# Veritas Storage Foundation™ and High Availability Release Notes

AIX

6.0 Rolling Patch 1

August 2012

# Veritas Storage Foundation and High Availability Release Notes

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 6.0 Rolling Patch 1

Document version: 6.0RP1.3

## Legal Notice

THE 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 as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations.  Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

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 support 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/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

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

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

## Contacting Technical Support

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

www.symantec.com/business/support/contact_techsupp_static.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:

www.symantec.com/business/support/

## Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as 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 support 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

## Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support 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 |

## Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec Web site.

https://sort.symantec.com/documents

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:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation

## About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

http://www.symantec.com/connect/storage-management

# Contents

Chapter 1

# About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

## Introduction

This document provides information about the Veritas Storage Foundation and High Availability Solutions 6.0 Rolling Patch 1 release.

# About the installrp and the uninstallrp scripts

Veritas Storage Foundation and High Availability Solutions 6.0 RP1 provides an upgrade script.

See "Supported upgrade paths" on page 96.

Symantec recommends that you use the upgrade script. The `installrp` script allows you to upgrade all the patches associated with the packages installed, after which you can reboot to start all of the processes.

## The installrp script options

**Table 1-1** The command line options for the product upgrade script

| Command Line Option | Function |
|---|---|
| *system1 system2...* | Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name. |
| –precheck | Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product. |
| –postcheck | Checks any issues after installation or upgrading on the system. |
| –responsefile *response_file* | Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file. |
| –logpath *log_path* | Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved. |
| –tmppath *tmp_path* | Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where filesets are copied on remote systems before installation. |

**Table 1-1** The command line options for the product upgrade script *(continued)*

| Command Line Option | Function |
|---|---|
| –timeout *timeout_value* | The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the -timeout option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 will prevent the script from timing out. The -timeout option does not work with the -serial option |
| –keyfile *ssh_key_file* | Specifies a key file for secure shell (SSH) installs. This option passes `-i ssh_key_file` to every SSH invocation. |
| –hostfile *full_path_to_file* | Specifies the location of a file that contains a list of hostnames on which to install. |
| –rootpath *root_path* | Specifies an alternative root directory on which to install filesets. |
| –serial | Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems. |
| –rsh | Specifies this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. |
| –redirect | Displays progress details without showing the progress bar. |
| –pkgset | Discovers and displays the fileset group (minimum, recommended, all) and filesets that are installed on the specified systems. |
| –pkgtable | Displays product's filesets in correct installation order by group. |
| –pkginfo | Displays a list of filesets and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkginfo option with the installvcs script to display VCS filesets. |

**Table 1-1** The command line options for the product upgrade script *(continued)*

| Command Line Option | Function |
|---|---|
| –listpatches | The `–listpatches` option displays product patches in correct installation order. |
| –makeresponsefile | Use the `–makeresponsefile` option only to generate response files. No actual software installation occurs when you use this option. |
| –comcleanup | The `–comcleanup` option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated. |
| –version | Checks and reports the installed products and their versions. Identifies the installed and missing filesets and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing filesets and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available. |
| –nolic | Allows installation of product filesets without entering a license key. Licensed features cannot be configured, started, or used when this option is specified. |
| -rolling_upgrade | Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly. |
| –rollingupgrade_phase1 | The `–rollingupgrade_phase1` option is used to perform rolling upgrade Phase-I. In the phase, the product kernel filesets get upgraded to the latest version |
| –rollingupgrade_phase2 | The `–rollingupgrade_phase2` option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent filesets upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version." |

## The uninstallrp script options

Veritas Storage Foundation and High Availability Solutions 6.0 RP1 provides a new uninstallation script.

See About rolling back Veritas Storage Foundation and High Availability Solutions 6.0 RP1 for release versions and products that support rolling back.

Symantec recommends that you use the new uninstallation script. The `uninstallrp` script uninstalls all the patches associated with packages installed, and starts the processes.

**Table 1-2** The command line options for the product upgrade script

| Command Line Option | Function |
|---|---|
| *system1 system2...* | Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name. |
| –responsefile *response_file* | Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file. |
| –logpath *log_path* | Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved. |
| –tmppath *tmp_path* | Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where filesets are copied on remote systems before installation. |
| –timeout *timeout_value* | The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the -timeout option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 will prevent the script from timing out. The -timeout option does not work with the -serial option |

**Table 1-2** The command line options for the product upgrade script *(continued)*

| Command Line Option | Function |
|---|---|
| –keyfile *ssh_key_file* | Specifies a key file for secure shell (SSH) installs. This option passes `-i ssh_key_file` to every SSH invocation. |
| –hostfile *full_path_to_file* | Specifies the location of a file that contains a list of hostnames on which to install. |
| –serial | Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems. |
| –rsh | Specifies this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. |
| –redirect | Displays progress details without showing the progress bar. |
| –listpatches | The `-listpatches` option displays product patches in correct installation order. |
| –makeresponsefile | Use the `-makeresponsefile` option only to generate response files. No actual software installation occurs when you use this option. |
| –comcleanup | The `-comcleanup` option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated. |
| –version | Checks and reports the installed products and their versions. Identifies the installed and missing filesets and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing filesets and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available. |

# Overview of the installation and upgrade process

Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

**To install the Veritas software for the first time**

1. Skip this step if you are upgrading to 5.1 SP1 RP2. If you are installing 5.1 SP1 RP2 for the first time:

   - Download Storage Foundation and High Availability Solutions 5.1 SP1 from http://fileConnect.symantec.com.
   - Extract the tar ball into a directory called /tmp/sfha51sp1.
   - Check http://sort.symantec.com/patches to see if there are any patches available for the 5.1SP1 Installer. Download applicable P-patches and extract them to the /tmp directory.
   - Change the directory to /tmp/sfha51sp1:

     ```
     # cd  /tmp/sfha51sp1
     ```

   - Install the 5.1 SP1 software. Follow the instructions in the Installation Guide.

     ```
     ./installer -require complete_path_to_SP1_installer_patch
     ```

2. Download SFHA 5.1 SP1 RP2 from http://sort.symantec.com/patches and extract it to a directory called /tmp/sfha51sp1rp2.
3. Check http://sort.symantec.com/patches to see if there are patches available for the 5.1SP1RP2 installer. Download applicable P-patches and extract them to the /tmp directory.
4. Change the directory to /tmp/sfha51sp1rp2:

   ```
   #cd  /tmp/sfha51sp1rp2
   ```

5. Install 5.1SP1 RP2:

   ```
   ./installer -require complete_path_to_SP1RP2_installer_patch
   ```

# System requirements

This section describes the system requirements for this release.

## Supported AIX operating systems

This section lists the supported operating systems for this release of Veritas products.

Table 1-3 shows the supported operating systems for this release.

**Table 1-3** Supported operating systems

| Operating systems | Levels | Chipsets |
|---|---|---|
| AIX 7.1 | TL0 or later | Any chipset that the operating system supports |
| AIX 6.1 | TL5 or later | Power 5, Power 6 and Power 7 |

Be sure to install IBM APAR for AIX 6.1 TL6 and TL7, or AIX 7.1 TL0 and TL1. Contact IBM to get the necessary APAR for your level. For example, you may need APAR IVO3362.

For Storage Foundation for Oracle RAC, all nodes in the cluster must have the same operating system version and update level.

## AIX 7.1 support for virtual processors

Veritas Storage Foundation and High Availability supports up to 1024 virtual processors on AIX 7.1.

## Hardware compatibility list

The compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware go to the following URL:

http://www.symantec.com/docs/TECH170013

For information on specific HA setup requirements, see the *Veritas Cluster Server Installation Guide*.

## Veritas Storage Foundation for Database features supported in database environments

Veritas Storage Foundation for Database (SFDB) features are supported for the following database environments:

**Table 1-4** SFDB features database support for 6.0 RP1

| SFDB feature | DB2 | Oracle | Sybase |
|---|---|---|---|
| Oracle Disk Manager, Cached Oracle Disk Manager | No | Yes | No |
| Quick I/O, Cached Quick I/O | Yes | Yes | Yes |
| Concurrent I/O | Yes | Yes | Yes |
| Storage Checkpoints | Yes | Yes | Yes |
| Flashsnap | Yes | Yes | Yes |
| SmartTier | Yes | Yes | Yes |
| Database Storage Checkpoints | No | Yes | No |
| Database Flashsnap | No | Yes | No |
| SmartTier for Oracle | No | Yes | No |

Review current documentation for your database to confirm the compatibility of your hardware and software.

For the most current information on Storage Foundation products and single instance Oracle versions supported, see:

http://www.symantec.com/docs/DOC4039

## Veritas Storage Foundation memory requirements

Symantec recommends 2 GB of memory over the minimum requirement for the operating system.

## Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

# Fixed issues

This section covers the incidents that are fixed in this release.

## Veritas Volume Manager: Issues fixed in 6.0 RP1

This section describes the incidents that are fixed in Veritas Volume Manager (VxVM) in 6.0 RP1.

**Table 1-5** Veritas Volume Manager 6.0 RP1 fixed issues

| Fixed issues | Description |
|---|---|
| 2680604 | vxconfigbackupd does not work correctly with NUM_BK. |
| 2674465 | Data Corruption while adding/removing LUNs. |
| 2666163 | A small portion of possible memory leak introduced due to addition of enhanced messages. |
| 2657797 | Starting 32TB RAID5 volume fails with unexpected kernel error in configuration update. |
| 2649958 | vxdmpadm dumps core due to null pointer reference. |
| 2647795 | Intermittent data corruption after a vxassist move. |
| 2634072 | vxdisk output shows junk character. |
| 2627056 | vxmake -g <DGNAME> -d <desc-file> fails with very large configuration due to memory leaks. |
| 2626741 | Using vxassist -o ordered and mediatype:hdd options together do not work as expected. |
| 2621465 | When detached disk after connectivity restoration is tried to reattach gives 'Tagid conflict' error. |
| 2620556 | I/O hung after SRL overflow. |
| 2620555 | I/O hang due to SRL overflow and CVM reconfig. |
| 2610877 | In cluster configuration dg activation can hang due to improper handling of error codes. |
| 2610764 | In cluster configuration i/o can hang on master node after storage is removed. |
| 2608849 | Logowner local I/O starved with heavy I/O load from Logclient. |
| 2607519 | Secondary master panics in case of reconfig during autosync. |
| 2607293 | Primary master panic'ed when user deleted frozen RVG. |
| 2600863 | vxtune doesn't accept tunables correctly in human readable format. |
| 2591321 | while upgrading dg version if rlink is not up-to-date the vxrvg command shows error but dg version gets updated. |

**Table 1-5** Veritas Volume Manager 6.0 RP1 fixed issues *(continued)*

| Fixed issues | Description |
|---|---|
| 2590183 | write fails on volume on slave node after join which earlier had disks in "lfailed" state. |
| 2576602 | vxdg listtag should give error message and display correct usage when executed with wrong syntax. |
| 2575581 | vxtune -r option is printing wrong tunable value. |
| 2574752 | Support utility vxfmrmap (deprecating vxfmrshowmap) to display DCO map contents and verification against possible state corruptions. |
| 2565569 | read/seek i/o errors during init/define of nopriv slice. |
| 2562416 | vxconfigbackup throws script errors due to improper handling of arguments. |
| 2556467 | disabling all paths and rebooting host causes /etc/vx/.vxdmprawdev record loss. |
| 2530698 | after "vxdg destroy" hung (for shared DG), all vxcommands hang on master. |
| 2527289 | Both sites become detached after data/dco plex failue at each site, leading to i/o cluster wide outage. |
| 2526498 | Memory leaks seen in some I/O code path. |
| 2516584 | startup scripts use 'quit' instead of 'exit', causing empty directories in /tmp. |
| 2348180 | Failure during validating mirror name interface for linked mirror volume. |
| 1967512 | Need revisit of open/close ioctl implementation for DMPnode and its paths. |

## Veritas File System: Issues fixed in 6.0 RP1

This section describes the incidents that are fixed in Veritas File System (VxFS) in 6.0 RP1.

**Table 1-6** Veritas File System 6.0 RP1 fixed issues

| Fixed issues | Description |
|---|---|
| 2678096 | The fiostat command dumps core when the count value is 0. |
| 2663750 | Abrupt messages are seen in engine log after complete storage failure in cvm resiliency scenario. |

**Table 1-6** Veritas File System 6.0 RP1 fixed issues *(continued)*

| Fixed issues | Description |
|---|---|
| 2655786 | Shared' extents are not transferred as 'shared' by the replication process. |
| 2655754 | Deadlock because of wrong spin lock interrupt level at which delayed allocation list lock is taken. |
| 2653845 | When the fsckptadm(1M) command with the '-r' and '-R' option is executed, two mutually exclusive options gets executed simultaneously. |
| 2650330 | Accessing a file with O_NSHARE mode by multiple process concurrently on Aix could cause file system hang. |
| 2645441 | Native filesystem migrated to vxfs disk layout 8 where layout version 9 is the default. |
| 2645435 | The following error message is displayed during the execution of the fsmap(1M) command:'UX:vxfs fsmap: ERROR: V-3-27313'. |
| 2645112 | write operation on a regular file mapping to shared compressed extent results in corruption. |
| 2645109 | In certain rare cases after a successful execution of vxfilesnap command, if the source file gets deleted in a very short span of time after the filesnap operation, then the destination file can get corrupted and this could also lead to setting of VX_FULLFSCK flag in the super block. |
| 2645108 | In certain cases write on a regular file which has shared extent as the last allocated extent can fail with EIO error. |
| 2630954 | The fsck(1M) command exits during an internal CFS stress reconfiguration testing. |
| 2626390 | New tunable - chunk_inval_size and few more option with 'chunk_flush_size'. |
| 2624459 | Listing of a partitioned directory using the DMAPI does not list all the entries. |
| 2613884 | Metadata corruption may be seen after recovery. |
| 2609002 | The De-duplication session does not complete. |
| 2599590 | Expanding or shrinking a DLV5 file system using the fsadm(1M)command causes a system panic. |
| 2583197 | Upgrade of a file system from version 8 to 9 fails in the presence of partition directories and clones. |
| 2563251 | fsmigadm "commit/status" error messages should be clear. |

**Table 1-6** Veritas File System 6.0 RP1 fixed issues *(continued)*

| Fixed issues | Description |
|---|---|
| 2552095 | The system may panic while re-organizing the file system using the fsadm(1M) command. |
| 2536130 | The fscdsconv(1M) command which is used to convert corrupted or non-VxFS file systems generates core. |
| 2536054 | A hang may be seen because VxFS falsely detect low pinnable memory scenario. |
| 2389318 | Enabling delayed allocation on a small file system sometimes disables the file system. |

## Veritas Cluster Server: Issues fixed in 6.0 RP1

This seciton describes the incidents that are fixed in Veritas Cluster Server (VCS) in 6.0 RP1.

**Table 1-7** Veritas Cluster Server 6.0 RP1 fixed issues

| Fixed issues | Description |
|---|---|
| 2684822 | If a pure local attribute like PreOnline is specified before SystemList in `main.cf`, then it gets rejected when HAD is started. |
| 2653668 | The high availability daemon (HAD) process unexpectedly terminates. |
| 2644483 | "VCS ERROR V-16-25-50036 The child service group came online (recovered) before the parent was offlin ed." message is logging as ERROR message. |
| 2635211 | AMF calls VxFS API with spinlock held. |
| 2616497 | Fault propagation does not work if the parent is in faulted state on one or more nodes. |

## Veritas Enterprise Administrator: Issues fixed in 6.0 RP1

This seciton describes the incidents that are fixed in Veritas Enterprise Administrator (VEA) in 6.0 RP1.

**Table 1-8** Veritas Enterprise Administrator 6.0 RP1 fixed issues

| Fixed issues | Description |
|---|---|
| 2672039 | vxsvc process crashes and dumps core. |
| 2677191 | File placement policy operations are not available in VEA GUI. |

# Known issues

This section covers the known issues in this release.

## Installation known issues

This section describes the known issues in this release of installation.

### Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups (2591399)

The service groups freeze due to upgrading using the product installer if you stopped the installer after the installer already stopped some of the processes and then resumed the upgrade.

Workaround: You must unfreeze the service groups manually after the upgrade completes.

**To unfreeze the service groups manually**

1 List all the frozen service groups

```
# hagrp -list Frozen=1
```

2 Unfreeze all the frozen service groups:

```
# haconf -makerw
# hagrp -unfreeze service_group -persistent
# haconf -dump -makero
```

### Ignore VRTSgms request to boot during installation (2143672)

During installation, you may see this error which you can ignore.

```
VRTSgms: old driver is still loaded...
VRTSgms: You must reboot the system after installation...
```

## Web installer does not ask for authentication after the first session if the browser is still open (2509330)

If you install or configure SFHA and then close the Web installer, if you have other browser windows open, the Web installer does not ask for authentication in the subsequent sessions. Since there is no option to log out of the Web installer, the session remains open as long as the browser is open on the system.

Workaround: Make sure that all browser windows are closed to end the browser session and subsequently log in again.

## The VRTSvxvm fileset fails to install on a few cluster nodes because the template file is corrupted (2348780)

The installer debug log displays the failure of the `errupdate` command as following: `errupdate -f /usr/lpp/VRTSvxvm/inst_root/VRTSvxvm.err`. The `errupdate` command gets invoked through `/usr/lib/instl/install` by the operating system. The command also fails for the VRTSvxfs, VRTSglm, and VRTSgms packages.

The `errupdate` command generally creates a `*.undo.err` file to remove entries from the Error Record Template Repository in case of failed installation or cleanup. However, in this case the `*.undo.err` file does not get generated as the `errupdate` command fails. Also, it is not possible to manually remove entries from the Error Record Template Repository in order to undo the changes made by the failed installation, because the file is corrupted.

**Workaround:** Save a copy of the`/var/adm/ras/errtmplt` and `/etc/trcfmt` files before you install the product. Replace `/var/adm/ras/errtmplt` and `/etc/trcfmt` files with the ones that you saved, when the installation fails because the template file is corrupted. Uninstall all the packages you installed and reinstall.

## After finishing a kernel upgrade on a master node the cvm group on a slave node does not come online (2439439)

After successfully finishing a kernel upgrade on one node, the cvm group does not come online on the second node.

Workaround: Check that your cluster is not in a jeopardy state before you perform a rolling upgrade.

## sfmh discovery issue when you upgrade your Veritas product to 6.0 (2622987)

If a host is not reporting to any management server but sfmh discovery is running before you upgrade to 6.0, sfmh-discovery may fail to start after the upgrade.

Workaround:

If the host is not reporting to VOM, stop sfmh-discovery manually before upgrading to 6.0 by executing the following command on the managed host:

```
/opt/VRTSsfmh/adm/vxvmdiscovery-ctrl.sh stop
```

### When you uninstall CommandCentral Storage Managed Host from a system where Veritas Storage Foundation 6.0 is installed, SF 6.0 reconfiguration or uninstallation fails (2631486)

On a system where Veritas Storage Foundation (SF) 6.0 is installed, if you uninstall CommandCentral Storage (CCS) Managed Host (MH) using the installer script from the CCS media, the installer script removes the contents of `/opt/VRTSperl`. As a result, SF 6.0 reconfiguration or uninstallation using `/opt/VRTS/install/install_`*`sf_product_name`* or `/opt/VRTS/install/uninstall_`*`sf_product_name`* fails, because the installer script removed the contents of `/opt/VRTSperl`.

**Workaround:** To uninstall CCS MH from a system where SF 6.0 is installed, before you perform the uninstallation, perform the procedure in the following CCS TechNote:

http://www.symantec.com/business/support/index?page=content&amp;id=HOWTO36496

### Incorrect server names sometimes display if there is a clock synchronization issue (2627076)

When you install a cluster with the Web-based installer, you choose to synchronize your systems with an NTP server due to a clock synchronization issue, you may see the NTP server name in messages instead of your server names.

Workaround:

Ignore the messages. The product is still installed on the correct servers.

### Manual upgrade of VRTSvlic fileset loses keyless product levels [2115662]

If you upgrade the `VRTSvlic` fileset manually, the product levels that were set using `vxkeyless` may be lost. The output of the `vxkeyless display` command will not display correctly. To prevent this, perform the following steps while manually upgrading the VRTSvlic fileset.

1. Note down the list of products configured on the node for keyless licensing.

```
# vxkeyless display
```

2. Set the product level to NONE.

   ```
   # vxkeyless set NONE
   ```

3. Upgrade the VRTSvlic fileset.

   ```
   # installp -u VRTSvlic
   ```

   This step may report a dependency, which can be safely overridden.

   ```
   # installp -acgX -d pathname VRTSvlic
   ```

4. Restore the list of products that you noted in step 1.

   ```
   # vxkeyless set product[|,product]
   ```

## Issues with keyless licensing reminders after upgrading VRTSvlic [2141446]

After upgrading from 5.1 to higher versions of VCS, some keyless licenses may be left in the system. As a result, you may see periodic reminders being logged if the VOM server is not configured.

This happens if you are using keyless licenses before upgrading to 5.1SP1 or higher versions of VCS. After the upgrade, you install real keys and run `vxkeyless set NONE`. In this case, the keyless licenses may not be completely removed and you see warning messages being logged after two months (if VOM server is not configured). This does not result in any functionality impact.

To resolve this issue, perform the following steps:

1. Note down the list of products configured on the node for keyless licensing. Run `vxkeyless display` to display the list.

2. Set the product level to *NONE* with the command:

   ```
   # vxkeyless set NONE
   ```

3. Find and delete the keyless licenses left over in the system. To do this, perform the following steps for every key stored in `/etc/vx/licenses/lic`:

   - Verify if the key has `VXKEYLESS` feature Enabled using the following command:

     ```
     # vxlicrep -k <license_key> | grep VXKEYLESS
     ```

   - Delete the key if and only if `VXKEYLESS` feature is Enabled.

   **Note:** When performing the search, do not include the .vxlic extension as part of the search string.

4. Restore the previous list of products with the command:

```
# vxkeyless set product1[|,product]
```

## Secure WAC communication needs to be disabled explicitly [2392568]

If you have WACs communicating securely where VCS is configured in secure mode and if you disable the VCS security, the WAC where VCS security is disabled continues attempting to communicate securely without success. Therefore, you need to explicitly disable WAC security when you disable VCS security.

