Veritas Storage Foundation™ and High Availability Release Notes

Linux

6.0 Rolling Patch 1

August 2012

Symantec™

Veritas Storage Foundation and High Availability Release Notes

Product version: 6.0 Rolling Patch 1

Document version: 6.0RP1.4

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's 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 script

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

See "Supported upgrade paths" on page 108.

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 RPMs 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 RPMs.
–kickstart *dir_path*	Produces a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of Symantec rpms in the correct order for installing, in a format that can be used for Kickstart installations. The *dir_path* indicates the path to the directory in which to create the file.
-yumgroupxml	The `-yumgroupxml` option is used to generate a yum group definition XML file. The `createrepo` command can use the file on Redhat Linux to create a yum group for automated installation of all rpms for a product. An available location to store the XML file should be specified as a complete path. The `-yumgroupxml` option is supported on Redhat Linux only.
–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.

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

Command Line Option	Function
–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 RPM group (minimum, recommended, all) and RPMs that are installed on the specified systems.
–pkgtable	Displays product's RPMs in correct installation order by group.
–pkginfo	Displays a list of RPMs 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 RPMs.
–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 RPMs and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing RPMs and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available.

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

Command Line Option	Function
–nolic	Allows installation of product RPMs 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 RPMs 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 RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version."

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 Linux operating systems

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

Table 1-2 shows the supported operating systems for this release.

Table 1-2 Supported operating systems

Operating systems	Levels	Kernel version	Chipsets
Red Hat Enterprise Linux 6	6.1 6.2	2.6.32-131.0.15.el6 2.6.32-220.el6	64-bit x86, EMT*/Opteron 4.1 64-bit only
Red Hat Enterprise Linux 5	Update 5 Update 6 Update 7	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5	64-bit x86, EMT*/Opteron 4.1 64-bit only

Table 1-2 Supported operating systems *(continued)*

Operating systems	Levels	Kernel version	Chipsets
SUSE Linux Enterprise 11	SP1	2.6.32.12-0.7	64-bit x86, EMT*/Opteron 4.1 64-bit only
SUSE Linux Enterprise 10	SP4	2.6.16.60-0.85.1	64-bit x86, EMT*/Opteron 4.1 64-bit only
Oracle Enterprise Linux 6	**6.1 **6.2	2.6.32-131.0.15.el6 2.6.32-220.el6	64-bit x86, EMT*/Opteron
Oracle Enterprise Linux 5	**Update 5 **Update 6 **Update 7	2.6.18-194.el5 2.6.18-238.el5 2.6.18-274.el5	64-bit x86, EMT*/Opteron

* Extended Memory Technology

** RHEL-compatible mode only.

Note: Only 64-bit operating systems are supported.

If your system is running an older version of either Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Oracle Enterprise Linux, upgrade it before attempting to install the Veritas software. Consult the Red Hat, SUSE, or Oracle documentation for more information on upgrading or reinstalling your operating system.

Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel Application Binary Interface (ABI) compatibility.

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

Mandatory patch required for Oracle Bug 4130116

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

For more information, refer to the following TechNote:

http://www.symantec.com/docs/HOWTO19718

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*.

VMware Environment

For information about the use of this product in a VMware Environment, refer to http://entsupport.symantec.com/docs/289033

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 1-3 SFDB features database support for 6.0 RP1

SFDB feature	DB2	Oracle	Sybase
Oracle Disk Manager, Cached Oracle Disk Manager	No	Yes	No
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 1-4 Veritas Volume Manager 6.0 RP1 fixed issues

Fixed issues	Description
2674465	Data Corruption while adding/removing LUNs.
2666163	A small portion of possible memory leak introduced due to addition of enhanced messages.
2660151	vxconfigd is generating a series of LVM header messages for devices (CLONES/replicated devices)Secondary EMC MirrorView LUNS in an error state.
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.
2629429	vxunroot does not set original menu.lst and fstab files, SUSE 10.0 NETAPP FAS3000 ALUA SANBOOT.

Table 1-4 Veritas Volume Manager 6.0 RP1 fixed issues *(continued)*

Fixed issues	Description
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.
2605702	Bail out initialising disk with large sector size and its foreign device.
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.
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.

