurring during CFS recovery. Log replay, which is performed as the first step during CFS recovery when a node crashes, was found to block in certain scenarios on an inode that the replay is trying to process. The problem has been resolved by freeing all inodes in the chunk one at a time without marking other inodes as busy.

Incident	Description
771970	When a file is truncated, the truncate request is shipped to the cluster member that has the current ownership of the inode. There were certain inconsistencies in pre-splitting a buffer in the presence of clones, when the truncate request was shipped to another node of the cluster leading to an assertion. The problem has been resolved by changing the truncation code paths to handle the pre-splitting of buffers at a later stage during truncation, for both locally mounted and cluster mounted file systems.
771999	A node in a cluster was hanging under heavy CFS activity in the cluster. When all the worker threads servicing callbacks from a queue were blocked, the callbacks that were still lying in the same queue faced a deadlock. The problem has been resolved by queuing the callbacks that need to be processed to unblock other callbacks into a different queue, and processing them at a higher priority.
775611	Upon encountering an I/O error, the file system is correctly disabled but an internal inconsistency check fails. The check failure in this instance is quite benign and cannot cause customer data corruption since the file system is safely disabled upon encountering an I/O error of this nature. To further diagnose this condition, a new consistency check has been added.
782024	Under heavy inode cache stress, file system recovery on a cluster file system was hanging after an I/O error. The problem has been resolved by ensuring that the code paths that attempt to reuse an inode do not wait for a dirty/bad/inactive inode to become free if recovery initiated by a disabled file system is currently in progress. If there are no inodes available for re-use, a new inode will be allocated to ensure that the recovery proceeds to completion.
793030	The vxfsutil.h was using 'struct timeval' in one of the function declarations without including time.h resulting in a warning message during compilation. The problem has been resolved by including time.h in vxfsutil.h.

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/283857>

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*.

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

## Oracle-Linux

### 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 `rpm -e` command.

```
# rpm -e VRTSorgui-common VRTSdbed-common VRTSdbcom-common \  
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]

## DB2-Linux

### Problems uninstalling or upgrading Veritas Storage Foundation for DB2 when Veritas Storage Foundation Cluster File System is installed on the same system

If Veritas Storage Foundation for DB2 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 DB2 packages manually.

#### To uninstall the Veritas Storage Foundation for DB2 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:

```
# hstop -local
```

To stop VCS processes on all systems:

```
# hstop -all
```

## 3 Remove the Veritas Storage Foundation for DB2 packages using the `rpm -e` command.

```
# rpm -e VRTSd2gui-common VRTSdb2ed-common VRTSdbcom-common \ VRTSdbdoc
```

If Veritas Storage Foundation for DB2 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 DB2.

If the second upgrade fails, remove the Veritas Storage Foundation for DB2 packages as described above, then run the installer to upgrade Veritas Storage Foundation for DB2. [840486]