Workaround: No workaround. Secure WAC communication needs to be disabled explicitly.

## Web installer has no option to remove node from a cluster

Web Installer does not provide the option to remove node from a cluster.

Workaround: Manually remove nodes from a cluster. There is no option to remove nodes available from Web Installer or CPI.

## Perl messages seen in engine log during rolling upgrade [2627360]

While performing a rolling upgrade from VCS 5.1SP1 to 6.0 with MultiNICA resource configured, if VRTSperl fileset is upgraded but VRTSvcsag fileset is not yet upgraded on the system, Perl code related messages may be seen. The messages seen are similar to the following:

```
Using a hash as a reference is deprecated at MultiNICA.pm line 108.
```

Workaround: Complete the rolling upgrade to VCS 6.0.

## After a locale change restart the vxconfig daemon (2417547)

You need to restart the vxconfig daemon you change the locale of nodes that use it. The vxconfig daemon starts at boot. If you have changed locale, you need to restart the daemon.

**Workaround:**

Refer to the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for the section, "vxconfigd daemon recovery."

### After performing a rolling upgrade, make sure the CVM is online on all nodes without errors (2595441)

Make sure that the CVM is online on all nodes without errors after you perform the first phase of a rolling upgrade. The CVM protocol version will not upgrade successfully on the nodes where CVM is offline or has errors.

If the CVM protocol version does note upgrade successfully, upgrade the CVM protocol on the CVM master node.

**To upgrade the CVM protocol on the CVM master node**

1 Find out which node is the CVM master. Enter the following:

```
# vxdctl -c mode
```

2 On the CVM master node, upgrade the CVM protocol. Enter the following:

```
# vxdctl upgrade
```

### Unable to stop some SFHA processes (2329580)

If you install and start SFHA, but later configure SFHA using `installvcs`, some drivers may not stop successfully when the installer attempts to stop and restart the SFHA drivers and processes. The reason the drivers do not stop is because some dependent SFHA processes may be in the running state.

Workaround: To re-configure the product, use the corresponding `install`*`product`* command to re-configure the product. Otherwise some processes may fail to stop or start.

For example, use `installsfrac` to re-configure SFHA rather than using `installvcs`.

## Veritas Dynamic Multi-pathing known issues

This section describes the known issues in this release of Veritas Dynamic Multi-pathing.

### Some paths in DMP can get DISABLED if LVM volume group is created on OS device path (1978941)

On AIX, when an LVM volume group is created directly on the OS device path, the SCSI driver performs SCSI2 reservation on the rest of the paths to that LUN. As a result, some of the paths of the corresponding DMP devices may be disabled, as shown by the `vxdmpadm getsubpaths` command output. For some arrays, the `vxdisk list` command shows the device in the '`error`' state.

This issue is not seen when LVM volume groups are created on the DMP devices.

Example of this issue:

```
# vxdisk list | grep emc0_00bc
emc0_00bc    auto:none       -            -            online invalid

# vxdmpadm getsubpaths dmpnodename=emc0_00bc
NAME     STATE[A]  PATH-TYPE[M] CTLR-NAME ENCLR-TYPE ENCLR-NAME ATTRS
======================================================================
hdisk110 ENABLED(A)   -          fscsi0    EMC        emc0        -
hdisk123 ENABLED(A)   -          fscsi0    EMC        emc0        -
hdisk136 ENABLED(A)   -          fscsi1    EMC        emc0        -
hdisk149 ENABLED(A)   -          fscsi1    EMC        emc0        -

# vxdisk rm emc0_00bc

# mkvg -y dmxvg hdisk110
dmxvg

# lspv | egrep "hdisk110|hdisk123|hdisk136|hdisk149"
hdisk110        00c492ed6fbda6e3          dmxvg           active
hdisk123        none                      None
hdisk136        none                      None
hdisk149        none                      None

# vxdisk scandisks

# vxdmpadm getsubpaths dmpnodename=emc0_00bc
NAME     STATE[A] PATH-TYPE[M] CTLR-NAME ENCLR-TYPE ENCLR-NAME ATTRS
=====================================================================
hdisk110 ENABLED(A) -           fscsi0    EMC       emc0        -
hdisk123 DISABLED   -           fscsi0    EMC       emc0        -
hdisk136 DISABLED   -           fscsi1    EMC       emc0        -
hdisk149 DISABLED   -           fscsi1    EMC       emc0        -
```

### To recover from this situation

1. Varyoff the LVM volume group:

   ```
   # varyoffvg dmxvg
   ```

2. Remove the disk from VxVM control.

   ```
   # vxdisk rm emc0_00bc
   ```

3. Trigger DMP reconfiguration.

   ```
   # vxdisk scandisks
   ```

4. The device which was in DISABLED state now appears as ENABLED.

   ```
   # vxdmpadm getsubpaths dmpnodename=emc0_00bc
   NAME     STATE[A] PATH-TYPE[M] CTLR-NAME ENCLR-TYPE ENCLR-NAME ATTRS
   ===================================================================
   hdisk110 ENABLED(A)   -        fscsi0    EMC        emc0       -
   hdisk123 ENABLED(A)   -        fscsi0    EMC        emc0       -
   hdisk136 ENABLED(A)   -        fscsi1    EMC        emc0       -
   hdisk149 ENABLED(A)   -        fscsi1    EMC        emc0       -
   ```

## I/O fails on some paths after array connectivity is restored, due to high restore daemon interval (2091619)

If a path loses connectivity to the array, the path is marked as suspected to fail and hence is not used for I/O. After the connectivity is restored, the restore daemon detects that the path is restored when the restore daemon probes the paths. The restore daemon makes the path available for I/O. The restore daemon probes the paths at the interval set with the tunable parameter dmp_restore_interval. If you set the dmp_restore_interval parameter to a high value, the paths are not available for I/O until the next interval.

## Changes in enclosure attributes are not persistent after an upgrade to VxVM 6.0 RP1 (2082414)

The Veritas Volume Manager (VxVM) 6.0 RP1 includes several array names that differ from the array names in releases prior to release 5.1SP1. Therefore, if you upgrade from a previous release to VxVM 6.0 RP1, changes in the enclosure attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0 RP1. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-9 shows the Hitachi arrays that have new array names.

**Table 1-9** Hitachi arrays with new array names

| Previous name | New name |
|---|---|
| TagmaStore-USP | Hitachi_USP |
| TagmaStore-NSC | Hitachi_NSC |

**Table 1-9** Hitachi arrays with new array names *(continued)*

| Previous name | New name |
|---|---|
| TagmaStoreUSPV | Hitachi_USP-V |
| TagmaStoreUSPVM | Hitachi_USP-VM |
| <New Addition> | Hitachi_R700 |
| Hitachi AMS2300 Series arrays | New array names are based on the Model Number 8x. For example, AMS_100, AMS_2100, AMS_2300, AMS_2500, etc. |

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.0 RP1. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

## DS4K series array limitations

In case of DS4K array series connected to AIX host(s), when all the paths to the storage are disconnected and reconnected back, the storage does not get discovered automatically. To discover the storage, run the `cfgmgr` OS command on all the affected hosts. After the `cfgmgr` command is run, the DMP restore daemon brings the paths back online automatically in the next path restore cycle. The time of next path restore cycle depends on the restore daemon interval specified (in seconds) by the tunable dmp_restore_interval.

```
# vxdmpadm gettune dmp_restore_interval
        Tunable               Current Value  Default Value
------------------------     -------------  -------------
dmp_restore_interval               300            300
```

On DS4K array series connected to AIX host(s) DMP is supported in conjunction with RDAC. DMP is not supported on DS4K series arrays connected to AIX hosts In MPIO environment.

### vxconfigd hang with path removal operation while IO is in-progress (1932829)

In AIX with HBA firmware version SF240_320, vxdisk scandisks (device discovery) takes a long time when a path is disabled from the switch or from the array.

**Workaround:**

To resolve this issue, upgrade the HBA firmware version to SF240_382.

### Adding a DMP device or its OS device path as a foreign disk is not supported (2062230)

When DMP native support is enable, adding a DMP device or its OS device path as a foreign disk using the `vxddladm addforeign` command is not supported. Using this command can lead to unexplained behavior.

### The secondary path of DMP nodes cannot be added after rebooting VIOS (2564356)

When you check the path status using command `vxdmpadm getsubpaths ctlr=fscsi*` and then reboot VIOS, addition of the secondary path of DMP nodes is not accepted.

When you check the path status using command `vxdmpadm getsubpaths ctlr=fscsi*` and then execute `vxdisks Scandisks`, addition of the secondary path of DMP nodes is accepted.

### DMP native support is not persistent after upgrade to 6.0 (2526709)

The DMP tunable parameter dmp_native_support is not persistent after upgrade to DMP 6.0. After you upgrade, set the tunable parameter using the following command:

```
# vxdmpadm settune dmp_native_support=on
```

### Devices unmanaged from PowerPath go into error state (2482308)

After unmanaging devices from PowerPath, devices go into an error state.

**Work-around:**

Reboot the system to enabled DMP to claim the devices.

### Array controller reboot on CLARiiON storage in failovermode 1 or 4 on AIX 6.1 (2418875)

On an AIX 6.1 host, when you reboot the array controller on a CLARiiON array that is configured in failovermode 1 or 4, the dmpnode may go in failed state, resulting in I/O failures on the LUN.

## Veritas Storage Foundation known issues

This section describes the known issues in this release of Veritas Storage Foundation.

### Veritas Storage Foundation known issues

This section describes the known issues in this release of Veritas Storage Foundation (SF).

#### AT Server crashes when authenticating unixpwd user multiple times (1705860)

There is a known issue in the AIX kernel code that causes '`getgrent_r`' function to corrupt the heap. This issue is present in AIX 5.3 and AIX 6.1 Refer to IBM's Web site for more information:

http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52585

AT uses `getgrent_r` function to get the groups of the authenticated user.

IBM has released the fix as a patch to fileset bos.rte.libc. There are different patches available for different version of bos.rte.libc. You need to check the version of `bos.rte.libc` (For example: `lslpp -l | grep bos.rte.libc`) and apply the appropriate IBM patch:

- For version 6.1.3.1:
  http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52959
  For the fix:
  ftp://ftp.software.ibm.com/aix/efixes/iz52959/
- For version 6.1.2.4:
  http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52720
  For the fix:
  ftp://ftp.software.ibm.com/aix/efixes/iz52720/
- For version 6.1.2.5 :
  http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52975
  For the fix:
  ftp://ftp.software.ibm.com/aix/efixes/iz52975/

There are IBM patches for only certain version of bos.rte.libc that are available. If your system has a different bos.rte.libc version, you may have to upgrade to a higher version where the fix is available. If your version is not available, you may have to contact IBM.

### Not all the objects are visible in the VOM GUI (1821803)

After upgrading SF stack from 5.0MP3RP2 to 5.1, the volumes are not visible under the Volumes tab and the shared diskgroup is discovered as Private and Deported under the Disgroup tab in the SFM GUI.

**Workaround:**

**To resolve this known issue**

◆ On each manage host where VRTSsfmh 2.1 is installed, run:

```
# /opt/VRTSsfmh/adm/dclisetup.sh -U
```

### A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a dynamic storage tiering placement policy (1880622)

A volume's placement class tags are not visible in the Veritas Enterprise Administrator (VEA) GUI when you are creating a SmartTier placement policy if you do not tag the volume with the placement classes prior to constructing a volume set for the volume.

**Workaround:** To see the placement class tags in the VEA GUI, you must tag the volumes prior to constructing the volume set. If you already constructed the volume set before tagging the volumes, restart vxsvc to make the tags visible in the GUI.

## Veritas Volume Manager known issues

The following are the Veritas Volume Manager known issues for this release.

### Tunable values doesn’t change as per the values applied through vxtune (2698035)

The vxtune command displays tunable values in terms of default unit. When a new value is provided to a tunable, vxtune accepts it in terms of default unit. For example, voldrl_min_regionsz is stored in terms of blocks (OS block size). So if a value of 1024 is provided, then vxtune interprets it as 1024 blocks; if a value of 2M is provided, it interprets it as 2M blocks (2097152 blocks), not 4096 blocks (2MB).

**Workaround**:

It is recommended to provide exact tunable value without any suffix. This will be addressed in future releases.

### The failure of one disk causes all the sub-disks to be relocated (2641510)

In a campus cluster environment, failure of one disk causes all the sub-disks, on the storage devices tagged with the specific site, to be relocated.

Ideally, the disks that encounter an I/O failure only are marked for hot-relocation (RLOC). However, the vxrelocd daemon does not try to recover the volumes that become available after the devices at the specific sites are partially reattached. It checks if the sub-disks/plexes are still detached and then relocates all the sub-disks.

There is no workaround for this issue.

### Node join can lead to hang if an upgrade of the cluster protocol version is in progress (2103567)

If you attempt to join a node to the cluster while Cluster Volume Manager (CVM) is upgrading the cluster protocol version, the system may hang. This issue occurs if the node is attempting to join the cluster after you issue the `vxdctl upgrade` command to upgrade the CVM cluster.

**Work-around:**

Avoid joining a new node to the cluster until the CVM cluster upgrade is completed.

### The cluster may hang if a node goes down (1835718)

The cluster may hang if a node goes down while one array is disabled or offline in a mirror=enclosure configuration.

This may occur, if a node panics or loses power while one array of a mirror=enclosure configuration is offline or disabled, then the cluster, fencing, I/O loads, and VxVM transactions hang.

**Workaround:** There is no workaround for this issue.

### vxconvert failures if PowerPath disks are formatted as simple disks (857504)

If a PowerPath disk is formatted as a simple disk (a foreign device), then the vxconvert utility may fail during conversion of LVM to VxVM. To view the format of the disk, use the vxdisk list command. This issue may also occur if the /etc/vx/darecs file contains an hdiskpower disk entry. This entry may be present if PowerPath disks were configured as foreign disks in Storage Foundation 4.0, and the entry was not changed after subsequent upgrades.

### Veritas Volume Manager (VxVM) might report false serial split brain under certain scenarios (1834513)

VxVM might detect and report a false serial split brain when all of the following conditions are met:

- One or more arrays that provide the shared storage for the cluster are being powered off
- At the same time when the arrays are being powered off, an operation that requires an internal transaction is initiated (such as VxVM configuration commands)

In such a scenario, disk group import will fail with a split brain error and the vxsplitlines output will show 0 or 1 pools.

**Workaround:**

**To recover from this situation**

1. Retrieve the disk media identifier (dm_id) from the configuration copy:

   ```
   # /etc/vx/diag.d/vxprivutil dumpconfig device-path
   ```

   The dm_id is also the serial split brain id (ssbid)

2. Use the dm_id in the following command to recover from the situation:

   ```
   # /etc/vx/diag.d/vxprivutil set device-path ssbid=dm_id
   ```

### vxdisk -f init can overwrite some of the public region contents (1190117)

If a disk was initialized by a previous VxVM version or defined with a smaller private region than the new default of 32 MB, then the public region data will be overridden.

**Workaround:**

Specify explicitly the length of privoffset, puboffset, publen, and privlen while initializing the disk.

### The relayout operation fails when there are too many disks in the disk group. (2015135)

The attempted relayout operation on a disk group containing approximately more than 300 LUNs or disks may fail with the following error:

```
Cannot setup space
```

### Co-existence check might fail for CDS disks

In Veritas Volume Manager (VxVM) 5.1 SP1, VxVM introduces the ability to support Cross-platform Data Sharing (CDS) on disks larger than 1 TB. VxVM uses the SUN VTOC Table to initialize the cdsdisk layout on devices up to 1 TB. VxVM uses the GUID Partition Table (GPT) to initialize the cdsdisk layout on devices larger than 1 TB.

In layouts where SUN VTOC Table is used for initialization (typically, when the disk size has never exceeded 1 TB), the AIX co-existence label can be found at sector 7 and VxVM ID block (also known as HP co-existence label) can be found at sector 16.

In layouts where GPT is used for initialization (typically, when the disk size is currently greater than or had earlier exceeded 1 TB), the AIX co-existence label is placed at sector 55 and VxVM ID block (also known as HP co-existence label) is placed at sector 64. Consequently, AIX utilities would not be able to recognize a cdsdisk initialized using GPT to be a valid VxVM disk. Symantec is working with IBM and third party OEMs to enhance the co-existence check in these utilities.

**Workaround:** There is no workaround for this issue.

### vxsnap addmir command sometimes fails under heavy I/O load (2441283)

The `vxsnap addmir` command sometimes fails under heavy I/O load and produces multiple errors.

**Workaround:** Rerun the `vxsnap addmir` command.

### The vxassist maxsize option fails to report the maximum size of the volume that can be created with given constraints when the disk group has the siteconsistent flag set (2563195)

The `vxassist maxsize` option fails to report the maximum size of volume that can be created with given constraints when the disk group has the siteconsistent flag set. The following error is reported:

```
# vxassist -g dgname maxsize
VxVM vxassist ERROR V-5-1-752 No volume can be created within the given
constraints
```

**Workaround:**

Specify the size explicitly to the `vxassist make` command.

### Hardware paths for operating system paths have changed in DMP 6.0 (2410716)

In DMP 6.0, the hardware paths for operating system paths have changed. After upgrading to DMP 6.0, path attributes are reset to the default values. You must reconfigure any path-level attributes that were defined in the /etc/vx/dmppolicy.info file.

**Workaround:**

**To configure path-level attributes**

1. Remove the path entries from the `/etc/vx/dmppolicy.info` file.
2. Reset the path attributes.

### Upgrading SFHA to version 6.0 marks vSCSI disks as cloned disks (2434444)

This issue is seen when you upgrade from a previous version of SFHA which has vSCSI disks included in a disk group. After upgrading SFHA to 6.0, the vSCSI disks that were included in a disk group are marked as cloned disks.

**Workaround:**

Use the following procedure to clear the clone disk flag.

**To clear the clone disk flag**

1. Remove the vSCSI devices that are in error state (ibm_vscsi#_#) using the following command:

   ```
   # vxdisk rm device_name
   ```

2. Deport the disk group.

   ```
   # vxdg deport dg_name
   ```

3. Re-import the disk group with a new udid.

   ```
   # vxdg -o updateid import dg_name
   ```

4. Display the devices that are part of the disk group.

   ```
   # vxdisk -g dg_name list
   ```

5. Clear the clone_disk tag from these devices.

   ```
   # vxdisk set device_name clone=off
   ```

### The vxsnap print command shows incorrect value for percentage dirty (2360780)

The `vxsnap print` command can display the percentage of regions that differ between snapshots, shown as the %dirty. In SFHA 6.0, if this command is run while the volumes are online and being actively used, the shown %dirty may lag from actual percentage dirty for instant snap data change object (DCO) volumes. That is, the command output may show less %dirty than actual.

### Encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path names are not the same (2607706)

The encapsulation of a multi-pathed root disk fails if the dmpnode name and any of its path name are not the same.

For example:

Dmpnode:sdh

Paths: sda sdb

**Work-around:**

Before running the encapsulation command (`vxencap`), run the following command:

```
# vxddladm assign names
```

### Recovery and rollback to original configuration may not succeed if the system reboots while the online migration setup is in partial state (2611423)

During online migration from LVM to VxVM volumes, if there is a system reboot when the migration setup is in partial state, that is, the start operation has not completed successfully, then the recover and abort operations might not be able to recover and rollback the configuration.

Workaround: This needs manual intervention for cleanup, depending on the state, to restore the original configuration.

## Veritas File System known issues

This section describes the known issues in this release of Veritas File System (VxFS).

### A cluster mounted file system may cause system panic showing vx_tflush_map in the stack trace (2558892)

VxFS causes the server to panic. The subroutine initiating the panic is `vx_tflush_map`.

There is no workaround for this issue.

### Full file system check takes over a week (2628207)

On a large file system with many checkpoints, a full file system check using the `fsck_vxfs`(1m) command may appear to be hung.

There is no workaround for this issue.

### umount(1m) on a CFS file system causes system panic (2107152)

In rare corner cases, the system panics while unmounting a cluster mounted file system.

There is no workaround for this issue.

### fsckptadm(1m) fails with ENXIO (1956458)

The `fsckptadm`(1m) fails with the ENXIO error and the file system is marked for full file system check.

There is no workaround for this issue.

### vxfsd consumes a lot of CPU resources after deleting some directories (2129455)

The vxfsd daemon consumes 100% CPU after some directories are deleted.

There is no workaround for this issue.

### The fsppadm command does not work while assigning a policy (2715414)

The `fsppadm`(1m) dumps core with SIGSEGV while assigning a policy.

**Workaround:** Increase the pthread stack size using the following command.

```
export PTHREAD_DEFAULT_STACK_SIZE=2048000
```

### vxfsckd resource fails for start when a cluster node is rebooted or when it is killed manually (2720034)

If you reboot a node in a cluster and kill the `vxfsckd` resource manually, vxfsckd does not come up and the cvm services are faulted.

**Workaround**:

Use the following commands for this situation:

```
hastop -local
rm /var/adm/cfs/vxfsckd-pid
```

Kill all `vxfsckd` processes:

```
fsclustadm cfsdeinit
hastart
```

### Cannot use some commands from inside an automounted Storage Checkpoint (2490709)

If your current work directory is inside an automounted Storage Checkpoint, for example `/mnt1/.checkpoint/clone1`, some commands display the following error:

```
can't find current directory
```

This issue is verified with the following commands:

- `cp -r`
- `du`

However, this issue might occur with other commands.

**Workaround:** Run the command from a different directory.

### Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system nears 100% usage even if other volumes have free space (2438368)

Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system is nearing 100% usage even if other volumes in the file system have free space.

**Workaround:** After sufficient space is freed from the volume, delayed allocation automatically resumes.

### A mutex contention in vx_worklist_lk() can use up to 100% of a single CPU (2086902)

A mutex contention in the `vx_worklist_lk`() call can use up to 100% of a single CPU.