Table 1-4 Veritas Volume Manager 6.0 RP1 fixed issues *(continued)*

Fixed issues	Description
2556467	disabling all paths and rebooting host causes /etc/vx/.vxdmprawdev record loss.
2535142	New crash was detected on RHEL6.1 during upgrade due to mod unload, possibly of vxspec.
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.
2402774	Install Upgrade : After upgrade to 6.0 encapsulated root disk is marked as 'clone_disk'.
2348180	Failure during validating mirror name interface for linked mirror volume.
2176084	Intermittent failure to start ESD on a node.
1765916	Correcting World Writable and unapproved file permissions.

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 1-5 Veritas File System 6.0 RP1 fixed issues

Fixed issues	Description
2679361	Network Customization screen doesn't show any NICs in I18N-level0 environment.
2678096	The fiostat command dumps core when the count value is 0.
2672201	Certain commands get blocked by kernel, causing EACCES(ERRNO = 13).
2672148	vxdelestat (1M) when invoked with -v option goes into infinite loop.
2663750	Abrupt messages are seen in engine log after complete storage failure in cvm resiliency scenario.

Table 1-5 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.
2649367	Kernel crashes in vx_fopen because of NULL pointer dereference.
2646936	The replication process dumps core when shared extents are present in the source file system.
2646930	Permission denied errors(EACCES) seen while doing I/O's on nfs shared filesystem.
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.
2634483	On RHEL6U1 writing to VxFS /proc hidden interface fails with EINVAL.
2630954	The fsck(1M) command exits during an internal CFS stress reconfiguration testing.
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.

Table 1-5 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.
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 1-6 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	had core dumped with sigabrt when panic'ed node came back
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. Cluster nodes may panic randomly while online/offline/switch 'ing operations of different service groups.
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 1-7 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
```

Error messages in syslog (1630188)

If you install or uninstall a product on a node, you may see the following warnings in syslog: /var/log/message. These warnings are harmless and can be ignored.

```
Jul  6 10:58:50 swlx62 setroubleshoot: SELinux is preventing the
semanage from using potentially mislabeled files
```

```
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For
complete SELinux messages. run sealert -l ed8978d1-0b1b-4c5b-a086-
67da2a651fb3
Jul  6 10:58:54 swlx62 setroubleshoot: SELinux is preventing the
semanage from using potentially mislabeled files
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For
complete SELinux messages. run sealert -l ed8978d1-0b1b-4c5b-a086-
67da2a651fb3
Jul  6 10:58:59 swlx62 setroubleshoot: SELinux is preventing the
restorecon from using potentially mislabeled files
```

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.

Ignore certain errors after an operating system upgrade—after a product upgrade with encapsulated boot disks (2030970)

Ignore certain errors after an operating system upgrade after a product upgrade with encapsulated boot disks.

You can ignore the following errors after you upgrade the operating system after a product upgrade that occurred with an encapsulated boot disk. Examples of the errors follow:

```
The partioning on disk /dev/sda is not readable by
The partioning tool parted, which is used to change the
partition table.
You can use the partitions on disk /dev/sda as they are.
You can format them and assign mount points to them, but you
cannot add, edit, resize, or remove partitions from that
disk with this tool.
```

Or

```
Root device: /dev/vx/dsk/bootdg/rootvol (mounted on / as reiserfs)
Module list: pilix mptspi qla2xxx silmage processor thermal fan
reiserfs aedd (xennet xenblk)
```