**Workaround:** There is no workaround for this issue.

### Performance on a VxFS file system can be slower than on a JFS file system (2511432)

There are two causes for the performance degradation:

- There are unnecessary page faults that set the write permissions on mapped pages.
- The entire file flushes.

The issue of unnecessary page faults has been fixed, which has improved the performance considerably. However, the performance on a VxFS file system can still be slower than on a JFS file sometimes due to the entire file flushing.

**Workaround:** There is no workaround for this issue.

### Deduplication can fail with error 110 (2591473)

In some cases, data deduplication fails with a message similar to the following example:

```
Saving    Status    Node            Type        Filesystem
---------------------------------------------------------------------
00%       FAILED    node01          MANUAL      /data/fs1
        2011/10/26 01:38:58 End full scan with error
```

In addition, the deduplication log contains an error similar to the following example:

```
2011/10/26 01:35:09 DEDUP_ERROR AddBlock failed. Error = 110
```

These errors indicate that the deduplication process is running low on space and needs more free space to complete.

**Workaround:** Make more space available on the file system.

### You are unable to unmount the NFS exported file system on the server if you run the fsmigadm command on the client (2355258)

Unmounting the NFS-exported file system on the server fails with the "Device busy" error when you use the `fsmigadm` command on the NFS client.

**Workaround:** Unexport the file system prior to unmounting.

### vxresize fails while shrinking a file system with the "blocks are currently in use" error (2437138)

The `vxresize` shrink operation may fail when active I/Os are in progress on the file system and the file system is being shrunk to a size closer to its current usage. You see a message similar to the following example:

```
UX:vxfs fsadm: ERROR: V-3-20343: cannot shrink /dev/vx/rdsk/dg1/vol1 -
blocks are currently in use.
VxVM vxresize ERROR V-5-1-7514 Problem running fsadm command for volume
vol1, in diskgroup dg1
```

**Workaround:** Rerun the shrink operation after stopping the I/Os.

## Replication known issues

This section describes the replication known issues in this release of Veritas Storage Foundation and High Availability.

### vradmin syncvol command compatibility with IPv6 addresses (2075307)

The `vradmin syncvol` command does not work with the compressed form of IPv6 addresses. In IPv6 environments, if you run the `vradmin syncvol` command and identify the target host using compressed form of the IPv6 address, the command fails with following error message:

```
# vradmin -s -full syncvol vol1 fe80::221:5eff:fe49:ad10:dg1:vol1
VxVM VVR vradmin ERROR V-5-52-420 Incorrect format for syncvol.
```

Also, if you run the `vradmin addsec` command and you specify the Secondary host using the compressed IPv6 address, the `vradmin syncvol` command also fails – even if you specify the target as `hostname`.

**Workaround:** When you use the `vradmin addsec` and `vradmin syncvol` commands, do not specify compressed IPv6 addresses; instead, use hostnames.

### RVGPrimary agent operation to start replication between the original Primary and the bunker fails during failback (2054804)

The RVGPrimary agent initiated operation to start replication between the original Primary and the bunker fails during failback – when migrating back to the original Primary after disaster recovery – with the error message:

```
VxVM VVR vxrlink ERROR V-5-1-5282 Error getting information from
remote host. Internal Error.
```

The issue applies to global clustering with a bunker configuration, where the bunker replication is configured using storage protocol. It occurs when the Primary comes back even before the bunker disk group is imported on the bunker host to initialize the bunker replay by the RVGPrimary agent in the Secondary cluster.

**Workaround:**

**To resolve this issue**

1. Before failback, make sure that bunker replay is either completed or aborted.
2. After failback, deport and import the bunker disk group on the original Primary.
3. Try the start replication operation from outside of VCS control.

### Bunker replay did not occur when the Application Service Group was configured on some of the systems in the Primary cluster, and ClusterFailoverPolicy is set to "AUTO" (2047724)

The time that it takes for a global cluster to fail over an application service group can sometimes be smaller than the time that it takes for VVR to detect the configuration change associated with the primary fault. This can occur in a bunkered, globally clustered configuration when the value of the `ClusterFailoverPolicy` attribute is `Auto` and the `AppGroup` is configured on a subset of nodes of the primary cluster.

This causes the RVGPrimary online at the failover site to fail. The following messages appear in the VCS engine log:

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname could not be
imported on bunker host hostname. Operation failed with error 256
and message VxVM VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote
server unreachable... Timestamp VCS ERROR V-16-2-13066 (hostname)
Agent is calling clean for resource(RVGPrimary) because the resource
is not up even after online completed.
```

**Workaround:**

**To resolve this issue**

◆ When the configuration includes a bunker node, set the value of the `OnlineRetryLimit` attribute of the RVGPrimary resource to a non-zero value.

### The RVGPrimary agent may fail to bring the application service group online on the new Primary site because of a previous primary-elect operation not being run or not completing successfully (2043831)

In a primary-elect configuration, the RVGPrimary agent may fail to bring the application service groups online on the new Primary site, due to the existence of previously-created instant snapshots. This may happen if you do not run the `ElectPrimary` command to elect the new Primary or if the previous `ElectPrimary` command did not complete successfully.

**Workaround:** Destroy the instant snapshots manually using the `vxrvg -g` *`dg`* `-P` *`snap_prefix`* `snapdestroy` *`rvg`* command. Clear the application service group and bring it back online manually.

### A snapshot volume created on the Secondary, containing a VxFS file system may not mount in read-write mode and performing a read-write mount of the VxFS file systems on the new Primary after a global clustering site failover may fail (1558257)

**Issue 1:**

When the `vradmin ibc` command is used to take a snapshot of a replicated data volume containing a VxFS file system on the Secondary, mounting the snapshot volume in read-write mode may fail with the following error:

```
UX:vxfs mount: ERROR: V-3-21268: /dev/vx/dsk/dg/snapshot_volume
is corrupted. needs checking
```

This happens because the file system may not be quiesced before running the `vradmin ibc` command and therefore, the snapshot volume containing the file system may not be fully consistent.

**Issue 2:**

After a global clustering site failover, mounting a replicated data volume containing a VxFS file system on the new Primary site in read-write mode may fail with the following error:

```
UX:vxfs mount: ERROR: V-3-21268: /dev/vx/dsk/dg/data_volume
is corrupted. needs checking
```

This usually happens because the file system was not quiesced on the original Primary site prior to the global clustering site failover and therefore, the file systems on the new Primary site may not be fully consistent.

**Workaround:** The following workarounds resolve these issues.

For issue 1, run the `fsck` command on the snapshot volume on the Secondary, to restore the consistency of the file system residing on the snapshot.

For example:

```
# fsck -V vxfs /dev/vx/dsk/dg/snapshot_volume
```

For issue 2, run the `fsck` command on the replicated data volumes on the new Primary site, to restore the consistency of the file system residing on the data volume.

For example:

```
# fsck -V vxfs /dev/vx/dsk/dg/data_volume
```

### In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon

Issue: After upgrading VVR to an IPv6-only environment in 6.0 release, vradmin commands may not work when a colon is specified in the RVG, data volume(s) and/or SRL name. It is also possible that after upgrading VVR to an IPv6-only environment, vradmin createpri may dump core when provided with RVG, volume and/or SRL names containing a colon in it.

**Workaround:** Make sure that colons are not specified in the volume, SRL and RVG names in the VVR configuration

### vradmin commands might fail on non-logowner node after logowner change (1810827)

When VVR is used for replicating shared disk groups in an SFCFS or SFRAC environment consisting of three or more nodes, a logowner change event might, in rare instances, render `vradmin` commands unusable on some or all of the cluster nodes. In such instances, the following message appears in the "Config Errors:" section of the output of the `vradmin repstatus` and `vradmin printrvg`commands:

```
vradmind not reachable on cluster peer
```

In addition, all other `vradmin` commands (except `vradmin printvol`) fail with the error:

```
"VxVM VVR vradmin ERROR V-5-52-488 RDS has configuration error related
to the master and logowner."
```

This is due to a defect in the internal communication sub-system, which will be resolved in a later release.

**Workaround:** Restart `vradmind` on all the cluster nodes using the following commands:

```
# /lib/svc/method/vras-vradmind.sh stop
# /lib/svc/method/vras-vradmind.sh start
```

### While vradmin commands are running, vradmind may temporarily lose heart beats (2162625, 2275444)

This issue may occasionally occur when you use `vradmin` commands to administer VVR. While the `vradmin` commands run, `vradmind` may temporarily lose heartbeats, and the commands terminate with the following error message:

```
VxVM VVR vradmin ERROR V-5-52-803 Lost connection to host host;
terminating command execution.
```

**Workaround:**

**To resolve this issue**

1. Depending on the application I/O workload and network environment, uncomment and increase the value of the `IPM_HEARTBEAT_TIMEOUT` variable in the `/etc/vx/vras/vras_env` on all the hosts of the RDS to a higher value. The following example increases the timeout value to 120 seconds.

   ```
   export IPM_HEARTBEAT_TIMEOUT
   IPM_HEARTBEAT_TIMEOUT=120
   ```

2. Restart `vradmind` on all the hosts of the RDS to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following on all the hosts of the RDS:

   ```
   # /etc/init.d/vras-vradmind.sh stop
   # /etc/init.d/vras-vradmind.sh start
   ```

### vxassist relayout removes the DCM (2162522)

If you perform a relayout that adds a column to a striped volume that has a DCM, the DCM is removed. There is no message indicating that this has happened. To replace the DCM, enter the following:

```
#vxassist -g diskgroup addlog vol logtype=dcm
```

### vxassist and vxresize operations do not work with layered volumes that are associated to an RVG (2162579)

This issue occurs when you try a resize operation on a volume that is associated to an RVG and has a striped-mirror layout.

**Workaround:**

**To resize layered volumes that are associated to an RVG**

1 Pause or stop the applications.

2 Wait for the RLINKs to be up to date. Enter the following:

```
# vxrlink -g diskgroup status rlink
```

3 Stop the affected RVG. Enter the following:

```
# vxrvg -g diskgroup stop rvg
```

4 Disassociate the volumes from the RVG. Enter the following:

```
# vxvol -g diskgroup dis vol
```

5 Resize the volumes. In this example, the volume is increased to 10 GB. Enter the following:

```
# vxassist -g diskgroup growto vol 10G
```

6 Associate the data volumes to the RVG. Enter the following:

```
# vxvol -g diskgroup assoc rvg vol
```

7 Start the RVG. Enter the following:

```
# vxrvg -g diskgroup start rvg
```

8 Resume or start the applications.

### Creating a primary diskgroup fails if there is no extra LUN to mirror the data change map (2478684)

Creating a primary diskgroup fails if there is no extra LUN to mirror the data change map (DCM), even if you have enough disk space.

**Workaround:** Add a LUN to the diskgroup before creating the primary diskgroup.

### verifydata operation fails when replicating between versions 5.1 and 6.0 (2360713)

When replicating in a cross-version VVR environment consisting of hosts running Storage Foundation 5.1 and hosts running Storage Foundation 6.0, the `vradmin verifydata` command fails with the following error:

```
VxVM VVR vxrsync ERROR V-5-52-2222 [from host]: VxVM in.vxrsyncd
ERROR V-5-36-2125 Server volume access error during [assign volids]
volume path: [/dev/vx/dsk/dg/snapshot_volume] reason: [this could be
because a target volume is disabled or an rlink associated with a
target volume is not detached during sync operation].
```

**Workaround:** There are two workarounds for this issue.

- Upgrade the hosts running Storage Foundation 5.1 to Storage Foundation 5.1SP1 or later and re-run the `vradmin verifydata` command.
- Follow the offline verification procedure in the "Verifying the data on the Secondary" section of the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*. This process requires ensuring that the secondary is up-to-date, pausing replication, and running the `vradmin syncrvg` command with the `-verify` option.

### Replication hang when VVR logowner is on CVM slave node (2405943)

When VVR is used for asynchronous replication in shared disk group environment, one of the nodes of the cluster at the primary site is chosen as the logowner. When the logowner node is on a node which is a slave node for the underlying CVM cluster, in the presence of heavy I/O from a node that is not the logowner, it is possible to get into a replication hang. This is due to an internal defect which will be fixed in later releases.

**Workaround:** Enable the PreOnline trigger of the RVGLogOwner agent so that the VVR logowner will always reside on the CVM master node. For the detailed procedure, refer to the RVGLogowner agent notes section in the *Veritas Cluster Server Bundled Agents Reference Guide*.

### Cannot relayout data volumes in an RVG from concat to striped-mirror (2162537)

This issue occurs when you try a relayout operation on a data volume which is associated to an RVG, and the target layout is a striped-mirror.

**Workaround:**

**To relayout a data volume in an RVG from concat to striped-mirror**

1. Pause or stop the applications.
2. Wait for the RLINKs to be up to date. Enter the following:

   ```
   # vxrlink -g diskgroup status rlink
   ```

3. Stop the affected RVG. Enter the following:

   ```
   # vxrvg -g diskgroup stop rvg
   ```

4. Disassociate the volumes from the RVG. Enter the following:

   ```
   # vxvol -g diskgroup dis vol
   ```

5. Relayout the volumes to striped-mirror. Enter the following:

   ```
   # vxassist -g diskgroup relayout vol layout=stripe-mirror
   ```

6. Associate the data volumes to the RVG. Enter the following:

   ```
   # vxvol -g diskgroup assoc rvg vol
   ```

7. Start the RVG. Enter the following:

   ```
   # vxrvg -g diskgroup start rvg
   ```

8. Resume or start the applications.

## Veritas Storage Foundation for Databases (SFDB) tools known issues

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

### SFDB commands do not work with the ZHS16GBK character set (2715323)

SFDB commands do not work if the character set of the Oracle database is set to ZHS16GBK. This occurs because SFDB commands are not supported with multi-byte character sets except AL32UTF8 and ZHS16GBK is a multi-byte character set.

There is no workaround for this issue.

### Database Storage Checkpoints created by using `dbed_ckptcreate` may not be visible after upgrading to 6.0 RP1 (2626248)

After upgrading from a 5.0 release to 6.0 RP1, the Database Storage Checkpoints created earlier using `dbed_ckptcreate` may not be migrated.

#### Workaround

Perform the following steps to make the old Database Storage Checkpoints visible.

**To resolve the issue**

1. Remove the new repository.
   - Examine the contents of the `/var/vx/vxdba/rep_loc`file to determine the location of the 6.0 RP1 repository.
   - Remove the `.sfae` directory specified as the `location` attribute.
2. Remove the repository location file: `/var/vx/vxdba/rep_loc`.
3. Create a symlink `/var/vx/vxdba/<SID>/.sfdb_rept` pointing to the `.sfdb_rept` directory created in the same location as the `.sfae` directory removed earlier.

   ```
   $ ln -s <location>/.sfdb_rept /var/vx/vxdba/<SID>/.sfdb_rept
   ```

   This step creates a symlink to the old repository.
4. Import repository data by running the `dbed_update` command.

   This step imports the data from the old repository.

The old Database Storage Checkpoints are now visible.

### Database Storage Checkpoint unmount may fail with device busy (2591463)

In some cases, when a database that is cloned using a Database Storage Checkpoint is shut down, an error similar to the following may occur:

```
SFAE Error:0457: Failed to unmount device
/dev/vx/dsk/datadg/datavol:Ckpt_1317707593_rw_1317708154.
Reason: VxFS returned error : umount: /tmp/clonedb/data: device is busy
```

### Workaround

As an Oracle user, force shut down the clone database if it is up and then retry the unmount operation.

### Incorrect error message if wrong host name is provided (2585643)

If you provide an incorrect host name with the `-r` option of `vxsfadm`, the command fails with an error message similar to one of the following:

```
FSM Error: Can't use string ("") as a HASH ref while "strict refs"
in use at /opt/VRTSdbed/lib/perl/DBED/SfaeFsm.pm line 776.
```

```
SFDB vxsfadm ERROR V-81-0609 Repository location  is invalid.
```

The error messages are unclear.

### Workaround

Provide the name of a host that has the repository database, with the `-r` option of `vxsfadm`.

### FlashSnap validate reports snapshot unsplittable (2534422)

The FlashSnap validation operation fails with the following error if the mirrors for data volumes and archive log volumes share the same set of disks:

```
SFAE Error:0642: Storage for diskgroup oradatadg is not splittable.
```

### Workaround

Ensure that snapshot plexes for data volumes and snapshot plexes for archive log volumes reside on separate set of disks.

### Attempt to use SmartTier commands fails (2332973)

The attempts to run SmartTier commands such as `dbdst_preset_policy` or`dbdst_file_move` fail with the following error:

```
fsppadm: ERROR: V-3-26551: VxFS failure on low level mechanism
with message - Device or resource busy
```

This error occurs if a sub-file SmartTier command such as `dbdst_obj_move` has been previously run on the file system.

There is no workaround for this issue. You cannot use file-based SmartTier and sub-file SmartTier simultaneously.

### `dbed_vmclonedb` ignores new clone SID value after cloning once (2580318)

After you have done FlashSnap cloning using a snapplan, any further attempts to create a clone from the same snapplan using the `dbed_vmclonedb` continue to use the original clone SID, rather than the new SID specified using the *new_sid* parameter.

This issue is also observed when you resynchronize the snapplan, take a snapshot again without specifying the new clone SID, and then try to clone with the new SID.

#### Workaround

You can use one of the following workarounds:

- After the snapshot is resynchronized, delete the snapplan using the `dbed_vmchecksnap -o remove` command. You can then use a new clone SID by creating a new snapplan, which may have the same name, and using the snapplan for taking more snapshots.
- Use the `vxsfadm` command to take the snapshot again and specify the clone SID with the snapshot operation so that the clone operation can be done with the new clone SID.

### Attempt to use certain names for tiers results in error (2581390)

If you attempt to use certain names for tiers, the following error message is displayed:

```
SFORA dbdst_classify ERROR V-81-6107 Invalid Classname BALANCE
```

This error occurs because the following names are reserved and are not permitted as tier names for SmartTier:

- BALANCE
- CHECKPOINT
- METADATA

#### Workaround

Use a name for SmartTier classes that is not a reserved name.

### User authentication fails (2579929)

The `sfae_auth_op -o auth_user` command, used for authorizing users, fails with the following error message:

```
SFDB vxsfadm ERROR V-81-0384 Unable to store credentials for <username>
```

Reattempting the operation fails with the following error message:

```
SFDB vxsfadm ERROR V-81-0372 AT broker failed to start:
```

The authentication setup might have been run with a strict umask value, which results in the required files and directories being inaccessible to the non-root users.

### Workaround

If you have not done authentication setup, set umask to a less strict value before running the `sfae_auth_op -o setup` or `sfae_auth_op -o import_broker_config` commands.

**To set umask to a less strict value**

- ◆ Use the command:

  ```
  # umask 022
  ```

If you have already done authentication setup, perform the following steps.

**To resolve the problem if you have already done authentication setup**

1. Shut down the authentication broker, if it is running.

   ```
   # /opt/VRTSdbed/at-broker/bin/sfaeatd.sh stop
   ```

2. Change the permissions for files and directories that are required to be readable by non-root users.

   ```
   # chmod o+r /etc/vx/vxdbed/admin.properties
   # chmod o+rx /var/vx/vxdba/auth/users
   # find /opt/VRTSdbed/at-broker -type d -exec chmod o+rx {} \;
   ```

## Clone operation failure might leave clone database in unexpected state (2512664)

If the clone operation fails, it may leave the clone database in an unexpected state. Retrying the clone operation might not work.

### Workaround

If retrying does not work, perform one the following actions depending on the point-in-time copy method you are using:

- For FlashSnap, resync the snapshot and try the clone operation again.
- For FileSnap and Database Storage Checkpoints, destroy the clone and create the clone again.

- For space-optimized snapshots, destroy the snapshot and create a new snapshot.

Contact Symantec support if retrying using the workaround does not succeed.

### FlashSnap resync fails if there is an existing space-optimized snapshot (2479901)

If you try a FlashSnap resync operation when there is an existing space-optimized snapshot, the resync operation fails with the following error:

```
Error: VxVM vxdg ERROR V-5-1-4597 vxdg join FS_oradg oradg failed
datavol_snp : Record already exists in disk group
archvol_snp : Record already exists in disk group
```

#### Workaround

Destroy the space-optimized snapshot first and then perform the FlashSnap resync operation.

### Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0 RP1 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 or 5.0MP3 to SFHA or SF for Oracle RAC 6.0 RP1.

When upgrading from SFHA version 5.0 or 5.0MP3 to SFHA 6.0 RP1 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbms3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File:  is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

#### Workaround

Before running `sfua_rept_migrate`, rename the startup script NO_S*vxdbms3 to S*vxdbms3.

### Clone command fails if PFILE entries have their values spread across multiple lines (1764885)

If you have a `log_archive_dest_1` in single line in the init.ora file, then `dbed_vmclonedb` will work but `dbed_vmcloneb` will fail if you put in multiple lines for `log_archive_dest_1`.

#### Workaround

There is no workaround for this issue.

## Veritas Cluster Server known issues

This section describes the known issues in this release of Veritas Cluster Server.

### Getting some error messages when rolling back the VRTSamf patch from VCS 6.0RP1 to 6.0 (2694345)

When rolling back the VRTSamf patch from VCS 6.0RP1 to 6.0, some error messages display.

There is no functionality impact.

**Workaround**: There is no workaround.

### NFS cluster I/O fails when storage is disabled

The I/O from the NFS clusters are saved on a shared disk or a shared storage. When the shared disks or shared storage connected to the NFS clusters are disabled, the I/O from the NFS Client fails and an I/O error occurs.

Workaround: If the application exits (fails/stops), restart the application.

### Operational issues for VCS

#### Connecting to the database outside VCS control using sqlplus takes too long to respond

Connecting to start the database outside VCS control, using sqlplus takes more than 10 minutes to respond after pulling the public network cable. [704069]

#### Some VCS components do not work on the systems where a firewall is configured to block TCP traffic

The following issues may occur if you install and configure VCS on systems where a firewall is installed:

- If you set up Disaster Recovery using the Global Cluster Option (GCO), the status of the remote cluster (cluster at the secondary site) shows as "initing".
- If you configure fencing to use CP server, fencing client fails to register with the CP server.
- Setting up trust relationships between servers fails.

Workaround:

- Ensure that the required ports and services are not blocked by the firewall. Refer to the *Veritas Cluster Server Installation Guide* for the list of ports and services used by VCS.
- Configure the firewall policy such that the TCP ports required by VCS are not blocked. Refer to your respective firewall or OS vendor documents for the required configuration.

## Issues related to the VCS engine

### Extremely high CPU utilization may cause HAD to fail to heartbeat to GAB

When CPU utilization is very close to 100%, HAD may fail to heartbeat to GAB. [1818687]

### The hacf -cmdtocf command generates a broken main.cf file

The `hacf -cmdtocf` command used with the `-dest` option removes the include statements from the types files. [1728738]

Workaround: Add include statements in the main.cf files that are generated using the `hacf -cmdtocf` command.

### Agent framework can reject `hares -action` command

When a probed resource is disabled and later enabled then, the agent framework can reject `hares -action` command till the agent successfully monitors the resource.

### VCS fails to validate processor ID while performing CPU Binding [2441022]

If you specify an invalid processor number when you try to bind HAD to a processor on a remote system, HAD does not bind to any CPU. However, the command displays no error to indicate that the specified CPU does not exist. The error is logged on the node where the binding has failed and the values are reverted to default.

Workaround: Symantec recommends that you modify CPUBinding from the local system.

### Trigger does not get executed when there is more than one leading or trailing slash in the triggerpath [2368061]

The path specified in TriggerPath attribute must not contain more than one leading or trailing '\' character.

Workaround: Remove the extra leading or trailing '\' characters from the path.

### Service group is not auto started on the node having incorrect value of EngineRestarted [2397532]

When HAD is restarted by `hashadow` process, the value of EngineRestarted attribute is temporarily set to 1 till all service groups are probed. Once all service groups are probed, the value is reset. If HAD on another node is started at roughly the same time, then it is possible that it does not reset the value of EngineRestarted attribute. Therefore, service group is not auto started on the new node due to mismatch in the value of EngineRestarted attribute.

Workaround: Restart VCS on the node where EngineRestarted is set to 1.

### Group is not brought online if top level resource is disabled [2486476]

If the top level resource which does not have any dependancy is disabled then the other resources do not come online and the following message is displayed:

```
VCS NOTICE V-16-1-50036 There are no enabled
 resources in the group cvm to online
```

Workaround: Online the child resources of the topmost resource which is disabled.

### NFS resource goes offline unexpectedly and reports errors when restarted [2490404]

VCS does not perform resource operations, such that if an agent process is restarted multiple times by HAD, only one of the agent process is valid and the remaining processes get aborted, without exiting or being stopped externally. Even though the agent process is running, HAD does not recognize it and hence does not perform any resource operations.

Workaround: Forcefully stop the agent process.

### Parent group does not come online on a node where child group is online [2489053]

This happens if the AutostartList of parent group does not contain the node entry where the child group is online.

Workaround: Bring the parent group online by specifying the name of the system then use the `hargp -online [parent group] -any` command to bring the parent group online.

### Cannot modify temp attribute when VCS is in LEAVING state [2407850]

An `ha` command to modify a temp attribute is rejected if the local node is in a LEAVING state.

Workaround: Execute the command from another node or make the configuration read-write enabled.

### If secure and non-secure WAC are connected the engine_A.log receives logs every 5 seconds [1539646]

Two WACs in GCO must always be started either in secure or non-secure mode. The secure and non-secure WAC connections cause log messages to be sent to engine_A.log file.

Workaround: Make sure that WAC is running in either secure mode or non-secure mode on both the clusters in GCO.

### Oracle group fails to come online if Fire Drill group is online on secondary cluster [2556835]

If a parallel global service group faults on the local cluster and does not find a failover target in the local cluster, it tries to failover the service group to the remote cluster. However, if the firedrill for the service group is online on a remote cluster, offline local dependency is violated and the global service group is not able to failover to the remote cluster.

Workaround: Offline the Firedrill service group and online the service group on a remote cluster.

### POSTONLINE and POSTOFFLINE triggers are not enabled by default [2567387]

Before VCS 6.0, POSTONLINE and POSTOFFLINE triggers were enabled by default, so the triggers got executed whenever a service group came online. In VCS 6.0, you must explicitly enable the POSTONLINE and POSTOFFLINE triggers whenever you upgrade to VCS 6.0.

**Alternatively, if you want the triggers to execute after the upgrade:**

1. Before upgrade, set vcs_start = 0 in /etc/default/vcs

   so that HAD does not start after the upgrade.
2. Upgrade the existing VCS to VCS 6.0.
3. Set vcs_start = 1 in /etc/default/vcs
4. Start VCS on each node using `hastart`.
5. Set TriggersEnabled in main.cf for required groups as ollows:

   ```
   TriggersEnabled @<systemname>={POSTONLINE, POSTOFFLINE}
   ```

Example of trigger behavior:

```
group scriptfileonoff (
        SystemList = { vcssx235 = 0, vcssx236 = 1 }
        AutoStartList = { vcssx235, vcssx236 }
        TriggersEnabled @vcssx235 = { POSTONLINE }
```

```
        )
        MyFileOnOff MFileOnOff (
                PathName = "/tmp/mf1"
                )
        MyFileOnOff MFileOnOff1 (
                PathName = "/tmp/mf2"
```

### Two CmdServer instances seen running on a node [2399292]

You may see two instances of CmdServer running on a node. One of these using IPv4 and the other IPv6.

This does not impact functionality in any way.

Workaround: No workaround.

### Service group may fail to come online after a flush and a force flush operation [2616779]

A service group may fail to come online after flush and force flush operations are executed on a service group where offline operation was not successful.

Workaround: If the offline operation is not successful then use the force flush commands instead of the normal flush operation. If a normal flush operation is already executed then to start the service group use `-any` option.

## Issues related to the bundled agents

### VCS resources may time out if NFS server is down [2129617]

The VCS resources may time out if the server NFS mounted file system and the NFS server is down or inaccessible. This behavior is exclusive to AIX platform.

Workaround: You must unmount the NFS mounted file system to restore the cluster to sane condition.

### MultiNICB resource may show unexpected behavior with IPv6 protocol [2535952]

When using IPv6 protocol, set the LinkTestRatio attribute to 0. If you set the attribute to another value, the MultiNICB resource may show unexpected behavior.

Workaround: Set the LinkTestRatio attribute to 0.

### Application agent cannot handle a case with user as root, envfile set and shell as csh [2584285]

Application agent does not handle a case when the user is root, envfile is set, and shell is csh. The application agent uses the `system` command to execute the `Start`/`Stop`/`Monitor`/`Clean` Programs for the root user. This executes

Start/Stop/Monitor/Clean Programs in sh shell, due to which there is an error when root user has csh shell and EnvFile is written accordingly.

Workaround: Do not set csh as shell for root user. Use sh as shell for root instead.

### IMF registration fails for Mount resource if the configured MountPoint path contains spaces [2442598]

If the configured MountPoint of a Mount resource contains spaces in its path, then the Mount agent can online the resource correctly, but the IMF registration for ONLINE monitoring fails. This is due to the fact that the AMF driver does not support spaces in the path. Leading and trailing spaces are handled by the Agent and IMF monitoring can be done for such resources.

Workaround: Symantec recommends to turn off the IMF monitoring for a resource having spaces in its path. For information on disabling the IMF monitoring for a resource, refer to Veritas Cluster Server Administrator's Guide.

### Bringing the LPAR resource offline may fail [2418615]

Bringing the LPAR resource offline may fail with the following message in the engine_A.log file.

```
<Date Time> VCS WARNING V-16-10011-22003 <system_name>
LPAR:<system_name>:offline:Command failed to run on MC
<hmc_name> with error HSCL0DB4 An Operating System
Shutdown can not be performed because the operating system image
running does not support remote execution of this task from the HMC.
This may be due to problem in communication with
MC <hmc_name>
```

This is due to RMC failure between HMC and management LPAR. Since the LPAR could not be shutdown gracefully in offline, the LPAR is shutdown forcefully in the clean call, hence it shows as Faulted.

Workaround: In order to recycle the RSCT deamon for LPAR and HMC, refer the *Veritas Storage Foundation™ and High Availability Solutions Virtualization Guide*.

### LPAR agent may not show the correct state of LPARs [2418615, 2425990]

When the Virtual I/O server (VIOS) gets restarted, LPAR agent may not get the correct state of the resource. In this case, the LPAR agent may not show the correct state of the LPAR.

Workaround: Restart the management LPAR and all the managed LPARs that depend on the VIOS.

### RemoteGroup agent does not failover in case of network cable pull [2588807]

A RemoteGroup resource with ControlMode set to OnOff may not fail over to another node in the cluster in case of network cable pull. The state of the RemoteGroup resource becomes UNKNOWN if it is unable to connect to a remote cluster.

Workaround:

- Connect to the remote cluster and try taking offline the RemoteGroup resource.
- If connection to the remote cluster is not possible and you want to bring down the local service group, change the ControlMode option of the RemoteGroup resource to MonitorOnly. Then try taking offline the RemoteGroup resource. Once the resource is offline, change the ControlMode option of the resource to OnOff.

### Concurrency violation in the service group [2555306]

Concurrency violation and data corruption of a Volume resource may occur, if storage connectivity is lost or all paths under VxDMP are disabled and PanicSystemOnDGLoss is set to 0

This happens when:

- In a cluster environment/configuration, if cluster wide UseFence attribute is set to SCSI3 and service group contains Volume resource and DiskGroup resource with the PanicSystemOnDGLoss attribute set to 0 (zero).
- If storage connectivity is lost or all paths under VxDMP are disabled, VCS fails over the service group. If storage connectivity is restored on the node on which the service group was faulted and DG is not deported manually, then volume may get started if disk group is not deported during the service group failover. So volume resource shows state as online on both the nodes and thus cause concurrency violation. This may lead to data corruption.

Workaround: Ensure that the disk group is deported soon after storage connectivity is restored.

You are recommended to always configure Volume resource whenever Disk group resources is configured and set the attribute PanicSystemOnDGLoss to 1 or 2 as per requirement.

### Coordpoint agent remains in faulted state [2555191]

The Coordpoint agent remains in faulted state because it detects `rfsm` to be in replaying state.

Workaround: Clear the fault and reconfigure fencing.

### No IPv6 support for NFS [2022174]

IPv6 is not supported for NFS.

Workaround: No workaround.

### MemCPUAllocator agent fails to come online if DLPAR name and hostname do not match [2407671]

If hostname of the DLPAR and name of DLPAR as seen from HMC are different, the MemCPUAllocator agent is unable to provide CPU or memory to the DLPAR.

Workaround: Change the name of DLPAR from HMC to match the hostname.

### VCS does not monitor applications inside an already existing WPAR [2494532]

If a WPAR is already present on the system at the time of VCS installation, and this WPAR or an application running inside this WPAR needs to be monitored using VCS, then VCS does not monitor the application running in that WPAR. This is because the VCS packages/files are not visible inside that WPAR.

Workaround: Run `syncwpar` command for that WPAR. This makes the VCS packages/files visible inside the WPAR and VCS can then monitor the applications running inside the WPAR.

### The hawparsetup.pl script does not check the service group status [2523171]

If the service group in which WPAR resource needs to be added already exists and is not in OFFLINE state, `hawparsetup.pl` script does not modify the ContainerInfo attribute for the system on which it is not completely OFFLINE.The hawparsetup.pl scrip, therefore it does not check whether the service group passed to it is completely OFFLINE or not if it already exists.

Workaround: The `hawparsetup.pl` script does not check whether the service group passed to it is completely OFFLINE or not if it already exists.

### NFS client reports I/O error because of network split brain [2564517]

When network split brain occurs, the failing node may take some time to panic. Thus, the service group on the failover node may fail to come online, as some of the resources (like IP resource) are still online on the failing node or disk group on the failing node may get disabled but IP resource on the same node continues to be online.

**Workaround: Configure the preonline trigger for the service group containing DiskGroup resouce on each system in the service group:**

1. Copy the preonline_ipc trigger from `/opt/VRTSvcs/bin/sample_triggers/VRTSvcs` to `/opt/VRTSvcs/bin/triggers/preonline/` as T0preonline_ipc.

   ```
   # cp /opt/VRTSvcs/bin/sample_triggers/VRTSvcs/preonline_ipc
    /opt/VRTSvcs/bin/triggers/preonline/T0preonline_ipc
   ```

2. Enable PREONLINE trigger for the service group.

   ```
   # hagrp -modify <group_name> TriggersEnabled PREONLINE
    -sys <node_name>
   ```

### Some agents may fail to come online after full upgrade to VCS 6.0 if they were online before the upgrade [2618482]

Resources of type NFSRestart and DNS do not come online automatically after a full upgrade to VCS 6.0 if they were previously online

Workaround: Online the resources manually after the upgrade, if they were online previously.

### Coexistence of Live Partition Mobility (LPM) and VCS failover of managed LPAR

Coexistence of LPM and VCS failover of manged LPAR may cause an issue. If you plan to do LPM on a managed LPAR, make sure you see the *Live partition mobility of managed LPARs* section in *Veritas Storage Foundation and High Availability Solutions Virtualization Guide* for AIX for the correct information.

### Error messages for wrong HMC user and HMC name do not communicate the correct problem

The wrong HMC user and wrong HMC name errorsd are not reflective of the correct problem. If you see the following errors in engine_A.log for LPAR resource, it means wrong HMC user:

```
Permission denied, please try again
Permission denied, please try again
```

If you see the following errors in engine_A.log for LPAR resource, it means wrong HMC name:

```
ssh: abc: Hostname and service name
not provided or found.
```

You must see the applicationha_utils.log file to confirm the same.

## Issues related to the VCS database agents

### Health check monitoring does not work with VCS agent for Oracle [2101570, 1985055]

The health check monitoring in Oracle agent for VCS does not work due to incompatibility of the health check APIs provided by Oracle.

Resolution: Disable health check monitoring by setting the MonitorOption attribute to 0 (zero).

### Intentional Offline does not work for VCS agent for Oracle [1805719]

Due to issues with health check monitoring, Intentional Offline does not work for VCS agent for Oracle.

### Make sure that the ohasd has an entry in the init scripts [1985093]

Make sure that the ohasd process has an entry in the init scripts so that when the process is killed or the machine is rebooted, this automatically restarts the process.

Workaround: Respawn off the `ohasd` process. Add the `ohasd process` in the`/etc/inittab` file to ensure that this process is automatically restarted when killed or the machine is rebooted.

### The ASMInstAgent does not support having pfile/spfile for the ASM Instance on the ASM diskgroups

The ASMInstAgent does not support having pfile/spfile for the ASM Instance on the ASM diskgroups.

Workaround:

Have a copy of the pfile/spfile in the default $GRID_HOME/dbs directory to make sure that this would be picked up during the ASM Instance startup.

### VCS agent for ASM: Health check monitoring is not supported for ASMInst agent

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

### NOFAILOVER action specified for certain Oracle errors

The Veritas High Availability agent for Oracle provides enhanced handling of Oracle errors encountered during detailed monitoring. The agent uses the reference file oraerror.dat, which consists of a list of Oracle errors and the actions to be taken.

See the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide* for a description of the actions.

Currently, the reference file specifies the NOFAILOVER action when the following Oracle errors are encountered:

```
ORA-00061, ORA-02726, ORA-6108, ORA-06114
```

The NOFAILOVER action means that the agent sets the resource's state to OFFLINE and freezes the service group. You may stop the agent, edit the oraerror.dat file, and change the NOFAILOVER action to another action that is appropriate for your environment. The changes go into effect when you restart the agent.

### ASM instance does not unmount VxVM volumes after ASMDG resource is offline

In configurations where ASMInstance resource is part of a separate parallel service group, the ASM instance does not unmount the volumes even after the ASMDG resource is taken offline. Therefore, the Volume resource cannot be taken offline. This issue occurs when you use VxVM volumes as ASM disk groups. [918022]

Workaround: Configure the ASMInstance resource as part of the failover service group where ASMDG resource is configured.

### IMF registration fails if sybase server name is given at the end of the configuration file [2365173]

AMF driver supports a maximum of 80 characters of arguments. In order for AMF to detect the start of the Sybase process, the Sybase server name must occur in the first 80 characters of the arguments.

Workaround: Must have the server name, -sSYBASE_SERVER, as the first line in the configuration file: `ASE-15_0/install/RUN_SYBASE_SERVER`.

## Issues related to the agent framework

### Agent framework cannot handle leading and trailing spaces for the dependent attribute

Agent framework does not allow spaces in the target resource attribute name of the dependent resource.

Workaround: Do not provide leading and trailing spaces in the target resource attribute name of the dependent resource.

### The agent framework does not detect if service threads hang inside an entry point [1511211]

In rare cases, the agent framework does not detect if all service threads hang inside a C entry point. In this case it may not cancel them successfully.

Workaround: If the service threads of the agent are hung, send a kill signal to restart the agent. Use the following command: `kill -9` *`hung agent's pid`*. The `haagent -stop` command does not work in this situation.

### IMF related error messages while bringing a resource online and offline [2553917]

For a resource registered with AMF, if you run `hagrp -offline` or `hagrp -online` explicitly or through a collective process to offline or online the resource respectively, the IMF displays error messages in either case.

The errors displayed is an expected behavior and it does not affect the IMF functionality in any manner.

Workaround: No workaround.

## Issues related to global clusters

### The engine log file receives too many log messages on the secure site in global cluster environments [1539646]

When the WAC process runs in secure mode on one site, and the other site does not use secure mode, the engine log file on the secure site gets logs every five seconds.

Workaround: The two WAC processes in global clusters must always be started in either secure or non-secure mode. The secure and non-secure WAC connections will flood the engine log file with the above messages.

### Application group attempts to come online on primary site before fire drill service group goes offline on the secondary site (2107386)

The application service group comes online on the primary site while the fire drill service group attempts to go offline at the same time, causing the application group to fault.

**Workaround:** Ensure that the fire drill service group is completely offline on the secondary site before the application service group comes online on the primary site.

## Issues related to LLT

This section covers the known issues related to LLT in this release.

### LLT may fail to make connections with LLT on peer nodes in virtual environment (2343451/2376822)

After you upgrade from 5.0 MP3 or earlier releases to version 6.0, LLT may fail to make connections with LLT on the peer nodes in AIX virtual environment.

This is a known IBM VIOS issue. Install APAR IV00776 on your VIOS server. Without this fix, VIOS fails to handle new LLT packet header and drops packets.

Workaround: Disable the largesend attribute of the SEA adapter. Check the properties of the SEA adapter (on which the virtual links are configured under LLT maps) on each VIOS using the following command:

```
#lsattr -El SEA
```

If the largesend is set to 1, then set it to 0 using the following command:

```
#chdev -l SEA -a largesend=0
```

### LLT port stats sometimes shows recvcnt larger than recvbytes (1788315)

With each received packet, LLT increments the following variables:

- recvcnt (increment by one for every packet)
- recvbytes (increment by size of packet for every packet)

Both these variables are integers. With constant traffic, recvbytes hits and rolls over MAX_INT quickly. This can cause the value of recvbytes to be less than the value of recvcnt.

This does not impact the LLT functionality.

### LLT may incorrectly declare port-level connection for nodes in large cluster configurations (1809827)

When ports get registered and unregistered frequently on the nodes of the cluster, LLT may declare that a port-level connection exists with another peer node. This occurs in some corner cases even though a port is not even registered on the peer node.

## Issues related to GAB

This section covers the known issues related to GAB in this release.

### While deinitializing GAB client, "gabdebug -R GabTestDriver" command logs refcount value 2 (2536373)

After you unregister the gtx port with `-nodeinit` option, the `gabconfig -C` command shows refcount as 1. But when forceful `deinit` option (`gabdebug -R GabTestDriver`) is run to deinitialize GAB client, then a message similar to the following is logged.

```
GAB INFO V-15-1-20239
Client GabTestDriver with refcount 2 forcibly deinited on user request
```

The `refcount` value is incremented by 1 internally. However, the refcount value is shown as 2 which conflicts with the `gabconfig -C` command output.

### Cluster panics during reconfiguration (2590413)

While a cluster is reconfiguring, GAB broadcast protocol encounters a race condition in the sequence request path. This condition occurs in an extremely narrow window which eventually causes the GAB master to panic.

## Issues related to I/O fencing

This section covers the known issues related to I/O fencing in this release.

### CP server repetitively logs unavailable IP addresses (2530864)

If coordination point server (CP server) fails to listen on any of the IP addresses that are mentioned in the `vxcps.conf` file or that are dynamically added using the command line, then CP server logs an error at regular intervals to indicate the failure. The logging continues until the IP address is bound to successfully.

```
CPS ERROR V-97-51-103 Could not create socket for host
10.209.79.60 on port 14250
CPS ERROR V-97-1400-791 Coordination point server could not
open listening port = [10.209.79.60]:14250
Check if port is already in use.
```

Workaround: Remove the offending IP address from the listening IP addresses list using the `rm_port` action of the `cpsadm` command.

See the *Veritas Storage Foundation and High Availability Administrator's Guide* for more details.

### Fencing port b is visible for few seconds even if cluster nodes have not registered with CP server (2415619)

Even if the cluster nodes have no registration on the CP server and if you provide coordination point server (CP server) information in the `vxfenmode` file of the cluster nodes, and then start fencing, the fencing port b is visible for a few seconds and then disappears.

Workaround: Manually add the cluster nodes' and users' information to the CP server to resolve this issue. Alternatively, you can use installer as the installer adds cluster nodes' and users' information to the CP server during configuration.

### The cpsadm command fails if LLT is not configured on the application cluster (2583685)

The `cpsadm` command fails to communicate with the coordination point server (CP server) if LLT is not configured on the application cluster node where you run the `cpsadm` command. You may see errors similar to the following:

```
# cpsadm -s 10.209.125.200 -a ping_cps
CPS ERROR V-97-1400-729 Please ensure a valid nodeid using
environment variable
CPS_NODEID
CPS ERROR V-97-1400-777 Client unable to communicate with CPS.
```

However, if you run the `cpsadm` command on the CP server, this issue does not arise even if LLT is not configured on the node that hosts CP server. The `cpsadm` command on the CP server node always assumes the LLT node ID as 0 if LLT is not configured.

According to the protocol between the CP server and the application cluster, when you run the `cpsadm` on an application cluster node, `cpsadm` needs to send the LLT node ID of the local node to the CP server. But if LLT is unconfigured temporarily, or if the node is a single-node VCS configuration where LLT is not configured, then the `cpsadm` command cannot retrieve the LLT node ID. In such situations, the `cpsadm` command fails.

Workaround: Set the value of the `CPS_NODEID` environment variable to 255. The `cpsadm` command reads the `CPS_NODEID` variable and proceeds if the command is unable to get LLT node ID from LLT.

### In absence of cluster details in CP server, VxFEN fails with pre-existing split-brain message (2433060)

When you start server-based I/O fencing, the node may not join the cluster and prints error messages in logs similar to the following:

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1043
Detected a preexisting split brain. Unable to join cluster.
```

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
operation failed.
CPS ERROR V-97-1400-446 Un-authorized user cpsclient@node1,
domaintype vx; not allowing action
```

The vxfend daemon on the application cluster queries the coordination point server (CP server) to check if the cluster members as seen in the GAB membership are registered with the CP server. If the application cluster fails to contact the CP server due to some reason, then fencing cannot determine the registrations on the CP server and conservatively assumes a pre-existing split-brain.

Workaround: Before you attempt to start VxFEN on the application, ensure that the cluster details such as cluster name, UUID, nodes, and privileges are added to the CP server.

### The vxfenswap utility does not detect failure of coordination points validation due to an RSH limitation (2531561)

The `vxfenswap` utility runs the `vxfenconfig -o modify` command over RSH or SSH on each cluster node for validation of coordination points. If you run the `vxfenswap` command using RSH (with the `-n` option), then RSH does not detect the failure of validation of coordination points on a node. From this point, `vxfenswap` proceeds as if the validation was successful on all the nodes. But, it fails at a later stage when it tries to commit the new coordination points to the VxFEN driver. After the failure, it rolls back the entire operation, and exits cleanly with a non-zero error code. If you run `vxfenswap` using SSH (without the `-n` option), then SSH detects the failure of validation of coordination of points correctly and rolls back the entire operation immediately.

Workaround: Use the `vxfenswap` utility with SSH (without the `-n` option).

### Fencing does not come up on one of the nodes after a reboot (2573599)

If VxFEN unconfiguration has not finished its processing in the kernel and in the meantime if you attempt to start VxFEN, you may see the following error in the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1007 Vxfen already configured
```

However, the output of the `gabconfig -a` command does not list port b. The `vxfenadm -d` command displays the following error:

```
VXFEN vxfenadm ERROR V-11-2-1115 Local node is not a member of cluster!
```

Workaround: Start VxFEN again after some time.

### The cpsadm command fails after upgrading CP server to 6.0 in secure mode (2478502)

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTSat fileset is not removed from the system, the `cpsadm` command loads the old security libraries present on the system. As the installer runs the `cpsadm` command on the CP server to add or upgrade the SFHA cluster (application cluster), the installer also fails.

Workaround : Perform the following steps on all the nodes of the CP server:

- Rename cpsadm to cpsadmbin.

  ```
  # mv /opt/VRTScps/bin/cpsadm /opt/VRTScps/bin/cpsadmbin
  ```

- Create a file `/opt/VRTScps/bin/cpsadm` with the following content:

  ```
  #!/bin/sh
  EAT_USE_LIBPATH=”/opt/VRTScps/lib”
  export EAT_USE_LIBPATH
  /opt/VRTScps/bin/cpsadmbin "$@"
  ```

- Provide the following permissions to the new file:

  ```
  # chmod 755 /opt/VRTScps/bin/cpsadm
  ```

### Server-based fencing comes up incorrectly if default port is not mentioned (2403453)

When you configure fencing in customized mode and do no provide default port, fencing comes up. However, the `vxfenconfig -l` command output does not list the port numbers.

Workaround: Retain the "port=<port_value>" setting in the `/etc/vxfenmode` file, when using customized fencing with atleast one CP server. The default port value is 14250.

### Secure CP server does not connect from localhost using 127.0.0.1 as the IP address (2554981)

The `cpsadm` command does not connect to the secure CP server on the localhost using 127.0.0.1 as the IP address

Workaround: Connect the secure CP server using any of the virtual IPs that is configured with the CP server and is plumbed on the local node.

### Unable to customize the 30-second duration (2551621)

When the vxcpserv process is not able to bind to an IP address during startup, it attempts to bind to that IP address at an interval of 30 seconds. This interval is not configurable.

Workaround: No workaround.

### NIC resource gets created with incorrect name while configuring CPSSG with the configure_cps.pl script (2585229)

The name of the NIC resource created by the `configure_cps.pl` script does not come out correct when, for example, $m^{th}$ VIP is mapped to $n^{th}$ NIC and every m is not equal to n. In this case, although CPSSG continues to function without any problem, when you unconfigure CPSSG using `configure_cps.pl`, it fails.

Workaround: To unconfigure CPSSG, you must remove the CPSSG configuration from the VCS configuration.

## Issues related to Intelligent Monitoring Framework (IMF)

### Registration error while creating a Firedrill setup [2564350]

While creating the Firedrill setup using the `Firedrill setup` utility, VCS encounters the following error:

```
AMF amfregister ERROR V-292-2-167
Cannot register mount offline event
```

During Firedrill operations, VCS may log error messages related to IMF registration failure in the engine log. This happens because in the firedrill service group, there is a second CFSMount resource monitoring the same MountPoint through IMF. Both the resources try to register for online/offline events on the same MountPoint and as a result, registration of one fails.

Workaround: No workaround.

### Pearl errors seen while using haimfconfig command

Pearl errors seen while using `haimfconfig` command:

```
Pearl errors seen while using haimfconfig command
```

This error is due to the absolute path specified in main.cf for type-specific configuration files. Currently, `haimfconfig` does not support absolute path for type-specific configuration file in main.cf.

Wrokaround: Replace the actual path with the actual file name and copy the file from its absolute location to `/etc/VRTSvcs/conf/config` directory.

For example, if OracleTypes.cf is included in main.cf as:

```
include "/etc/VRTSagents/ha/conf/Oracle/OracleTypes.cf"
```

It should be replaced as follows in main.cf:

```
include "OracleTypes.cf"
```

## Issues related to the Cluster Manager (Java Console)

This section covers the issues related to the Cluster Manager (Java Console).

### Some Cluster Manager features fail to work in a firewall setup [1392406]

In certain environments with firewall configurations between the Cluster Manager and the VCS cluster, the Cluster Manager fails with the following error message:

```
V-16-10-13 Could not create CmdClient. Command Server
may not be running on this system.
```

Workaround: You must open port 14150 on all the cluster nodes.

### VCS Cluster Manager (Java Console) does not encrypt Sybase and SybaseBk agent passwords [2379510]

If `isvcsagencrypt` flag is set to True in Sybase.xml and SybaseBk.xml files, the attribute values get encrypted. However, the password attributes of Sybase and SybaseBk agents do not have the `isvcsagentcrypt` flag set to True in Sybase.xml and SybaseBk.xml files.

Workaround: Sybase and SybaseBk agents are modified to encrypt the password by default. As a result, you need not encrypt passwords if you use the VCS Cluster Manager (Java Console) to configure attributes.

## Issues related to Virtual Business Services (VBS)

### Child service group fault missed in case of firm dependency (2673289)

In case of firm dependency, if the child service group has faulted, the parent service group is brought offline. Subsequently when the child service group recovers, the parent service group is brought online.

However, if the child service group faults again before the parent has recovered, this new child fault event is missed. This could happen if the parent recovery takes time. As a result, the parent service group may remain online.

**Workaround**: There is no workaround.

### VBS version mismatch across tiers may affect its functionality (2695412)

If you have different versions of VBS across various tiers, VBS may not function as expected. Therefore, you must ensure that when you upgrade a particular tier of VBS, you must also upgrade the other tiers to the same version, such that VBS version across all tiers is the same.

**Workaround**: Maintain the same VBS versions on all tiers.

### Fault propagation for Virtual Business Services with shared service groups and different controllers [2407832]

Fault propagation may not work for certain configurations having shared service groups and distinct controllers.

Workaround: No workaround.

### Virtual Business Services fail to start if a participating service group has multiple children with the LOCAL FIRM dependency type [2490098]

Virtual Business Services fail to start if a participating service group has multiple children with the LOCAL FIRM dependency type. This occurs because Vertias Cluster Server (VCS) does not support propagating dependencies.

**Workaround:** Pull the dependent VCS groups into the Virtual Business Services without any dependencies. The Virtual Business Services will recognize the VCS dependencies and treat them as soft Virtual Business Services dependencies.

# Veritas Storage Foundation and High Availability known issues

For known issues of Veritas Storage Foundation and High Availability, refer to the section called “Veritas Storage Foundation known issues” and the section called “Veritas Cluster Server known issues”.

# Veritas Storage Foundation Cluster File System High Availability known issues

This section describes the known issues in this release of Veritas Storage Foundation Cluster File System High Availability.

## Veritas Storage Foundation Cluster File System High Availability known issues

This section describes the known issues in this release of Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

### CFS commands might hang when run by non-root (2403263)

The CFS commands might hang when run by non-root.

#### Workaround

**To resolve this issue**

◆ Use halogin command to save the authentication information before running any CFS commands on a non-root sessions.

  When you run the halogincommand, VCS stores encrypted authentication information in the user's home directory.

### The installer may fail to mount some share disk groups (2167226)

The installer fails to mount some share disk goups if its name is a substring of other disk groups.

### Workaround

You need to manually add those share disk groups to the newly added nodes. Or avoid naming your share disk groups that could be substring of others.

### You sometimes receive shell error messages (2172138)

You sometimes receive shell error messages while adding a node into an existing SFHA cluster. The following is a sample of the shell error message that you can receive:

```
sh[2]: sw:  not found
```

You can safely ignore these error messages.

### Miscalculated file set usage (2123429)

When file set quotas are enabled, it may be possible for VxFS to get into a state where it thinks a very large number of blocks are allocated to Storage Checkpoints. This issue can be seen using the fsckptadm command:

```
# fsckptadm getquotalimit /mnt1
Filesystem   hardlimit   softlimit   usage   action_flag
/mnt1        10000       10000       18446744073709551614
```

This could cause writes to Storage Checkpoints to fail. It could also trigger the removal of removable Storage Checkpoints.

### Workaround

If this occurs, disabling and re-enabling file set quotas causes VxFS to recalculate the number of blocks used by Storage Checkpoints:

```
# fsckptadm quotaoff /mnt1
# fsckptadm quotaon /mnt1
# fsckptadm getquotalimit /mnt1
Filesystem   hardlimit   softlimit   usage   action_flag
/mnt1        10000       10000       99
```

### The cfsmntadm add command may fail with no errors (2169538)

The `cfsmntadm add` command fails, if one host name is a substring of another host name in the list.

**Note:** VOM is affected by this issue when adding a CFS mount to a cluster that has systems with host names that are substrings of each other.

#### Workaround

Run the `cfsmntadm` command with the "`all=`" option on one of the nodes in the CFS cluster to add the cfsmounts to all nodes.

### Multiple CFSmount resources are in a single service group they may not all come online after a reboot (2164670)

In some cases when multiple CFSmount resources are in a single service group, they all may not come online after a reboot. You will need to manually bring them online after a reboot.

#### Workaround

Create a resource dependency between the various CFSmount resources.

### CVMVolDg agent may fail to deport CVM disk group

The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the `CVMDeportOnOffline` attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported.

**Workaround:** If multiple CVMVolDg resources are configured for a shared disk group, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources.

### NFS issues with VxFS Storage Checkpoint (1974020)

NFS clients mounting VxFS Storage Checkpoints that are NFS-exported by SFHA cluster nodes using a Virtual IP may receive the following error message upon Virtual IP failover:

```
Stale NFS file handle
```

This is a result of major numbers of VxFS Storage Checkpoints not necessarily being the same on all SFHA cluster nodes.

Workaround: There is no workaround at this time.

### Panic due to null pointer de-reference in vx_bmap_lookup() (2582232)

A null pointer dereference in the `vx_bmap_lookup`() call can cause a panic.

**Workaround:** Resize the file system with the `fsadm` command from the primary node of the cluster.

### Multiple system panics upon unmounting a CFS file system (2107152)

There is a system panic when you unmount a `mntlock`-protected VxFS file system, if that device is duplicate mounted on different directories.

**Workaround:** There is no workaround for this issue.

### "Configuration must be ReadWrite : Use haconf -makerw" error message appears in VCS engine log when hastop -local is invoked (2609137)

A message similar to the following example appears in the `/var/VRTSvcs/log/engine_A.log` log file when you run the `hastop -local` command on any system in a SFHA cluster that has `CFSMount`resources:

```
2011/11/15 19:09:57 VCS ERROR V-16-1-11335 Configuration must be
ReadWrite : Use haconf -makerw
```

The `hastop -local` command successfully runs and you can ignore the error message.

**Workaround:** There is no workaround for this issue.

### Issues with the disk state on the CVM slave node when vxconfigd is restarted on all nodes (2615680)

When a CVM master node and a slave node have lost storage access, and `vxconfigd` is restarted on all nodes, the disk state on the CVM slave node shows as invalid.

**Work-around:**

#### To work around this issue

1. Restore storage connectivity.
2. Deport the disk group.
3. Import the disk group.

### Re-enabling connectivity if the disks are in local failed (lfailed) state (2425977)

In a Cluster Volume Manager (CVM) cluster, you can disable connectivity to the disks at the controller or enclosure level with the `vxdmpadm disable` command. In this case, CVM may place the disks into the `lfailed` state. When you restore connectivity with the `vxdmpadm enable` command, CVM may not automatically clear the `lfailed` state. After enabling the controller or enclosure, you must run disk discovery to clear the locally failed state.

#### To run disk discovery

◆ Run the following command:

```
# vxdisk scandisks
```

### The vxcdsconvert utility is supported only on the master node (2616422)

The `vxcdsconvert` utility should be run only from the master node, not from the slave nodes of the cluster.

### Issues if the storage connectivity to data disks is lost on a CVM slave node while vxconfigd was not running on the node (2562889)

If storage connectivity to data disks is lost on a CVM slave node while vxconfigd was not running on the node, this may result in following issues when `vxconfigd`comes up on this node:

- The shared disk groups on the disconnected storage are marked as dgdisabled on the slave node only.
- The shared disk groups are available to rest of the cluster nodes but no transactions, such as VxVM configuration changes, are possible on any shared disk group.
- Attempts to deport such shared disk groups will fail.

**Work-arounds:**

Use one of the following work-arounds:

- Remove the faulty slave node out of CVM cluster, restore storage connectivity, and rejoin the node to the cluster.

- Restart `vxconfigd` on the CVM master node.

### The "vxdg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set (235456)

In Cluster Volume Manager environment, "vxdg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set. This can happen on master/slave nodes.

**Workaround:**

Administrator has to run "vxdisk scandisks" or "vxdisk -o alldgs list" followed by "vxdg listclone" to get all the disks containing "clone_disk" or "udid_mismatch" flag on respective host.

### Failback to primary paths does not occur if the node that initiated the failover leaves the cluster (1856723)

When CVM is configured on non-A/A storage, if a node loses access to the storage through all the primary paths, then all the nodes in the cluster switches to the secondary paths. If the node which raised the protocol leaves the cluster and if all the rest of the nodes in the cluster are seeing the primary paths as healthy, then failback to primary paths never happens.

### Required attributes of LUNs for DMP devices with cluster set-up having fencing enabled (2521801)

When cluster set-up has fencing enabled, the following attributes are required to be set on the LUNs.

**Set the following attributes for LUNs**

1 Set the following attributes:

   - If the path has the `reserve_policy` attribute set, change the `reserve_policy` attribute to `no_reserve` for all the paths.

     ```
     # lsattr -E1 hdisk557 | grep res
     reserve_policy single_path
     Reserve Policy True

     # chdev -l hdisk557 -a reserve_policy=no_reserve -P
     hdisk557 changed
     ```

   - If the path has the `reserve_lock` attribute set, change the `reserve_lock` attribute to `no`.

```
# lsattr -El hdisk558 | grep reserve_lock
reserve_lock  yes
Reserve Device on open True
```

```
# chdev -l hdisk558 -a reserve_lock=no -P
hdisk558 changed
```

2 Reboot the system for the changes to take effect.

### A snapshot volume created on the Secondary, containing a VxFS file system may not mount in read-write mode and performing a read-write mount of the VxFS file systems on the new Primary after a global clustering site failover may fail (1558257)

**Issue 1:**

When the `vradmin ibc` command is used to take a snapshot of a replicated data volume containing a VxFS file system on the Secondary, mounting the snapshot volume in read-write mode may fail with the following error:

```
UX:vxfs mount: ERROR: V-3-21268: /dev/vx/dsk/dg/snapshot_volume
is corrupted. needs checking
```

This happens because the file system may not be quiesced before running the `vradmin ibc` command and therefore, the snapshot volume containing the file system may not be fully consistent.

**Issue 2:**

After a global clustering site failover, mounting a replicated data volume containing a VxFS file system on the new Primary site in read-write mode may fail with the following error:

```
UX:vxfs mount: ERROR: V-3-21268: /dev/vx/dsk/dg/data_volume
is corrupted. needs checking
```

This usually happens because the file system was not quiesced on the original Primary site prior to the global clustering site failover and therefore, the file systems on the new Primary site may not be fully consistent.

**Workaround:** The following workarounds resolve these issues.

For issue 1, run the `fsck` command on the snapshot volume on the Secondary, to restore the consistency of the file system residing on the snapshot.

For example:

```
# fsck -V vxfs /dev/vx/dsk/dg/snapshot_volume
```

For issue 2, run the `fsck` command on the replicated data volumes on the new Primary site, to restore the consistency of the file system residing on the data volume.

For example:

```
# fsck -V vxfs /dev/vx/dsk/dg/data_volume
```

### Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 [1985055]

Health check monitoring is not supported for Oracle database 11g R1 and 11g R2.

Workaround: Set MonitorOption attribute for Oracle resource to 0.

# Veritas Storage Foundation for Oracle RAC known issues

This section describes the known issues in this release of Veritas Storage Foundation for Oracle RAC.

## Oracle RAC issues

This section lists the known issues in Oracle RAC.

### Oracle Grid Infrastructure installation may fail with internal driver error

The Oracle Grid Infrastructure installation may fail with the following error:

```
[INS-20702] Unexpected Internal driver error
```

**Workaround**:

Perform one of the following steps depending on the type of installer you use for the installation:

- Script-based installer
  Export the `OUI_ARGS` environment variable, before you run the SFHA installation program:

  ```
  export OUI_ARGS=-ignoreInternalDriverError
  ```

  For more information, see the Oracle Metalink document: 970166.1

- Web-based installer
  When you run the Web-based installer, in the **Enter the arguments to be passed to the Oracle installer** text box, enter the value `-ignoreInternalDriverError`.

  For more information, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

### During installation or system startup, Oracle Grid Infrastructure may fail to start

After successful installation of Oracle RAC 11g Release 2 Grid Infrastructure, while executing the `root.sh` script, `ohasd` may fail to start. Similarly, during system startup, Oracle Grid Infrastructure may fail to start though the VCS engine logs may indicate that the cssd resource started Oracle Grid Infrastructure successfully.

For possible causes and workarounds, see the Oracle Metalink document: 1069182.1

### Oracle VIP Configuration Assistant fails with an error message (1182220)

During Oracle RAC 10g Release 2 installation, the VIP Configuration Assistant may fail with the following error message:

```
The given interface(s), "" is not public.
Public interfaces should be used to configure virtual IPs.
```

This message appears only when the VIP is not from the regular public IP range (for example, 200.).

Workaround: Invoke the `vipca` utility manually as the superuser.

```
# export DISPLAY=nebula:0.0
# $CRS_HOME/bin/vipca
```

### Oracle Cluster Verification utility displays a warning message

During the final stage of Oracle RAC 10g Release 2 installation, you may receive a warning message with the Oracle Cluster Verification utility.

For example:

```
Utility
============================================================
OUI-25031: Some of the configuration assistants failed. It is
strongly recommended that you retry the configuration
assistants at this time. Not successfully running any "
Recommended" assistants means your system will not be correctly
configured.
1. Check the Details panel on the Configuration Assistant Screen
to see the errors resulting in the failures.
2. Fix the errors causing these failures.
3. Select the failed assistants and click the 'Retry' button
to retry them.
============================================================
```

Workaround: You may safely ignore this message if the cluster is operating satisfactorily.

## SFHA issues

This section lists the known issues in SFHA for this release.

### SFHA installer does not support use of fully qualified domain names (2585899)

The SFHA installer does not support the use of fully qualified domain names (FQDN). Specifying the fully qualified domain name of a system results in the following error:

```
The node galaxy doesn’t seem to be part of the cluster,
or CVM is not running on the node galaxy.
```

**Workaround:** Use only the host name of the system when you specify the system name.

### PrivNIC and MultiPrivNIC agents not supported with Oracle RAC 11.2.0.2

The PrivNIC and MultiPrivNIC agents are not supported with Oracle RAC 11.2.0.2.

For more information, see the following Technote:

http://www.symantec.com/business/support/index?page=content&id=TECH145261

### Node fails to join the SFHA cluster if the file system containing Oracle Clusterware is not mounted (2611055)

The sequence number of the startup script for Oracle High Availability Services daemon (ohasd) is lower than some of the SFHA components such as VXFEN and VCS. During system startup, if the file system containing Oracle Clusterware does not get mounted before the ohasd startup script is executed, the script continuously waits for the file system to become available. As a result, the other scripts (including those of SFHA components) are not executed and the node being started does not join the SFHA cluster.