```
Kernel image; /boot/vmlinuz-2.6.16.60-0.54.5-smp
Initrd image: /boot/initrd-2.6.16.60-0.54.5-smp
```

The operating system upgrade is not failing. The error messages are harmless.

Workaround: Remove the /boot/vmlinuz.b4vxvm and /boot/initrd.b4vxvm files (from an un-encapsulated system) before the operating system upgrade.

Upgrading from Veritas Storage Foundation 5.1 Service Pack 1 Rolling Patch 2 to 6.0 with rootability enabled fails (2581313)

Upgrading from Veritas Storage Foundation (SF) 5.1 Service Pack (SP) 1 Rolling Patch (RP) 2 to 6.0 while using an encapsulated root disk fails because the post installation scripts of Veritas Volume Manager (VxVM) are unable to start the `initrd` daemon.

Workaround: To upgrade from 5.1 SP1 RP2 to 6.0 while using an encapsulated root disk, you must reinstall the nash utility on the system prior to the upgrade.

To upgrade from 5.1 SP1 RP2 to 6.0 while using an encapsulated root disk

1. Encapsulate the root disk.
2. Reinstall the nash utility.
3. Upgrade to the SF 6.0 release.

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&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.

Issue with soft links getting deleted in a manual upgrade (2115662)

While performing a manual upgrade (from 5.1 to 6.0) of the VRTSvlic RPM, some of the soft links created during your previous installation are deleted. As a result, `vxkeyless` binary is not found in its specified path.

To prevent this, use the `--nopreun` option.

For example: `rpm -Uvh --nopreun VRTSvlic-3.02.60.007-0.x86_64.rpm`

Manual upgrade of VRTSvlic RPM loses keyless product levels [2115662]

If you upgrade the `VRTSvlic` RPM 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 RPM.

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 RPM.

   ```
   # rpm -Uvh --nopreun VRTSvlic-3.02.60.007-0.x86_64.rpm
   ```

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.

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`.

The Configure Sybase ASE CE Instance in VCS option creates duplicate service groups for Sybase binary mount points (2560188)

The CPI installer does not check to see if Sybase binary mount points are already configured on systems, nor does it give an error message. It creates a duplicate service group for Sybase binary mount points.

This issue will be resolved in a later release.

Veritas Dynamic Multi-pathing known issues

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

DMP fast recovery not supported with RHEL 6 (2221507)

On RHEL6, DMP fast recovery caused issues with the following configurations:

- 3PAR A/A array plus QLogic 2462 HBA
- USPV A/A array plus QLogic 2462 HBA

Resolution:

Set the DMP tunable dmp_fast_recovery to off.

To disable dmp_fast_recovery, run the following command:

```
# vxdmpadm settune dmp_fast_recovery=off
```

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 1-8 shows the Hitachi arrays that have new array names.

Table 1-8 Hitachi arrays with new array names

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

DMP disables subpaths and initiates failover when an iSCSI link is failed and recovered within 5 seconds. (2100039)

When using iSCSI S/W initiator with an EMC CLARiiON array, iSCSI connection errors may cause DMP to disable subpaths and initiate failover. This situation occurs when an iSCSI link is failed and recovered within 5 seconds.

Work around:

When using iSCSI S/W initiator with an EMC CLARiiON array, set the node.session.timeo.replacement_timeout iSCSI tunable value to 40 secs or higher.

DMP marks the subpaths as DISABLED while these subpaths are accessible from OS level (2037222)

For iSCSI devices on SLES 10 SP3, the DMP tunable parameter dmp_fast_recovery needs to be turned off.

```
# vxdmpadm settune dmp_fast_recovery=off
```

DMP panics if a DDL device discovery is initiated immediately after loss of connectivity to the storage (2040929)

When using EMC Powerpath with VxVM 5.1SP1 on SLES11, set the fast_io_fail_tmo on the HBA port to any non-zero value that is less than the dev_loss_tmo value so as to avoid a panic in case a DDL device discovery is initiated by the `vxdisk scandisks` command or the `vxdctl enable` command immediately after loss of connectivity to the storage.

Upgrading the Linux kernel when the root volume is under DMP control

This section includes the procedures for upgrading the Linux kernel when the root volume is under DMP control.

Linux kernel can be upgraded on RHEL5 systems without turning off the DMP native support. Only one reboot is required to bring system LVM volume on DMP after kernel upgrade.

To update the kernel on a RHEL5 system

1 Update kernel with the rpm command.

```
# rpm -ivh kernel_rpm
```

2 Turn on the dmp_native_support tunable:

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

This enables booting with new kernel with LVM devices with DMP.

3 Reboot.

On SLES10 or SLES11

On SLES, the kernel can not be upgraded in a single reboot due to limitation in mkinitrd command.

To update the kernel on a SLES10 or SLES11 system

1 Turn off DMP native support

```
# vxdmpadm settune dmp_native_support=off
```

2 Reboot the system.

3 Upgrade kernel using the rpm command

```
# rpm -ivh kernel_rpm
```

4 Turn on DMP native support.

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

5 Reboot the system to bring the root LVM volume under DMP control.

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.

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
```

After rebooting the array controller for CX4-240-APF array, I/O errors occur on shared file systems (2616315)

For Linux hosts, rebooting the array controller for a CX4-240-APF array may result in I/O errors on shared file systems.