**Workaround:** If the rebooted node does not join the SFHA cluster, the cluster can be started manually using the following command:

```
# installsfrac -start node1 node2
```

### Issue with format of the last 8-bit number in private IP addresses (1164506)

The PrivNIC/MultiPrivNIC resources fault if the private IP addresses have a leading 0 in any of the octets that comprise the IP address, for example X.X.X.01 or X.X.0X.1. or X.0X.X.1 or 0X.X.X.1, where X is an octet of the IP address.

When you configure private IP addresses for Oracle Clusterware, ensure that the IP addresses have a format as displayed in the following two-node example:

- On galaxy: 192.168.12.1
- On nebula: 192.168.12.2

Confirm the correct format by viewing the PrivNIC or MultiPrivNIC resource in the `/etc/VRTSvcs/conf/config/main.cf` file.

### Policy-managed Oracle RAC databases fail to come online on some of the nodes in the server pool (2392741)

If the cardinality of a policy-managed Oracle RAC database is set to a number lesser than the number of nodes in the server pool, and if the Oracle agent tries to bring the database online on all the nodes in the server pool, the operation fails on some of the nodes in the server pool. The resource on respective nodes move to the faulted state.

### Removal of SAN cable from any node in a global cluster setup takes application service groups offline on all nodes (2580393)

In a replicated global cluster setup, the removal of SAN cable from any node in the cluster causes the CFS mount points to fault. As a result, dependent application groups are taken offline and replication to the secondary site is adversely affected.

## Software limitations

There are no software limitations in this release.

## List of patches

This section lists the patches for 6.0 RP1.

**Table 1-10** Patches for AIX

| BFF file | Size in bytes | Patches | Version |
|---|---|---|---|
| VRTSamf.bff | 9574400 | VRTSamf | 6.0.001.000 |
| VRTScavf.bff | 921600 | VRTScavf | 6.0.001.000 |
| VRTSfsadv.bff | 32716800 | VRTSfsadv | 6.0.001.000 |
| VRTSob.bff | 63027200 | VRTSob | 03.04.0526.0004 |
| VRTSsfcpi60.bff | 4403200 | VRTSsfcpi60 | 6.0.001.000 |

**Table 1-10** Patches for AIX *(continued)*

| BFF file | Size in bytes | Patches | Version |
|---|---|---|---|
| VRTSvbs.bff | 57702400 | VRTSvbs | 6.0.001.000 |
| VRTSvcs.bff | 353075200 | VRTSvcs | 6.0.001.000 |
| VRTSvxfs.bff | 43059200 | VRTSvxfs | 6.0.001.000 |
| VRTSvxvm.bff | 285900800 | VRTSvxvm | 6.0.001.000 |

**Note:** You can also view the list using the `installrp` command: `./installrp -listpatches`

# Documentation errata

The following sections cover additions or corrections for the product documentation. These additions or corrections may be included in later versions of the product documentation that can be downloaded from the Symantec Support website and the Symantec Operations Readiness Tools (SORT).

## Veritas Storage Foundation Cluster File System High Availability manual page

The following errata applies to the Veritas Storage Foundation Cluster File System High Availability 6.0 manual page:

### cfsshare manual page

The following argument is missing from the cfsshare manual page:

*network_hosts* Hosts on the network that receive pings to determine the state of the NIC. The specified hosts must be pingable.

The following example in the cfsshare manual page is missing the note:

To add a Virtual IP "10.192.111.161" with the netmask "255.255.240.0" on network interface "en0":

# cfsshare addvip en0 10.182.111.161 255.255.240.0

Note: If a NIC is a virtual NIC the network_hosts argument is a mandatory. For example: cfsshare addvip en0 10.209.145.40 255.255.252.0 10.209.116.1

# Veritas Storage Foundation Cluster File System High Availability Administrator's Guide

The following errata applies to the Veritas Storage Foundation Cluster File System High Availability 6.0 Administrator's Guide.

## "Adding a virtual IP address to VCS" section in the "Clustered NFS" chapter

The example is incorrect and should be:

cfsshare addvip [ -a *nodname* ] *device address netmask* [ *network_hosts* ]

Chapter 2

# Installing the products for the first time

This chapter includes the following topics:

- Installing the Veritas software using the script-based installer
- Installing Veritas software using the Web-based installer

## Installing the Veritas software using the script-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and install 6.0 RP1. Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

See the 6.0 *Installation Guide* and *Release Notes* for your product for more information.

**To install the Veritas software for the first time**

1. Download Storage Foundation and High Availability Solutions 6.0 from http://fileConnect.symantec.com.
2. Extract the tar ball into a directory called `/tmp/sfha6.0`.
3. Check http://sort.symantec.com/patches to see if there are any patches available for the 6.0 Installer. Download applicable P-patches and extract them to the `/tmp` directory.
4. Change to the `/tmp/sfha6.0` directory:

   ```
   # cd /tmp/sfha6.0
   ```

5 Run the installer to install SFHA 6.0.

See the *Installation Guide* for instructions on installing the 6.0 version of this product.

```
#./installer -require complete_path_to_6.0_installer_patch
```

---

**Note:** If the P-patch is not available for 6.0 installer, use the `installer` script without `-require` option.

---

6 Download SFHA 6.0 RP1 from http://sort.symantec.com/patches.

7 Extract it to a directory called `/tmp/sfha6.0RP1`.

8 Check http://sort.symantec.com/patches to see if there are patches available for the 6.0 RP1 installer. Download applicable P-patches and extract them to the `/tmp` directory.

9 Change to the `/tmp/sfha6.0RP1` directory:

```
# cd /tmp/sfha6.0RP1
```

10 Invoke the `installrp` script to install 6.0 RP1:

```
# ./installrp -require complete_path_to_6.0RP1_installer_patch
```

---

**Note:** Note: If the P-patch is not available for 6.0 RP1 installer, use the `installrp` script without `-require` option.

---

11 If you did not configure the product after the 6.0 installation, the installer prompts you to configure the product during RP installation. If you do not want to configure the product now, answer **n** when prompted. To configure the product in the future, run the product installation script from the 6.0 installation media or from `/opt/VRTS/install` directory with the `-configure` option.

## Installing Veritas software using the Web-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and then to install 6.0 RP1 using the Web-based installer. For detailed instructions on how to install 6.0

using the Web-based installer, follow the procedures in the 6.0 Installation Guide and Release Notes for your products.

See "Upgrading to 6.0 RP1" on page 96.

## Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

**To start the Web-based installer**

1 Start the Veritas XPortal Server process `xprtlwid`, on the installation server:

```
# ./webinstaller start
```

The webinstaller script displays a URL.

2 Start the Web browser on the system from which you want to perform the installation.

3 Navigate to the URL displayed from step 1.

4 The browser may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

5 When prompted, enter `root` and root's password of the installation server.

## Obtaining a security exception on Mozilla Firefox

You may need to get a security exception on Mozilla Firefox.

**To obtain a security exception**

1 Click **Or you can add an exception** link.

2 Click **Add Exception** button.

3 Click **Get Certificate** button.

4 Uncheck **Permanently Store this exception checkbox (recommended)**.

5 Click **Confirm Security Exception** button.

6 Enter root in User Name field and root password of the web server in the Password field.

## Installing 6.0 RP1 with the Veritas Web-based installer

This section describes installing SFHA with the Veritas Web-based installer.

**To install SFHA**

1. The 6.0 version of the Veritas product must be installed before upgrading to 6.0 RP1.

   See “Prerequisites for upgrading to 6.0 RP1” on page 95.

2. On the **Select a task and a product** page, select **Install 6.0 RP1** from the **Task** drop-down list, and click **Next**.

3. Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Next**.

4. After the validation completes successfully, click **Next** to install 6.0 RP1 patches on the selected system.

5. After the installation completes, you must choose your licensing method.

   On the license page, select one of the following tabs:

   - Keyless licensing

     ---

     **Note:** The keyless license option enables you to install without entering a key. However, in order to ensure compliance you must manage the systems with a management server.

     For more information, go to the following website:

     http://go.symantec.com/sfhakeyless

     ---

     Complete the following information:
     Choose whether you want to install Standard or Enterprise mode.
     Choose whether you want to enable Veritas Volume Replicator.
     For Storage Foundation High Availability, choose whether you want to enable Global Cluster option.
     Choose whether you want to enable Global Cluster option.
     Click Register.

   - Enter license key
     If you have a valid license key, select this tab. Enter the license key for each system. Click **Register**.

6 For Storage Foundation, click Next to complete the configuration and start the product processes.

For Storage Foundation High Availability, the installer prompts you to configure the cluster.

Note that you are prompted to configure only if the product is not yet configured.

If you select n, you can exit the installer. You must configure the product before you can use SFHA.

After the installation completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

7 The installer prompts you to configure the cluster.

If you select n, you can exit the installer. You must configure the product before you can use SFHA.

After the installation completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

8 Select the checkbox to specify whether you want to send your installation information to Symantec.

```
Would you like to send the information about this installation
to Symantec to help improve installation in the future?
```

Click **Finish**.

## Upgrading SFHA with the Veritas Web-based installer

This section describes upgrading SFHA with the Veritas Web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

### To upgrade SFHA

1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.

2 If you are upgrading a high availability (HA) product, take all service groups offline. To list all service groups:

```
# /opt/VRTSvcs/bin/hagrp -list
```

For each service group listed, take it offline:

```
# /opt/VRTSvcs/bin/hagrp -offline service_group -sys system_name
```

3 Start the Web-based installer.

See "Starting the Veritas Web-based installer" on page 91.

4 On the **Select a task and a product** page, select **Install 6.0 RP1** from the **Task** drop-down list, and click **Next**.

5 Stop all applications accessing the file system. Unmount all mounted filesystems before installation.

6 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Next**.

The installer detects the product that is installed on the specified system.

7 The installer stops the processes. Choose to restore and reuse the previous configuration on all systems. Click **Next** to start the processes.

8 Click **Next** to complete the upgrade.

After the upgrade completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

9 Click **Finish**. The installer prompts you for another task.

# Upgrading to 6.0 RP1

This chapter includes the following topics:

## Prerequisites for upgrading to 6.0 RP1

The following list describes prerequisites for upgrading to the 6.0 RP1 release:

- For any product in the Veritas Storage Foundation stack, you must have the 6.0 installed before you can upgrade that product to the 6.0 RP1 release.
- Each system must have sufficient free space to accommodate patches.
- The full list of prerequisites can be obtained by running `./installrp -precheck`.
- Make sure to download the latest patches for the installer.
  See "Downloading required software to upgrade to 6.0 RP1 " on page 95.

## Downloading required software to upgrade to 6.0 RP1

This section describes how to download the latest patches for the installer.

**To download required software to upgrade to 6.0 RP1**

1. Download SFHA 6.0 RP1 from http://sort.symantec.com/patches.
2. Extract it to a directory, say /tmp/sfha60rp1.

# Supported upgrade paths

This section describes the supported upgrade paths for this release:

- 6.0 to 6.0 RP1

# Upgrading to 6.0 RP1

This section describes how to upgrade from 6.0 to 6.0 RP1 on a cluster or a standalone system.

- Performing a full upgrade to 6.0 RP1 on a cluster
  Use the procedures to perform a full upgrade to 6.0 RP1 on a cluster that has Veritas Cluster Server (VCS), Veritas Storage Foundation and High Availability Solutions (SFHA), Veritas Storage Foundation Cluster File System High Availability (SFCFSHA), or Veritas Storage Foundation for Oracle RAC (SFRAC) installed and configured.
- Upgrading to 6.0 RP1 on a standalone system
  Use the procedure to upgrade to 6.0 RP1 on a system that has SF installed.
- Performing a rolling upgrade using the script-based installer
  Use the procedure to upgrade your Veritas product with a rolling upgrade.

See "Installing the Veritas software using the script-based installer" on page 89.

## Performing a full upgrade to 6.0 RP1 on a cluster

Performing a full upgrade on a cluster requires stopping cluster failover functionality during the entire procedure. However, if you use Veritas Storage Foundation Cluster File System (SFCFSHA) and Cluster Volume Manager (CVM), the SFCFSHA and CVM services remain available.

Depending on your cluster's configuration, select one of the following procedures to upgrade to 6.0 RP1:

- Performing a full upgrade to 6.0 RP1 on a Veritas Cluster Server
- Performing a full upgrade to 6.0 RP1 on an SFHA cluster
- Performing a full upgrade to 6.0 RP1 on an SFCFSHA cluster

- Performing a full upgrade to 6.0 RP1 on an SF Oracle RAC cluster
  See "Downloading required software to upgrade to 6.0 RP1 " on page 95.

## Performing a full upgrade to 6.0 RP1 on a Veritas Cluster Server

The following procedure describes performing a full upgrade on a Veritas Cluster Server (VCS) cluster.

**To upgrade VCS**

1 Make sure you have downloaded the latest VCS 6.0 RP1 patches required for the upgrade.

2 Log in as superuser.

**Note:** Upgrade the Operating System and reboot the systems if required.

3 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 6.0 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check:

```
# ./installrp -precheck node1 node2 ... nodeN
```

4 Resolve any issues that the precheck finds.

5 Start the upgrade:

```
# ./installrp node1 node2 ... nodeN
```

6 After the upgrade, review the log files for any issues.

## Performing a full upgrade to 6.0 RP1 on an SFHA cluster

The following procedure describes performing a full upgrade on a SFHA and VCS cluster.

**To perform a full upgrade to 6.0 RP1 on an SFHA cluster**

1 Make sure you have downloaded the latest software required for the upgrade.

2 Log in as superuser.

3 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.

4 On each node, enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

5 Unmount all Storage Checkpoints and file systems:

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

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the vxrvg stop command to stop each RVG individually:

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

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

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

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

7 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.

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

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

Verify that no volumes remain open:

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

9 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
```

10 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 6.0 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check.

```
# ./installrp -precheck [-rsh] node1 node2 ... nodeN
```

The program proceeds in a noninteractive mode to examine the systems for licenses, filesets, disk space, system-to-system communications, etc.

11 Review the output as the program displays the results of the check and saves the results of the check in a log file.

12 Make sure all your nodes meet the prerequisites, and all the issues reported by above pre-check process have been resolved.

13 Start the upgrade.

```
# ./installrp [-rsh] node1 node2 ... nodeN
```

Review the output.

14 Restart all the volumes. Enter the following command for each disk group:

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

15 If you stopped any RVGs in step 6, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

16 Remount all VxFS file systems on all nodes in the selected group:

```
# mount /filesystem
```

17 Remount all Storage Checkpoints on all nodes in the selected group:

```
# mount /checkpoint_name
```

18 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

## Performing a full upgrade to 6.0 RP1 on an SFCFSHA cluster

The following procedure describes performing a full upgrade on an SFCFSHA cluster.

### To perform a full upgrade to 6.0 RP1 on an SFCFSHA cluster

1 Make sure you have downloaded the latest software required for the upgrade.

2 Log in as superuser.

3 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.

4 On each node, enter the following command to check if any Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

If any Storage Checkpoints are mounted, on each node in the cluster unmount all Storage Checkpoints.

```
# umount /checkpoint_name
```

5 On each node, enter the following command to check if any VxFS file systems are mounted:

```
# mount | grep vxfs
```

- If any VxFS file systems are present, on each node in the cluster unmount all of the VxFS file systems:

  ```
  # umount /filesystem
  ```

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvg stop` command to stop each RVG individually:

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

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

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

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

7 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.

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

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

Verify that no volumes remain open:

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

9 If required, apply the OS kernel patches.

See IBM’s documentation for the procedures.

10 On each node, 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
```

11 From the directory that contains the extracted and untarred 6.0 RP1 rolling patch binaries, change to the directory that contains the installrp script.

```
# ./installrp node1 node2
```

where `node1` and `node2` are nodes which are to be upgraded.

12 After all the nodes in the cluster are upgraded, the processes restart. If the `installrp` script finds issues, it may require you to reboot the nodes.

13 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node.

14 Bring the CVM service group online on each node:

```
# hagrp -online cvm -sys nodename
```

15 Restart all the volumes by entering the following command for each disk group:

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

16 If you stopped any RVGs in step 6 , restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

17 Remount all VxFS file systems on all nodes:

```
# mount -V vxfs block_device_name mount_point
```

18 Remount all Storage Checkpoints on all nodes:

```
# mount -V vxfs -o ckpt=checkpoint_name block_device_name: checkpoint_name
```

19 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

## Performing a full upgrade to 6.0 RP1 on an SF Oracle RAC cluster

The following procedure describes performing a full upgrade on an SF for Oracle RAC cluster.

**To upgrade to 6.0 RP1 on a SF Oracle RAC cluster**

1 Make sure you have downloaded the latest software required for the upgrade.

2 Log in as superuser.

3 Verify that `/opt/VRTS/bin` is in your **PATH** so that you can execute all product commands.

4 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

And set the the AutoStart attribute of Oracle Agent to 0:

```
# hagrp -modify oracle_group AutoStart 0
# haconf -dump -makero
```

5 If the Oracle DB is not managed by VCS, prevent auto startup of Oracle DB:

```
# srvctl modify database -d db_name -y manual
```

6 Stop Oracle database on the cluster:

- If the Oracle RAC instance is managed by VCS:

  ```
  # hagrp -offline oracle_group -sys galaxy
  # hagrp -offline oracle_group -sys nebula
  ```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on all nodes in the cluster and shut down the instances:

  ```
  $ srvctl stop database -d db_name
  ```

7 Stop all applications on the cluster that are not configured under VCS. Use native application commands to stop the application.

8 Unmount the VxFS and CFS file systems that are not managed by VCS.

- Ensure that no processes are running that make use of mounted shared file system or shared volumes. To verify that no processes use the VxFS or CFS mount point, enter the following commands:

  ```
  # mount | grep vxfs
  # fuser -cu /mount_point
  ```

- Unmount the VxFS or CFS file system:

  ```
  # umount /mount_point
  ```

9 Stop all VxVM and CVM volumes for each diskgroup that are not managed by VCS on the cluster:

```
# vxvol -g disk_group stopall
```

Verify that no volumes remain open:

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

10 Stop VCS.

```
# hastop -all
```

11 From the directory that contains the extracted and untarred 6.0 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

```
# ./installrp node1 node2 ...
```

12 Manually mount the VxFS and CFS file systems that are not managed by VCS.

13 Start all applications on the cluster that are not configured under VCS. Use native application commands to start the application.

14 Relink the SF Oracle RAC libraries with Oracle.

Refer to the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* for 6.0 for more information.

15 Start Oracle Group on All nodes.

```
# hagrp -online oracle_group -any
```

16 If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on all nodes in the cluster and start the instances:

```
$ srvctl start database -d db_name
```

17
- If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

  ```
  # haconf -makerw
  # hagrp -modify oracle_group AutoStart 1
  # haconf -dump -makero
  ```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

  ```
  $ srvctl modify database -d db_name -y AUTOMATIC
  ```

## Upgrading to 6.0 RP1 on a standalone system

You can use this procedure to upgrade on a standalone system that runs SF.

**To upgrade to 6.0 RP1 on a standalone system**

1. Make sure you have downloaded the latest software required for the upgrade.
2. Log in as superuser.
3. Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
4. If required, apply the OS kernel patches.

   See IBM’s documentation for the procedures.
5. Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

   ```
   # mount | grep vxfs
   ```
6. Unmount all Storage Checkpoints and file systems:

   ```
   # umount /checkpoint_name
   # umount /filesystem
   ```
7. If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

   - Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
   - Use the `vxrvg stop` command to stop each RVG individually:

     ```
     # vxrvg -g diskgroup stop rvg_name
     ```
   - On the Primary node, use the vxrlink status command to verify that all RLINKs are up-to-date:

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

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

8 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.

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

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

Verify that no volumes remain open:

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

10 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
```

11 Copy the patch archive downloaded from the patch central to temporary location and untar the archive and browse to the directory containing the installrp installer script. Enter the `installrp` script:

```
# ./installrp nodename
```

12 If necessary, reinstate any missing mount points in the `/etc/filesystems` file.

13 Restart all the volumes by entering the following command for each disk group:

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

14 If you stopped any RVGs in step 7, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

15 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem
# mount /checkpoint_name
```

16 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

# Performing a rolling upgrade using the script-based installer

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

- About rolling upgrades
- Prerequisites for a rolling upgrades
- Performing a rolling upgrade on kernel packages for VCS, SFHA and SFCFSHA: phase 1
- Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSHA: phase 2
- Performing a rolling upgrade on kernel packages for SF Oracle RAC: phase 1
- Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2

## About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime.

Rolling upgrades take two discrete phases. In the first, you upgrade the kernel packages with exception of VCS packages and agent packages. In the second, you upgrade the non-kernel packages, which are VCS packages and agents packages.

You can perform a rolling upgrade for the following products:

- Veritas Cluster Server
- Storage Foundation and High Availability
- Storage Foundation Cluster File System and High Availability

You can perform a rolling upgrade from 6.0 to 6.0 RP1.

## Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrades.
- Split up your clusters into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you are logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.
- Make sure you have downloaded the latest software required for the upgrade.

**Limitation**: During VCS and agents upgrade, you must bring down the application High Availability (HA) for several minutes. This does not affect the application running on the cluster. You can restore the application's high availability after VCS and the agent packages are upgraded.

## Performing a rolling upgrade using the installer

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

### Performing a rolling upgrade on kernel packages for VCS, SFHA and SFCFSHA: phase 1

Note that in the following instructions a sub-cluster can represent one or more nodes in a full cluster, but is represented by nodeA.

**To perform the rolling upgrade on kernel: phase 1**

1. Browse to the directory that contains the `installrp` script.
2. On the first sub-cluster, start the installer for the rolling upgrade with the `-rollingupgrade_phase1` option.

   ```
   # ./installrp -rollingupgrade_phase1 nodeA
   ```

3. The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.
4. The installer loads new kernel modules.
5. The installer starts all the relevant processes and brings all the service groups online.

6 Before you proceed to phase 2, complete step 2 to step 5 on the second subcluster.

7 After phase 1 is completed on nodeA, the following message displays:

```
It is recommended to perform rolling upgrade phase 1 on the
systems nodeB in the next step.
```

```
Would you like to perform rolling upgrade phase 1 on the systems?
[y,n,q] (y)
```

If you choose `y`, it continues to run phase 1 of the rolling upgrade by itself on the next node B.

If you choose `n` or `q`, you need to complete step 2 to step 5 on the next node B.

After phase 1 is completed on all nodes in the cluster, the following message displays:

```
It is recommended to perform rolling upgrade phase 2 on all the
cluster systems in the next step.
```

```
Would you like to perform rolling upgrade phase 2 on the cluster?
[y,n,q] (y)
```

If you choose `y`, it continues to run phase 2 of the rolling upgrade by itself.

If you choose `n` or `q`, you need to use the following steps to finish phase 2 of the rolling upgrade.

### Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSHA: phase 2

In this phase installer installs all non-kernel filesets on all the nodes in cluster and restarts VCS cluster.

#### To perform the rolling upgrade on non-kernel packages: phase 2

1 Browse to the directory that contains the `installrp` script.

2 Start the installer for the rolling upgrade with the `--rollingupgrade_phase2` option. Specify all the nodes in the cluster:

```
# ./installrp -rollingupgrade_phase2 nodeA nodeB nodeC nodeD
```

3 The installer checks system communications, fileset versions, product versions, and completes prechecks. It verifies completion of phase 1.

4 The installer upgrades non-kernel filesets and will start HA daemon (had) on all nodes. HA will be available once HA daemon is up.

5. The installer loads the new kernel modules. It then starts all relevant processes and brings all the service groups online.

6. Verify the cluster's status:

   ```
   # hastatus -sum
   ```

## Performing a rolling upgrade on kernel packages for SF Oracle RAC: phase 1

Note that in the following instructions that a subcluster can represent one or more nodes in a full cluster, but is represented by nodeA, nodeB as subcluster1 and nodeC, nodeD as subcluster2.

### To perform the rolling upgrade on kernel: phase 1

1. Log in as superuser to one of the nodes in the cluster.

2. Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`.

   For example:

   ```
   # cp /etc/VRTSvcs/conf/config/main.cf \
         /etc/VRTSvcs/conf/config/main.cf.save
   # cp /etc/VRTSvcs/conf/config/types.cf \
         /etc/VRTSvcs/conf/config/types.cf.save
   # cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
         /etc/VRTSvcs/conf/config/OracleTypes.cf.save
   # cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
         /var/tmp/PrivNIC.cf.save
   # cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
         /var/tmp/MultiPrivNIC.cf.save
   ```

3 If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw
# hagrp -modify oracle_group AutoStart 0
# haconf -dump -makero
```

If the Oracle database is not managed by VCS, change the management policy for the database to manual. Execute the command with oracle database user credentials.

```
$ srvctl modify database -d db_name -y manual
```

4 If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to SF Oracle RAC 6.0.

5 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

- If the applications are under VCS control:

  ```
  # hagrp -offline grp_name -sys node_name
  ```

- If the applications are not under VCS control, use native application commands to stop the application.

6 For Oracle RAC 10g and Oracle RAC 11g:

Stop the Oracle RAC resources on each node.

- If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

  ```
  # hagrp -offline oracle_group -sys nodeA
  # hagrp -offline oracle_group -sys nodeB
  ```

- If the database instances are not managed by VCS, then run the following on one node:
  For Oracle RAC 11.2.0.2:

  ```
  $ srvctl stop instance -d db_name \
  -n node_name
  ```

  For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \
-i instance_name
```

7 Switch over all failover service groups to the other nodes in the cluster:

```
# hagrp -switch grp_name -to node_name
```

8 Take all the VCS service groups offline:

```
# hagrp -offline grp_name -sys node_name
```

9 Unmount all the VxFS file system which is not under VCS control on each node of subcluster.

```
# mount |grep vxfs
# fuser -c /mount_point
# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -c /mount_point
```

10 On subcluster, stop all VxVM and CVM volumes (for each disk group) that are not managed by VCS:

```
# vxvol -g disk_group stopall
```

Verify that no volumes remain open under the diskgroups which are not managed by VCS:

```
# vxprint -g disk_group -ht -e v_open
```

**Note:** Installer will automatically stop all the applications, database instances, filesystems and volumes which are under VCS control on nodes, while using the `rollingupgrade_phase1` option.

11 On the sub-cluster, start the installer for the rolling upgrade with the `-rollingupgrade_phase1` option.

```
# ./installrp -rollingupgrade_phase1 nodeA nodeB
```

**12** The installer checks system communications, fileset versions, product versions, and completes prechecks. It will stop/failover the applications, database instances, filesystems which are under VCS control. It then upgrades applicable product kernel.

**13** Relink the SF Oracle RAC libraries with Oracle.

Refer to the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide for 6.0* for more information.

**14** It is strongly recommended to reboot the nodes. Reboot the nodes as prompted by the installer.

---

**Note:** The Oracle service group at this point will be offline as the AutoStart attribute is set to 0 to prevent the service group from starting automatically.

---

**15** After phase 1 is completed on the subcluster1, you need to complete step 5 to step 13 on the next subcluster.

After phase 1 is completed on all the subclusters in the cluster, the following message displays:

```
It is recommended to perform rolling upgrade phase 2 on all the
cluster systems in the next step.
```

```
Would you like to perform rolling upgrade phase 2 on the cluster?
[y,n,q]  (y)
```

- If you choose `y`, it continues to run phase 2 of the rolling upgrade by itself. You don’t need to run phase 2: See “Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2” on page 114.
  After phase 2 upgrade, complete step 15 to step 21 (except step 19) and verify the cluster's status:

  ```
  # hastatus -sum
  ```

  If you want to upgrade CP server systems that use VCS or SFHA to 6.0, make sure that you upgraded all application clusters to version 6.0. Then, upgrade VCS or SFHA on the CP server systems.
  For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.
- If you choose `n` or `q`, you need to complete step 15 to step 21 (except step 19) and run phase 2 of the rolling upgrade: See “Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2” on page 114.

**16** Manually mount the VxFS and CFS file systems that are not managed by VCS.

17. Bring the Oracle database service group online.

    - If the Oracle database is managed by VCS:

      ```
      # hagrp -online oracle_group -sys nodeA
      # hagrp -online oracle_group -sys nodeb
      ```

    - If VCS does not manage the Oracle database:

      ```
      $ srvctl start instance -d db_name
      ```

18. Start all applications that are not managed by VCS. Use native application commands to start the applications.

19. Before you proceed to phase 2, complete step 5 to step 18 on the remaining subcluster.

20. Perform one of the following steps:

    - If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

      ```
      # haconf -makerw
      # hagrp -modify oracle_group AutoStart 1
      # haconf -dump -makero
      ```

    - If the VCS does not manage the Oracle database, change the management policy for the database to automatic:

      ```
      $ srvctl modify database -d db-name -y AUTOMATIC
      ```

21. Migrate the SFDB repository database.

### Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2

In this phase installer installs all non-kernel filesets on all the nodes in cluster and restarts VCS cluster.

**To perform the rolling upgrade on non-kernel packages: phase 2**

1 Start the installer for the rolling upgrade with the `-rollingupgrade_phase2` option. Specify all the nodes in the cluster:

```
./installrp -rollingupgrade_phase2 nodeA nodeB nodeC nodeD
```

2 The installer checks system communications, fileset versions, product versions, and completes prechecks. It verifies completion of phase 1.

3 Installer will start HA daemon (had) on all nodes, HA will be available once HA daemon is up.

4 Verify the cluster's status:

```
# hastatus -sum
```

5 If you want to upgrade CP server systems that use VCS or SFHA to 6.0, make sure that you upgraded all application clusters to version 6.0. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

# Performing a NIM upgrade to 6.0 RP1 on a cluster

The following sections describe how to perform a NIM upgrade on a cluster for Veritas Cluster Server (VCS), SFHA, SFCFSHA and SF Oracle RAC cluster.

## Performing a NIM upgrade to 6.0 RP1 on a Veritas Cluster Server cluster

The following procedures describe how to perform a NIM upgrade on a Veritas Cluster Server (VCS) cluster.

### Preparing the bundle and script resources on NIM server

You need to prepare the bundle and script resources on the NIM server before using NIM to install VCS filesets. The following actions are executed on the NIM server.

**Note:** Make sure that a NIM LPP_SOURCE is present on the NIM server.

### To prepare the bundle resources on the NIM server

1. Insert and mount the installation media.
2. Make sure that a NIM LPP_SOURCE is present on the NIM server.

   Choose an LPP source:

   ```
   # lsnim |grep -i lpp_source
   lpp-7100-00-03-1115          resources       lpp_source
   ```

3. Navigate to the product directory on the disc and run the `installrp` command to prepare the bundle and script resources:

   ```
   # ./installrp -nim lpp-7100-00-03-1115
   ```

4. Select a product to install, for example, VCS.

   The installation program copies the necessary filesets and patches to the LPP resource directory.

5. Enter a name for the bundle, for example, `VCS60_bundle`.
6. Run the `lsnim -l` command to check that the `install_bundle` resource is created successfully.

   ```
   # lsnim -l VCS60_bundle
   VCS60_bundle:
      class       = resources
      type        = installp_bundle
      Rstate      = ready for use
      prev_state  = unavailable for use
      location    = /opt/VRTS/nim/VCS60_bundle.bundle
      alloc_count = 0
      server      = master
   ```

## Prerequisite steps before upgrading VCS on the NIM client using SMIT

Perform below steps on each node that has VCS installed in the cluster.

### Prerequisite steps before upgrading VCS on the NIM client using SMIT

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

3 Make a backup copy of the current `main.cf` and all `types.cf` configuration files.

For example, on one node in the cluster:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/types.cf.save
```

4 List the service groups in your cluster and their status. On any node:

```
# hagrp -state
```

5 Take the service groups offline if they are running:

```
# hagrp -offline group -any
```

Repeat this step for each VCS service group.

6 On each node, enter the following command to check if any VxFS file systems or Storage CheckPoints are mounted.

```
# mount |grep vxfs
```

7 Unmount all Storage CheckPoints and file systems:

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

8 Stop all VxVM volumes by entering below command for each disk group:

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

Verify no volumes remain open:

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

9 Stop VCS manually:

```
# hastop -all
```

10 Stop I/O fencing、GAB、LLT:

```
# /etc/rc.d/rc2.d/S97vxfen stop
# /etc/rc.d/rc2.d/S92gab stop
# /etc/rc.d/rc2.d/S70llt stop
```

### Upgrading VCS on the NIM client using SMIT on the NIM server

You can upgrade VCS to 6.0 RP1 on each NIM client using the SMIT tool on the NIM server.

**Perform these steps on each node that has 6.0 installed in the cluster.**

1 On the NIM server, start `smitty`.

```
# smitty nim
```

2 In the menu, select **Perform NIM Software Installation and Maintenance Tasks**.

3 In the menu, select **Install and Update Software**.

4 In the menu, select **Install Software Bundle**.

5 In the menu, select a TARGET for the operation. Select a client.

6 In the menu, select the LPP_SOURCE containing the install images. In this example, specify **lpp-7100-00-03-1115**.

7 In the menu, select the bundle, for example, **VCS60_bundle**.

8 For the `installp` flags, specify that the **ACCEPT new license agreements** flag has a `yes` value.

**Note:** By default, all the software updates will be committed after upgrade, which means you cannot roll back to 6.0. If you want to only apply the upgrade, change **COMMIT software updates** to **No** and change **SAVE replace files** to **Yes**.

9 Press **Enter** to start the installation. Note that it may take some time to finish.

### Verifying software versions and restart all the processes manually on each node

You can verify the software versions and restart all the processes manually.

1. List the Veritas filesets installed on your systems:

   ```
   # lslpp -L VRTS*
   ```

   **Note:** The versions of the filesets that have been upgraded to 6.0 RP1 are 6.0.1.0.

2. Start LLT、GAB、I/O fencing and VCS on each node:

   ```
   # /etc/rc.d/rc2.d/S70llt start
   # /etc/rc.d/rc2.d/S92gab start
   # /etc/rc.d/rc2.d/S97vxfen start
   # hastart
   ```

3. Start all VxVM volumes for each disk group:

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

4. On each node, mount all VxFS file systems or Storage CheckPoints.
5. Check if the VxFS file system or Storage CheckPoints are mounted:

   ```
   # mount|grep vxfs
   ```

6. Verify the cluster's status:

   ```
   # hastatus -sum
   ```

## Performing a NIM upgrade to 6.0 RP1 on an SFHA cluster

The following procedures describe how to perform a NIM upgrade on a Veritas Storage Foundation and High Availability (SFHA) cluster.

### Preparing the bundle and script resources on NIM server

You need to prepare the bundle and script resources on the NIM server before using NIM to install SFHA filesets. The following actions are executed on the NIM server.

**Note:** Make sure that a NIM LPP_SOURCE is present on the NIM server.

**To prepare the bundle resources on the NIM server**

1 Insert and mount the installation media.

2 Make sure that a NIM LPP_SOURCE is present on the NIM server.

Choose an LPP source:

```
# lsnim |grep -i lpp_source
lpp-7100-00-03-1115          resources       lpp_source
```

3 Navigate to the product directory on the disc and run the `installrp` command to prepare the bundle and script resources:

```
# ./installrp -nim lpp-7100-00-03-1115
```

4 Select a product to install, for example, SFHA.

The installation program copies the necessary filesets and patches to the LPP resource directory.

5 Enter a name for the bundle, for example SFHA60RP1.

6 Run the `lsnim -l` command to check that the `install_bundle` resource is created successfully.

```
# lsnim -l SFHA60RP1
sfha60RP1:
   class       = resources
   type        = installp_bundle
   Rstate      = ready for use
   prev_state  = unavailable for use
   location    = /opt/VRTS/nim/SFHA60RP1.bundle
   alloc_count = 0
   server      = master
```

## Prerequisite steps before upgrading SFHA on the NIM client using SMIT

Perform below steps on each node that has SFHA 6.0 installed in the cluster.

1 Log in as superuser.

2 Verify that `/opt/VRTS/bin` is in your PATH, so that you can execute all product commands.

3 Make a backup copy of the current `main.cf` and all `types.cf` configuration files.

For example,on one node in the cluster:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/main.cf.save
```

```
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/types.cf.save
```

4 Unmount all Storage Checkpoints and file systems:

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

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

```
# mount | grep vxfs
```

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvg stop` command to stop each RVG individually:

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

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

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

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

7 Stop activity to all the VxVM volumes.

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

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

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

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

9 Stop VCS and its modules manually.

```
# hastop -all
```

10 Stop I/O fencing on each node:

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

11 Stop GAB:

```
# /etc/rc.d/rc2.d/S92gab stop
```

12 Stop LLT:

```
# /etc/rc.d/rc2.d/S70llt stop
```

13 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
```

## Upgrading SFHA on the NIM client using SMIT on the NIM server

You can upgrade SFHA to 6.0 RP1 on each NIM client using the SMIT tool on the NIM server.

### Perform these steps on each node that has 6.0 installed in the cluster

1 On the NIM server, start smitty.

```
# smitty nim
```

2 In the menu, select **Perform NIM Software Installation and Maintenance Tasks**.

3 In the menu, select **Install and Update Software**.

4 In the menu, select **Install Software Bundle**.

5 In the menu, select a TARGET for the operation. Select a client.

6 In the menu, select the LPP_SOURCE containing the install images. In this example, specify lpp-7100-00-03-1115.

7 In the menu, select the bundle, for example, SFHA60RP1.

8 For the `installp` flags, specify that the `ACCEPT new license agreements` flag has a `yes` value.

**Note:** By default, all the software updates will be committed after upgrade, which means you can't roll back to 6.0. If you want to only apply these upgrade, change **COMMIT software updates** to **No** and change **SAVE replace files** to **Yes**.

9 Press **Enter** to start the installation. Note that it may take some time to finish.

### Verifying software versions and restart all the processes manually on each node

1 List the Veritas filesets that are installed on your systems:

```
# lslpp -L VRTS*
```

**Note:** The versions of the filesets that have been upgraded to 6.0 RP1 are 6.0.1.0, and VRTSob version is 3.4.526.4.

2 Start LLT、GAB、I/O fencing and VCS on each node:

```
# /etc/rc.d/rc2.d/S70llt start
# /etc/rc.d/rc2.d/S92gab start
# /etc/rc.d/rc2.d/S97vxfen start
# hastart
```

3 Start all VxVM volumes for each disk group:

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

4 If you stopped any RVGs in the prerequisite step 6, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

5 Remount all VxFS file systems on all nodes in the selected group:

```
# mount /filesystem
```

6 Remount all the Storage Checkpoints on all nodes in the selected group:

```
# mount /checkpoint_name
```

## Performing a NIM upgrade to 6.0 RP1 on an SFCFSHA cluster

The following procedures describe how to perform a NIM upgrade on a Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) cluster.

### Preparing the bundle and script resources on NIM server

You need to prepare the bundle and script resources on the NIM server before using NIM to install SFCFSHA filesets. The following actions are executed on the NIM server.

**Note:** Make sure that a NIM LPP_SOURCE is present on the NIM server.

**To prepare the bundle resources on the NIM server**

1 Insert and mount the installation media.

2 Make sure that a NIM LPP_SOURCE is present on the NIM server.

Choose an LPP source:

```
# lsnim |grep -i lpp_source
lpp-7100-00-03-1115          resources        lpp_source
```

3 Navigate to the product directory on the disc and run the `installrp` command to prepare the bundle and script resources:

```
# ./installrp -nim lpp-7100-00-03-1115
```

4 Select a product to install, for example, SFCFSHA.

The installation program copies the necessary filesets and patches to the LPP resource directory.

5. Enter a name for the bundle, for example SFCFSHA60RP1.

6. Run the `lsnim -l` command to check that the `install_bundle` resource is created successfully.

   ```
   # lsnim -l SFCFSHA60RP1
   sfcfsha60RP1:
      class       = resources
      type        = installp_bundle
      Rstate      = ready for use
      prev_state  = unavailable for use
      location    = /opt/VRTS/nim/SFCFSHA60RP1.bundle
      alloc_count = 0
      server      = master
   ```

### Prerequisite steps before upgrading SFCFSHA on the NIM client using SMIT

Perform below steps on each node that has SFCFSHA installed in the cluster.

1. Log in as superuser.

2. Verify that `/opt/VRTS/bin` is in your PATH, so that you can execute all product commands.

3. Make a backup copy of the current `main.cf` and all `types.cf` configuration files.

   For example, on one node in the cluster:

   ```
   # cp /etc/VRTSvcs/conf/config/main.cf \
   /etc/VRTSvcs/conf/main.cf.save
   ```

   ```
   # cp /etc/VRTSvcs/conf/config/types.cf \
   /etc/VRTSvcs/conf/types.cf.save
   ```

4. List the service groups in your cluster and their status. On any node:

   ```
   # hagrp -state
   ```

5. Take the service groups offline if they are running, which contain VxFS and CFS resources:

   ```
   # hagrp -offline group -any
   ```

   Repeat this step for each SFCFSHA service group.

6 On each node, enter the following command to check if any VxFS file systems or Storage CheckPoints are mounted:

```
# mount|grep vxfs
```

7 Unmount all the Storage CheckPoints and file systems:

```
# umount /checkpoint_name
```

```
# umount /filesystem
```

8 Stop all VxVM volumes by entering below command for each disk group:

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

Verify no volumes remain open:

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

9 Stop VCS and its modules manually:

```
# hastop -all
```

10 Stop odm、I/O fencing、GAB、LLT:

```
# /etc/rc.d/rc2.d/S99odm stop
```

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

```
# /etc/rc.d/rc2.d/S92gab stop
```

```
# /etc/rc.d/rc2.d/S70llt stop
```

11 Unload the ODM module:

```
# genkex|grep odm
```

```
# vxkextadm vxodm unload
```

12 Check if VEA service is running, if the VEA service is running, stop it:

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

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

### Upgrading SFCFSHA on the NIM client using SMIT on the NIM server

You can upgrade SFCFSHA to 6.0 RP1 on each NIM client using the SMIT tool on the NIM server.

**Perform these steps on each node that has 6.0 installed in the cluster:**

1 On the NIM server, start smitty.

```
# smitty nim
```

2 In the menu, select **Perform NIM Software Installation and Maintenance Tasks**.

3 In the menu, select **Install and Update Software**.

4 In the menu, select **Install Software Bundle**.

5 In the menu, select a TARGET for the operation. Select a client.

6 In the menu, select the LPP_SOURCE containing the install images. In this example, specify lpp-7100-00-03-1115.

7 In the menu, select the bundle, for example, SFCFSHA60RP1.

8 For the installp flags, specify that the **ACCEPT new license agreements** flag has a **yes** value.

**Note:** By default, all the software updates will be committed after upgrade, which means you can't roll back to 6.0. If you want to only apply these upgrade, change **COMMIT software updates** to **No** and change **SAVE replace files** to **Yes**.

9 Press **Enter** to start the installation. Note that it may take some time to finish.

### Verifying software versions and restart all the processes manually on each node

1. To list the Veritas filesets installed on your systems:

   ```
   # lslpp -L VRTS*
   ```

   **Note:** The versions of the filesets that have been upgraded to 6.0 RP1 are 6.0.1.0, and VRTSob version is 3.4.526.4.

2. Start LLT、GAB、odm、I/O fencing and VCS on each node:

   ```
   # /etc/rc.d/rc2.d/S70llt start
   ```

   ```
   # /etc/rc.d/rc2.d/S92gab start
   ```

   ```
   # /etc/rc.d/rc2.d/S99odm start
   ```

   ```
   # /etc/rc.d/rc2.d/S97vxfen start
   ```

   ```
   # hastart
   ```

3. Start all VxVM volumes for each disk group:

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

4. On each node, mount all the VxFS file systems or Storage CheckPoints.

5. Check if the VxFS file system or Storage CheckPoints are mounted:

   ```
   # mount|grep vxfs
   ```

6. Verify the cluster's status:

   ```
   # hastatus -sum
   ```

## Performing a NIM upgrade to 6.0 RP1 on an SF Oracle RAC cluster

You can use this procedure to perform a NIM upgrade to 6.0 RP1 on a Storage Foundation Oracle RAC (SFRAC) cluster.

### Preparing the bundle and script resources on NIM server

You need to prepare the bundle and script resources on the NIM server before using NIM to install SF Oracle RAC filesets. The following steps are executed on the NIM server.

**Note:** Make sure that the NIM LPP_SOURCE is present on the NIM server.

**To prepare the bundle resources on the NIM server:**

1 Insert and mount the installation media.

2 Make sure that the NIM LPP_SOURCE is present on the NIM server.

Choose an LPP source:

```
# lsnim | grep -i lpp_source
lpp-7100-00-03-1115          resources        lpp_source
```

3 Navigate to the product directory on the disc and run the `installrp` command to prepare the bundle and script resources:

```
# ./installrp -nim lpp-7100-00-03-1115
```

4 Select a product to install, for example, SF Oracle RAC.

The installation program copies the necessary filesets and patches to the LPP resource directory.

5 Enter a name for the bundle, for example SFRAC60RP1.

6 Run the `lsnim -l` command to check that the `install_bundle` resource is created successfully.