Work-around:

To work around this issue

- Set the tunable parameter dmp_lun_retry_timeout to 120 seconds before rebooting the array controller.

  ```
  # vxdmpadm settune dmp_lun_retry_timeout=120
  ```

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).

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.

Unable to upgrade the kernel on an encapsulated boot disk (2612301)

Upgrading the kernel on an encapsulated boot disk does not work if the device name changes across a reboot.

Workaround: Perform following procedure on the system with the encapsulated root disk to upgrade kernel.

To upgrade the kernel on a system with an encapsulated root disk

1 Unroot the encapsulated root disk:

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

2 Upgrade the kernel:

```
# rpm -Uvh Kernel-upgrade_version
```

3 Reboot the system.

4 Re-encapsulated the root disk:

```
# /etc/vx/bin/vxencap -c -g root_diskgroup rootdisk=root_disk
```

vxvmconvert displays an error message while converting an LVM volume (2682491)

The `vxvmconvert` command displays the following error message while converting an LVM volume:

```
Preparing VG aaa for conversion, (possibly migrating extents) please wait.
.ERROR!
```

This message only displays on RHEL5.7 and RHEL5.

Workaround: Upgrade the operating system from RHEL5 to RHEL6 or a later version.

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.

device.map must be up to date before doing root disk encapsulation (2202047)

If you perform root disk encapsulation while the `device.map` file is not up to date, the `vxdiskadm` command displays the following error:

```
VxVM vxencap INFO V-5-2-5327 Missing file: /boot/grub/device.map
```

Workaround:

Before you perform root disk encapsulation, run the the following command to regenerate the device.map file:

```
grub-install --recheck /dev/sdb
```

Post encapsulation of the root disk, system comes back up after first reboot unencapsulated (2119038)

In some cases, after encapsulating the root disk and rebooting the system, it may come up without completing the encapsulation. This happens because the `vxvm-reconfig` startup script is unable to complete the encapsulation process.

Workaround

Reboot the system or run the following command.

```
# service vxvm-reconfig start
```

This will reboot the system and complete the remaining stages of encapsulation.

Required attributes for Veritas Volume Manager (VxVM) devices to avoid boot failures (1411526)

To support iSCSI devices, Veritas Volume Manager (VxVM) does not start non-root devices until runlevel2. The boot process expects all local (non-NFS) mount points in the `/etc/fstab` file to be present at boot time. To avoid boot failures, all VxVM entries in the `/etc/fstab` file must have the _netdev attribute, and must not have the fsck required flag set. These attributes enable VxVM to defer mounting of VxVM devices until after VxVM has started.

vxrestored daemon fails to restore disabled paths (1663167)

The `vxrestored` daemon fails to restore disabled paths on RHEL 5 with direct attached disks.

Workaround:

Enable the `mpt_disable_hotplug_remove` tunable so that path level failover and failback function properly on RHEL 5 machines with direct attached disks.

To enable the mpt_disable_hotplug_remove tunable

1 Edit the /etc/modprobe.conffile and add the following line to the end of the file:

```
options mptsas mpt_disable_hotplug_remove=0
```

2 Rebuild the initrd image:

```
# mkinitrd -f /boot/initrd-`uname -r`.img `uname -r`
```

3 Reboot the system.

System hangs or panics after disabling 3 of 4 arrayside ports (1724260)

The system hangs or panics after you disable 3 of 4 arrayside ports.

Workaround:

This issue is fixed with a Novell patch for SLES 11 as indicated in Bugzilla ID 524347:

https://bugzilla.novell.com/show_bug.cgi?id=524347

Machine fails to boot after root disk encapsulation on servers with UEFI firmware (1842096)

Certain new servers in the market such as IBM x3650 M2, Dell PowerEdge T610, come with support for the UEFI firmware. UEFI supports booting from legacy MBR type disks with certain restrictions on the disk partitions. One of the restrictions is that each partition must not overlap with other partitions. During root disk encapsulation, it creates an overlapping partition that spans the public region of the root disk. If the check for overlapping partitions is not disabled from the UEFI firmware, then the machine fails to come up following the reboot initiated after running the commands to encapsulate the root disk.

Workaround:

The following workarounds have been tested and are recommended in a single-node environment.

For the IBM x3650 series servers, the UEFI firmware settings should be set to boot with the "Legacy Only" option.

For the Dell PowerEdge T610 system, set "Boot Mode" to "BIOS" from the "Boot Settings" menu.

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
   ```

Root disk encapsulation issue (1603309)

Encapsulation of root disk will fail if it has been assigned a customized name with vxdmpadm(1M) command. If you wish to encapsulate the root disk, make sure that you have not assigned a customized name to its corresponding DMP node.

See the vxdmpadm(1M) manual page.

See the "Setting customized names for DMP nodes" section of the *Veritas Storage Foundation Administrator's Guide*.

VxVM starts before OS device scan is done (1635274)

While working with some arrays, VxVM may start before all devices are scanned by the OS. This slow OS device discovery may result in malfunctioning of VM, fencing and VCS due to partial disks seen by VxVM.

Workaround:

After the fabric discovery is finished, issue the `vxdisk scandisks` command to bring newly discovered devices into the VxVM configuration.

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
```

Converting LVM volumes to VxVM volumes fails when multipathed storage devices are present (1471781, 1931727)

The `vxvmconvert` utility cannot convert LVM volumes to VxVM volumes when multipathed storage devices are present. This issue occurs with SLES 11 and RHEL5, due to changes in the LVM utilities. If multipathed devices are detected, the `vxvmconvert` utility exits with the following error:

```
vxvmconvert cannot convert multipathed devices on SLES11 systems.
... Exiting.
```

Workaround: There is no workaround for this issue.

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.

Removing a volume from a thin LUN in an alternate boot disk group triggers disk reclamation (2080609)

If you remove a volume from an alternate boot disk group on a thin LUN, this operation triggers thin reclamation, which may remove information required for the disk to be bootable. This issue does not affect the current boot disk, since VxVM avoids performing a reclaim on disks under the bootdg.

Workaround: If you remove a volume or plex from an alternate boot disk group with the `vxedit` command, specify the `-n` option to avoid triggering thin reclamation. For example:

```
# vxedit -g diskgroup -rfn rm volumename
```

The vxdiskunsetup operation fails the first attempt on EMC powerpath devices (2424845)

Performing `vxdiskunsetup` for the first time on EMC powerpath devices displays an error "Internal Configuration daemon error : disk destroy failed."

Workaround: Retry `vxdiskunsetup` using the same command to resolve the 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.

System panic occurs on RHEL6.0 when VxVM is installed on EMC PowerPath managed devices (2573229)

System panic occurs on RHEL6.0 due to page cache issue when VxVM is installed on top of EMC PowerPath managed devices.

Workaround :

Install VxVM before installing and configuring EMC PowerPath

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.

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
```

On SLES11, sometimes encap fails in partition phase during first reboot (2598859)

On SLES11, sometimes the encapsulation of a root disk fails during the re-partitioning phase, with the following error message:

```
Starting vxvm-reconfig
The Volume Manager is now reconfiguring (partition phase)...
Volume Manager: Partitioning sda as an encapsulated disk.
VxVM vxedpart ERROR V-5-1-13204 partition pre-verification failed : Device
resource busy

VxVM vxedroot ERROR V-5-2-3684 re-partitioning of rootdisk failed!
The partitioning of sda failed.
The following disks failed encapsulation:
sda
All changes the Volume Manager has made during this reconfiguration will be
reversed.
undo sda....
```

There is no work-around for this issue.

During system boot, some VxVM volumes fail to mount (2622979)

During system boot, some VxVM volumes that exist in the `/etc/fstab` file fail to mount with the following error messages:

```
# fsck
Checking all file systems.
  error on stat() /dev/vx/dsk//volume: No such
file or directory
```

The load order of kernel modules in Linux results in the VxFS file system driver loading late in the boot process. Since the driver is not loaded when the `/etc/fstab` file is read by the operating system, file systems of the type vxfs will not mount.

Workaround:

To resolve the failure to mount VxFS file systems at boot, specify additional options in the `/etc/fstab` file. These options allow the filesystems to mount later in the boot process. An example of an entry for a VxFS file system:

```
/dev/vx/dsk/testdg/testvolume /mountpoint  vxfs  _netdev,hotplug  1 1
```

To resolve the issue, the fstab entry for VxVM data volumes should be as per following template:

```
/dev/vx/dsk/testdg/testvol    /testmnt    vxfs   _netdev     0 0
```

System booting can take longer than expected (2486301)

System booting can sometimes take longer than expected due to hwpath generation.

Workaround: There is no workaround for this issue.

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.

Performance degradation seen on a CFS filesystem while reading from a large directory (2644485)

Performance degradation is seen on a CFS filesystem while reading from a large directory.

Workaround: There is no workaround.

On Linux system unable to select ext4 from the FileSystem (2691654)

On Linux system unable to select ext4 from the FileSystem.

Workaround: There is no workaround.

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
```

Taking a FileSnap over NFS multiple times with the same target name can result in the 'File exists' error (2353352)

The "File exists" error occurs as a result of the caching behavior of the NFS client. Because the link operation is successful, the NFS client assumes that a file with the specified target name, such as `file2::snap:vxfs:`, was created.. As a result, the NFS client caches a file with this name.

Workaround: Remove the target file after a snapshot is created. This forces the NFS client to remove the name from the cache. For example:

```
# ln file1 file2::snap:vxfs:
# rm file2::snap:vxfs:
```

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.

Task blocked messages display in the console for RHEL6 (2560360)

On RHEL6, the kernel occasionally displays messages in the console similar to the following example:

```
INFO: task seq:16957 blocked for more than 120 seconds.
```

These messages display because the task is blocked for a long time on sleep locks. However, the task is not hung and the messages can be safely ignored.

Workaround: You can disable these messages by using the following command:

```
# echo 0 > /proc/sys/kernel/hung_task_timeout_secs
```

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.

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.

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.

vfradmin abortjob does not kill the file replication process (2527916)

The `vfradmin abortjob` command does not kill the file replication process and may leave the checkpoints mounted.

Workaround: There is no workaround for this issue.

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 -t 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 -t vxfs /dev/vx/dsk/dg/data_volume
```

Running SUSE Linux and using Novell's YaST tool to configure an IPv6 address may result in an error (1679261)

When Novell's YaST tool is invoked to configure an IPv6 address on a different network interface and if:

- the host name, the DNS server name and domain name are specified to the YaST tool.
- IPv6 address is assigned by the Dynamic Host Configuration Protocol (DHCP).
- the "Write Hostname to /etc/hosts" option is selected (this is selected by default).

This results in the `vradmin` command returning the following error:

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

This happens because the YaST tool can replace the `/etc/hosts` entry containing `127.0.0.2` from the IPv4 host name to the specified new IPv6 host name. For example:

```
127.0.0.2 v6hostname.space.ipv6.com v6hostname
```

Workaround: The following procedure resolves this issue.

To resolve this issue

1. Edit the `/etc/hosts` file to specify the correct IPv6 address.
2. Restart the `vradmind` daemon on all VVR hosts:

   ```
   # /etc/init.d/vras-vradmind.sh restart
   ```

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:

```
# /etc/init.d/vras-vradmind.sh restart
```

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 restart
```

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 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 to SFHA or SF for Oracle RAC 6.0 RP1.

When upgrading from SFHA version 5.0 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.

Stale NFS file handle on the client across failover of a VCS service group containing LVMLogicalVolume resource

A VCS service group for a LVM volume group will be online automatically after a failover. However, the client applications may fail or be interrupted by stale NFS file handle error.

Workaround: To avoid the stale NFS file handle on the client across service group failover, specify "fsid=" in the Options attribute for Share resources.

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.

Migration of guest VM on native LVM volume may cause libvirtd process to terminate abruptly [2582716]

When the guest VM image is on native LVM volume, then the migration of that guest initiated by the administrator may cause `libvirtd` process to terminate abruptly.

Workaround: Start the `libvird` process manually.