```
# lsnim -l SFRAC60RP1
```

sfrac60RP1:

```
   class       = resources
   type        = installp_bundle
   Rstate      = ready for use
   prev_state  = unavailable for use
   location    = /opt/VRTS/nim/SFRAC60RP1.bundle
   alloc_count = 0
   server      = master
```

### Prerequisite steps before upgrading SF Oracle RAC on the NIM client using SMIT

Perform below steps on each node that has SF Oracle RAC 6.0 installed in the cluster

**Back up the following configuration files on your system:**

1. Log in as superuser to one of the nodes in the cluster.
2. Back up the following configuration files on your system:

   `main.cf, types.cf, CVMTypes.cf, CFSTypes.cf, OracleTypes.cf, OracleASMTypes.cf,PrivNIC.cf, MultiPrivNIC.cf, /etc/llttab, /etc/llthosts, /etc/gabtab, /etc/vxfentab, /etc/vxfendg, /etc/vxfenmode`.

   For example:

   ```
   # cp /etc/VRTSvcs/conf/config/main.cf \
   /etc/VRTSvcs/conf/config/main.cf.save

   # cp /etc/VRTSvcs/conf/config/types.cf \
   /etc/VRTSvcs/conf/config/types.cf.save

   # cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
   /etc/VRTSvcs/conf/config/OracleTypes.cf.save

   # cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
   /var/tmp/PrivNIC.cf.save

   # cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
   /var/tmp/MultiPrivNIC.cf.save
   ```

### Stopping SF Oracle RAC manually on each node

Perform the steps in the following procedures to stop SF Oracle RAC manually on each node.

### To stop SF Oracle RAC manually on each node

1. Stop the Oracle database.

   If the Oracle RAC instance is managed by VCS, log in as the root user and take the service group offline:

   ```
   # hagrp -offline oracle_group -sys node_name
   ```

   If the Oracle database instance is not managed by VCS, log in as the Oracle user on one of the nodes and shut down the instance:

   For Oracle RAC 11.2.0.2:

   ```
   $ srvctl stop instance -d db_name -n node_name
   ```

   For Oracle RAC 11.2.0.1 and earlier versions:

   ```
   $ srvctl stop instance -d db_name -i instance_name
   ```

2. Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.

3. Unmount the VxFS file systems that are not managed by VCS.

   Make sure that no processes are running, in case they make use of mounted shared file system or shared volumes:

   ```
   # mount -v | grep vxfs
   ```

   ```
   # fuser -cu /mount_point
   ```

   Unmount the VxFS file system:

   ```
   # umount /mount_point
   ```

4. Stop all the VxVM and CVM volumes for each diskgroup that are not managed by VCS on the cluster:

   ```
   # vxvol -g disk_group stopall
   ```

   Verify that no volumes remain open:

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

5 Take the VCS service groups offline:

```
# hagrp -offline group_name -sys node_name
```

Verify that the VCS service groups are offline:

```
# hagrp -state group_name
```

6 Stop VCS, CVM and CFS:

```
# hastop -local
```

Verify that the ports 'f', 'u', 'v', 'w', 'y', and 'h' are closed:

7 Stop ODM:

```
# /etc/rc.d/rc2.d/S99odm stop
```

8 Stop VCSMM:

```
# /etc/init.d/vcsmm.rc stop
```

9 Stop fencing:

```
# /etc/init.d/vxfen.rc stop
```

10 Stop GAB:

```
# /etc/init.d/gab.rc stop
```

11 Stop LMX modules:

```
# /etc/init.d/lmx.rc stop
```

```
# /usr/lib/methods/lmxext -stop
```

12 Stop LLT:

```
# /etc/init.d/llt.rc stop
```

13 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
```

### Upgrading SF Oracle RAC on the NIM client using SMIT on the NIM server

You can upgrade SF Oracle RAC to the 6.0 RP1 on each NIM client using the SMIT tool on the NIM server.

**Perform these steps on each node that has 6.0 installed in the cluster:**

1 On the NIM server, start smitty.

```
# smitty nim
```

2 In the menu, select **Perform NIM Software Installation and Maintenance Tasks**.

3 In the menu, select **Install and Update Software**.

4 In the menu, select **Install Software Bundle**.

5 In the menu, select a TARGET for the operation. Select a client.

6 In the menu, select the LPP_SOURCE containing the install images. In this example, specify lpp-7100-00-03-1115.

7 In the menu, select the bundle, for example, **SFRAC60RP1**.

8 For the `installp` flags, specify that the **ACCEPT new license agreements** flag has a **yes** value.

**Note:** By default, all the software updates will be committed after upgrade which means you can't roll back to 6.0.If you want to only apply these upgrade, change **COMMIT software updates** to **No** and change **SAVE replace files** to **Yes**.

9 Press **Enter** to start the installation. Note that it may take some time to finish.

### Verifying software versions and restart all the processes manually on each node

To list the Veritas filesets installed on your systems:

```
# lslpp -L VRTS*
```

**Note:** The versions of the filesets that have been upgraded to 6.0 RP1 are 6.0.1.0, and the VRTSob version is 3.4.526.4.

### Starting SF Oracle RAC manually on each node

Perform the steps in the following procedures to start SF Oracle RAC manually on each node.

**To start SF Oracle RAC manually on each node:**

1. Log into each node as the root user.
2. Start LLT:

   ```
   # /etc/init.d/llt.rc start
   ```

3. Start GAB:

   ```
   # /etc/init.d/gab.rc start
   ```

4. Start fencing:

   ```
   # /etc/init.d/vxfen.rc start
   ```

5. Start VCSMM:

   ```
   # /etc/init.d/vcsmm.rc start
   ```

6. Start LMX modules:

   ```
   # /usr/lib/methods/lmxext -start
   ```

   ```
   # /etc/init.d/lmx.rc start
   ```

7. Start ODM:

   ```
   # /etc/rc.d/rc2.d/S99odm start
   ```

8. Start VCS, CVM, and CFS:

   ```
   # hastart
   ```

9 Verify that all GAB ports are up and running:

```
# gabconfig -a

GAB Port Memberships
===============================================================
Port a gen 564004 membership 01
Port b gen 564008 membership 01
Port d gen 564009 membership 01
Port f gen 564024 membership 01
Port h gen 56401a membership 01
Port o gen 564007 membership 01
Port u gen 564021 membership 01
Port v gen 56401d membership 01
Port w gen 56401f membership 01
Port y gen 56401c membership 01
```

10 Manually mount the VxFS and CFS file systems that are not managed by VCS.

11 Relink the SF Oracle RAC libraries with Oracle.

Refer to Veritas Storage Foundation for OracleRAC 6.0 or later Installation and Configuration Guide for more information.

12 Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

```
# hagrp -online oracle_group -any
```

- If the Oracle database is not managed by VCS:

```
$ srvctl start database -d db_name
```

13 Start all applications that are not managed by VCS. Use native application commands to start the applications.

# Verifying software versions

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

```
# lslpp -L VRTS*
```

The output version for 6.0 RP1 is `6.0.1.0`, and the VRTSob version is `3.4.526.4`.

Chapter 4

# Rolling back and removing Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

## About rolling back Veritas Storage Foundation and High Availability Solutions 6.0 RP1

This section describes how to roll back either by using the `uninstallrp` script or manually.

## Rolling back using the uninstallrp script

Use the following procedure to roll back from any Veritas product to the previous version using the `uninstallrp` script.

**To roll back**

1 Browse to the directory that contains the uninstallrp script.

2 Unmount all the Storage Checkpoints and the file systems.

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

Verify that you unmounted the the Storage Checkpoints and the file systems.

```
# mount | grep vxfs
```

3 Stop all the processes and services accessing the file systems. For each disk group enter:

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

Verify that no volumes remain open.

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

4 Run the uninstallrp script to rollback patches, type:

```
# ./uninstallrp
```

5 The uninstallrp script checks whether the patches are at 6.0 (or later) or 6.0 PR1 commited level, and 6.0 RP1 applied level. If this is not the case, error messages showing the list of packages and commit levels will be shown.

6 The uninstallrp script removes 6.0 RP1 patches. After patch rollback completes, modules are loaded and processes are restarted. `uninstallrp` will also report any warning happened during uninstallation.

- For other products:

1 Run the `uninstallrp` command, type:

```
# ./uninstallrp system_list
```

2 If you performed a roll back on a system that has an encapsualted boot disk, you must reboot the system. After reboot, you may need to run `hagrp -list Frozen=1` to get the frozen SG list . Then run `hagrp -unfreeze <group> -persistent` to unfreeze all the frozen SGs manually.

# Rolling back manually

Use one of the following procedures to roll back to 6.0 manually.

- Rolling back Storage Foundation or Storage Foundation and High Availability manually
- Rolling back Storage Foundation Cluster File System High Availability manually
- Rolling back Veritas Cluster Server manually
- Rolling back Dynamic Multi-Pathing manually

**Note:** You must reboot systems that you roll back manually at the end of the roll back procedure.

## Rolling back Storage Foundation or Storage Foundation and High Availability manually

Use the following procedure to roll back to 6.0 manually.

**To roll back SF or SFHA**

1. Log in as superuser.
2. Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
3. Unmount all Storage Checkpoints and file systems:

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

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

   ```
   # mount | grep vxfs
   ```

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

   - Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
   - Use the `vxrvg stop` command to stop each RVG individually:

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

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

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

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

6 Stop activity to all VxVM 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, enter the following command:

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

8 Stop VCS and its modules manually.

```
# hastop -all -force
```

9 Stop I/O fencing on each node:

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

10 Stop GAB:

```
# /etc/rc.d/rc2.d/S92gab stop
```

11 Stop LLT:

```
# /etc/rc.d/rc2.d/S70llt stop
```

12 Unmount `/dev/odm`:

```
# umount /dev/odm
```

13 Unload the ODM module:

```
# genkex | grep odm
# vxkextadm vxodm unload
```

14 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
```

15
- Create a file that contains all the 6.0 RP1 patches. In this example, it is called `/reject.list`.
- Reject each patch from the patch list file, for example:

```
# installp -rBXf /reject.list
```

16 Reboot the systems. On each system, run the following command.

```
# /usr/sbin/shutdown -r
```

## Rolling back Storage Foundation Cluster File System High Availability manually

Use the following procedure to roll back to 6.0 manually.

**To roll back SFCFSHA manually**

1 Log in as superuser.

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

3 Unmount all Storage Checkpoints and file systems:

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

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

```
# mount | grep vxfs
```

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvg stop` command to stop each RVG individually:

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

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

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

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

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, enter the following command:

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

8 Stop VCS along with all the resources. Then, stop the remaining resources manually:

```
# /etc/rc.d/rc2.d/S99vcs stop
```

9 Unmount `/dev/odm`:

```
# umount /dev/odm
```

10 Unload the ODM module:

```
# genkex | grep odm
# vxkextadm vxodm unload
```

11 Stop I/O fencing on each node:

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

12 Stop GAB:

```
# /etc/rc.d/rc2.d/S92gab stop
```

13 Stop LLT:

```
# /etc/rc.d/rc2.d/S70llt stop
```

14 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
```

15 Remove the Storage Foundation Cluster File System High Availability 6.0 RP1 patches.

- Create a file that contains all the 6.0 RP1 patches. In this example, it is called `/reject.list`.
- Reject each patch from the patch list file, for example:

  ```
  # installp -rBXf /reject.list
  ```

16 Reboot the systems. On each system, run the following command.

```
# /usr/sbin/shutdown -r
```

## Rolling back Veritas Cluster Server manually

Use the following procedure to roll back VCS 6.0 RP1 to VCS 6.0 on your cluster manually. To uninstall VCS, see the *Veritas Cluster Server Installation Guide*.

**Note:** Use this procedure only when rolling back VCS. Do not roll back VCS when it is part of other products that rely on VCS, for example Storage Foundation Clustered File System or Storage Foundation for Oracle RAC.

**To roll back 6.0 RP1:**

1 Verify that all of the VCS 6.0 RP1 patches are in the APPLIED state. Create a text file called filesets.to.reject that contains the name and version of each fileset, one per line, exactly as shown below.

```
VRTSamf                    6.0.1.0
VRTSvcs                    6.0.1.0
VRTSvbs                    6.0.1.0
```

2 On each node, make a local copy of filesets.to.reject and then type:

```
# nohdr='^Z$'
  # while read pkg ver; do
  lslpp -l  $pkg | egrep -v "$nohdr"
  nohdr='^  Fileset +Level  State '
  done  < filesets.to.reject
```

**Note:** Any updates that are in COMMITTED state cannot be rejected (undone). You must remove each one and then re-install it.

3 List the service groups in your cluster and their status. On any node, type:

```
# hagrp -state
```

4 If there is some service group is created, for example ClusterService service group, take the ClusterService service group offline if it is running. On any node, type:

```
# hagrp -offline -force ClusterService -any
```

5 Make the VCS configuration writable. On any node, type:

```
# haconf -makerw
```

6 Freeze all service groups except the ClusterService service group. On any node, type:

```
# hagrp -list | sort -u +0b -1 | \
   while read grp sys ; do
       hagrp -freeze $grp -persistent
   done
```

You can safely ignore the warning about the failure to freeze the ClusterService group.

7 Save the configuration (main.cf) file with the groups frozen. On any node, type:

```
# haconf -dump -makero
```

8 Make a backup copy of the current main.cf and all types.cf configuration files. For example, on one node in the cluster, type:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
   /etc/VRTSvcs/conf/main.cf.save
  # cp /etc/VRTSvcs/conf/config/types.cf \
   /etc/VRTSvcs/conf/types.cf.save
```

9 Shut down VCS. On any node, type:

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

10 Shut down CmdServer. On each node, type:

```
# /opt/VRTSvcs/bin/CmdServer -stop
```

11 Verify that VCS has shut down.

- On any node, type:

  ```
  # /sbin/gabconfig -a
  ```

  The output resembles:

  ```
  GAB Port Memberships
     Port a gen 23dc0001 membership 01
  ```

  Output for membership for port h does not appear.

- On each node, run the command:

  ```
  # ps -ef | egrep "had|hashadow|CmdServer"
  ```

Terminate any instances of had, hashadow, or CmdServer that still run after 60 seconds.

12 Stop AMF, fencing, GAB, and LLT.

```
# /etc/init.d/amf.rc stop
# /etc/init.d/vxfen.rc stop
# /etc/methods/vxfenext -stop
# /etc/init.d/gab.rc stop
# /etc/methods/gabkext -stop
# /etc/init.d/llt.rc stop
```

13 Preview the patch removal selection and validity tests. On each node, type:

```
# installp -pr -gXv -f filesets.to.reject
```

Confirm that the patches to be removed are exactly the same as those listed in the filesets.to.reject file that you created in step 1.

14 Perform the patch removal. On each node, type:

```
# installp -r -gXv -f filesets.to.reject
```

Review the summaries at the end of each run and confirm that all of the intended patches removed successfully.

15 Reboot all nodes in the cluster.

16 After VCS has started, perform the following steps:

- Verify all resources have been probed. On any node, type:

```
# hastatus -summary
```

- Unfreeze all service groups. On any node, type:

```
# haconf -makerw
   # hagrp -list | sort -u +0b -1 | \
    while read grp sys ; do
     hagrp -unfreeze $grp -persistent
        done
   # haconf -dump -makero
```

You can safely ignore the warning about the failure to unfreeze the ClusterService group.

17 Bring the ClusterService service group online, if necessary. On any node, type:

```
# hagrp -online ClusterService -sys system
```

where system is the node name.

## Rolling back Dynamic Multi-Pathing manually

Use the following procedure to roll back to 6.0 manually.

**To roll back DMP manually**

1 Stop activity to all VxVM volumes.

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

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

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

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

3 Perform the following commands to determine whether root support or DMP native support is enabled.

- `# vxdmpadm gettune dmp_native_support`

  If the command returns an "on" value, DMP native support is enabled on the system. If the command returns any other value, DMP native support is disabled.

- `# vxdmpadm native list vgname=rootvg`

  If the output is a list of hdisks, root support is enabled on this system. If the command returns any other value, root support is disabled.

- Once you have determined if root support or DMP native support is enabled, go to step 4.

- Once you have determined that root support and DMP native support is not enabled, go to step 5.

4 If root support or DMP native support is enabled:

- You must disable DMP native support.

Run the following command to disable DMP native support and to disable root support:

```
# vxdmpadm settune dmp_native_support=off
```

- If only root support is enabled, run the following command to disable root support:

  ```
  # vxdmpadm native disable vgname=rootvg
  ```

- Reboot the system:

  ```
  # shutdown -r now
  ```

- Before backing out patch, stop the VEA server's vxsvc process:

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

- Create a file that contains all the 6.0 RP1 patches. In this example, it is called `/reject.list`. The content of reject.list for rejecting Dynamic multipathing is `VRTSvxvm`.
- Reject each patch from the patch list file, for example:

  ```
  # installp -rBFX /reject.list
  ```

- Reboot the system:

  ```
  # shutdown -r now
  ```

- Enable DMP native support, this also enables root support:

  ```
  # vxdmpadm settune dmp_native_support=on
  ```

- Reboot the system:

  ```
  # shutdown -r now
  ```

- Verify DMP native or root support is enabled:

  ```
  # vxdmpadm gettune dmp_native_support
  ```

5 If root support or DMP native support is not enabled:

- Before you back out the patch, kill the VEA Server's vxsvc process:

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

- To reject the patch if it is in APPLIED state

  ```
  # installp -r patch_name
  ```

- Reboot the system:

  ```
  # shutdown -r now
  ```

# Removing the Veritas product

Use one the following procedures to remove the Veritas product.

## Removing 6.0 RP1 on Veritas Storage Foundation or Veritas Storage Foundation Cluster File System High Availability

You can use the following procedure to uninstall 6.0 RP1 on SF or SFCFSHA.

**To uninstall 6.0 RP1 on Veritas Storage Foundation or Veritas Storage Foundation Cluster File System High Availability**

1. Log in as superuser.
2. Verify that /opt/VRTS/bin is in your PATH so you can execute all product commands.
3. Unmount all Storage Checkpoints and file systems:

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

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

   ```
   # mount | grep vxfs
   ```

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

   - Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
   - Use the vxrvg stop command to stop each RVG individually:

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

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

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

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

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, enter the following command:

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

8 Stop VCS and its modules manually.

```
# hastop -all -force
```

9 Stop VXFEN:

```
# /etc/init.d/vxfen.rc stop
# /etc/methods/vxfenext -stop
```

10 Unload the ODM module:

```
# genkex | grep odm
# vxkextadm vxodm unload
```

11 Stop GAB:

```
# /etc/init.d/gab.rc stop
# /etc/methods/gabkext -stop
```

12 Check if the LLT is running:

```
# lltconfig
# lltstat -nvv
```

If the LLT is running, stop it:

```
# /etc/init.d/llt.rc stop
```

13 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
```

14 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 or Veritas Storage Foundation Cluster File System High Availability, enter the following commands:

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

You can use this command to remove the packages from one or more systems. For other products, substitute the appropriate script for `uninstallsf` such as `uninstallsfcfsha` for the Storage Foundation Cluster File System High Availability software. 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 nodes of the sub-cluster.

15 After uninstalling the Veritas software, refer to the appropriate product's 6.0 Installation Guide document to reinstall the 6.0 software.

## Removing 6.0 RP1 on Veritas Storage Foundation for Oracle RAC

You can use the following procedure to uninstall the 6.0 RP1 on Storage Foundation for Oracle RAC systems.

**Note:** This procedure will remove the complete SF for Oracle RAC stack from all nodes.

**To uninstall the 6.0 RP1 on Veritas Storage Foundation for Oracle RAC**

1. Stop Oracle and CRS on each node of the cluster.

   - If CRS is controlled by VCS, log in as superuser on each system in the cluster and enter the following command:

     ```
     # hastop -all
     ```

   - If CRS is not controlled by VCS, enter the following command on each node of the cluster to stop CRS:

     ```
     # /opt/crs/oracleversion/bin/crsctl stop crs
     ```

     Unmount all VxFS file system used by a database or application and enter the following command to each node of the cluster:

     ```
     # hastop -local
     ```

2. Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped. In the `gabconfig -a` command output, the VCS engine or high availability daemon (HAD) port h is not displayed. This indicates that VCS has been stopped.

   ```
   # /sbin/gabconfig -a
   ```

   Sample output:

   ```
   GAB Port Memberships
     ===============================
     Port a gen 5c3d0b membership 01
     Port b gen 5c3d10 membership 01
     Port d gen 5c3d0c membership 01
     Port o gen 5c3d0f membership 01
   ```

3 Uninstall Storage Foundation for Oracle RAC.

```
# cd /opt/VRTS/install
# ./uninstallsfrac MyNode1 MyNode2
```

See the *Veritas Storage Foundation for Oracle RAC 6.0 Installation and Configuration Guide* for more information.

4 After uninstalling the packages, refer to the *Veritas Storage Foundation for Oracle RAC 6.0 Installation and Configuration Guide* to reinstall the 6.0 software.