Operational issues for VCS

LVMLogicalVolume online entry point stops responding and times out for mirrored volumes on SLES10 [2077294]

`LVMLogicalVolume` uses `lvchange` command to activate the logical volumes. In case of mirrored volumes, the `lvchange` command itself stops responding when it is invoked through a script. This causes online entry point to time out and the online entry point of `LVMLogicalVolume` resource stops responding. This is an issue with the SLES10.

LVM SG transition fails in all paths disabled status [2081430]

If you have disabled all the paths to the disks, the `LVM2 vg` commands stop responding and wait until at least one path to the disks is restored. As `LVMVolumeGroup` agent uses `LVM2` commands, this behavior causes online and offline entry points of `LVMVolumeGroup` agent to time out and `clean EP` stops responding for an indefinite time. Because of this, the service group cannot fail over to another node.

Workaround: You need to restore at least one path.

SG goes into Partial state if Native LVMVG is imported and activated outside VCS control

If you import and activate LVM volume group before starting VCS, the `LVMVolumeGroup` remains offline though the `LVMLogicalVolume` resource comes online. This causes the service group to be in a partial state.

Workaround: You must bring the VCS `LVMVolumeGroup` resource offline manually, or deactivate it and export the volume group before starting VCS.

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]

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.

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

LVM Logical Volume will be auto activated during I/O path failure [2140342]

LVM Logical Volume gets auto activated during the I/O path failure. This causes the VCS agent to report "Concurrency Violation" errors, and make the resource groups offline/online temporarily. This is due to the behavior of Native LVM.

Workaround: Enable the LVM Tagging option to avoid this issue.

System panics after starting KVM virtualized guest or initiating KVMGuest resource online [2337626]

System panics when the KVM guest is started or when the KVMGuest resource online is initiated. This issue is rarely observed.

The issue is observed due to the file descriptor leak in the libvirtd process. The maximum file open limit of file descriptor for libvirtd process is 1024. You may sometimes observe that more than 1024 file descriptors are opened when the KVM guest is started. Therefore, if the maximum file open limit is crossed, any attempt to start the KVM guest or to open a new file causes the system to panic. VCS cannot control this behavior as it suspects a file descriptor leak in the libvirtd process.

Workaround: There is no definite resolution for this issue; however, you can check the number of files opened by the libvirtd process in `/proc/<pid of libvirtd>/fd/`. If the file count exceeds 1000, restart libvirtd with the following command:

```
/etc/init.d/libvirtd restart
```

KVMGuest monitor entry point reports resource ONLINE even for corrupted guest or with no OS installed inside guest [2394235]

The VCS KVMGuest monitor entry point reports resource state as ONLINE in spite of the operating system inside the guest being corrupted or even if no operating system is installed inside the guest. The VCS KVMGuest agent uses `virsh` utility to determine the state of the guest. When the guest is started, the `virsh` utility reports the state of the running guest as running. Based on this running state, VCS KVMGuest agent monitor entry point reports the resource state as ONLINE.

In case the operating system is not installed inside the guest or the installed operating system is corrupted, `virsh` utility still reports the guest state as running. Thus, VCS also reports the resource state as ONLINE. Since RedHat KVM does not provide the state of the operating system inside guest, VCS cannot detect the guest state based on the state of the operating system.

Workaround: No workaround for this known issue.

LVM logical volume may get stuck with reiserfs file system on SLES11 [2120133]

LVM logical volume may get stuck with reiserfs file system on SLES11 if the service group containing the logical volume is switched continuously between the cluster node.

This issue may be observed:

- During the continuous switching of the service group having the LVM logical volume with reiserfs file system.
- On SLES11 and with reiserfs file system only.
- Due to the behavior of device-mapper on SLES11.

However, the issue is not consistent. Sometimes, the device-mapper gets stuck while handling the logical volumes and causes the logical volume to hang. In such a case, LVM2 commands also fail to clear the logical volume. VCS cannot handle this situation as the LVM2 commands are unable to deactivate the hung logical volume.

Resolution: You must restart the system on which the logical volumes are stuck in this situation.

KVMGuest resource comes online on failover target node when started manually [2394048]

The VCS KVMGuest resource comes online on failover target node when VM guest started manually, even though the resource is online on the primary node.

Red Hat kernel-based virtual machine (KVM) allows you to start the guest using same guest image on multiple nodes. The guest image is residing on the cluster file system. If the guest image is stored on the cluster file system, then it becomes available on all the cluster nodes simultaneously.

If the KVMGuest resource of VCS has made the guest online on one node by starting it using the guest image on cluster file system and if you manually start the same guest on the other node, Red Hat KVM does not prevent you from doing so. However, as this particular guest is under VCS control, VCS does not allow the resource to be ONLINE on multiple nodes simultaneously (unless it is in parallel service group configuration). VCS detects this concurrency violation and brings down the guest on the second node.

Note: This issue is also observed with CVM raw volume.

Workaround: No workaround required in VCS. VCS concurrent violation mechanism handles this scenario appropriately.

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.

DiskReservation agent may call clean if you configure large number of its resources in a single service group [2336391]

DiskReservation agent may call clean if you configure large number of DiskReservation resource (more than 400 resources) in a single service group and try to offline the service group.

In a single service group configuration with more than 400 DiskReservation resources and equal number of Mount resources, the service group offline may

cause the DiskReservation agent to call clean entry point. This issue is not observed if you configure about 150 resources.

Workaround: No workaround.

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.

DiskGroup agent is unable to offline the resource if volume is unmounted outside VCS

DiskGroup agent is unable to offline the resource if volume is unmounted using the `umount -l` command outside VCS.

A service group contains DiskGroup, Volume and Mount resources and this service group is online. Volume is mounted by Mount resource with VxFSMountLock enabled. An attempt to manually unmount the volume using `umount -l` system command causes the mount point to go away; however, the file system lock remains as it is. The volume cannot be stopped as it is mount locked and hence the disk group cannot be imported. This causes the disk group resource to go into UNABLE to OFFLINE state. Also, any attempt to again mount the file system fails, because it is already mount locked. This issue is due to file system behavior on Linux.

Workaround: Do not use `umount -l` command to unmount the VxFS file system when the mount lock is enabled. Instead, first unlock the mount point using the `/opt/VRTS/bin/fsadm` command and then unmount the file system.

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.

VVR setup with FireDrill in CVM environment may fail with CFSMount Errors [2564411]

When you try to bring the FireDrill service group online through Java Console or `hagrp -online` command, the CFSMount resource goes into faulted state.

Workaround: Run the `fsck` command.. You can find these commands in the engine logs.

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.

RVGsnapshot agent does not work with volume sets created using vxvset [2553505]

RVGsnapshot agent does not work with volume sets created using `vxvset`. This happens during FireDrill in a VVR einviroment.

Workaround: No workaround.

No log messages in engine_A.log if VCS does not find the Monitor program [2563080]

No message is logged in the engine_A.log, when VCS cannot find the Monitor program with KVM guest with service group online.

Workaround: In case resource state is unknown , also refer to agent log files for messages.

No IPv6 support for NFS [2022174]

IPv6 is not supported for NFS.

Workaround: No workaround.

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, DNS and LogicalVolumeGroup 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.

Guest virtual machine may fail on RHEL 6.1 if KVM guest image resides on CVM-CFS

If a KVM guest image file resides on CVM-CFS, the migration of that guest virtual machine may fail with "Permission Denied" error on RHEL 6.1. This causes guest virtual machine to go in "shut-off" state on both source and destination node, and the associated VCS KVMGuest. [2615089]

Workaround: Make sure that the guest image file is having 777 permission.

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.

No health check monitoring for Oracle agent on SLES11 platform [1938167]

Oracle agent does not support health check monitoring on SLES11 platform.

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.

Issues related to the agent framework

Agent may fail to heartbeat under heavy load [2073018]

An agent may fail to heart beat with the VCS engine under heavy load.

This may happen when agent does not get enough CPU to perform its tasks and when the agent heartbeat exceeds the time set in the AgentReplyTimeout attribute. The VCS engine therefore stops the agent and restarts it. The VCS engine generates a log when it stops and restarts the agent.

Workaround: If you are aware that the system load is likely to be high, then:

- The value of AgentReplyTimeout attribute can be set to a high value
- The scheduling class and scheduling priority of agent can be increased to avoid CPU starvation for the agent, using the AgentClass and AgentPriority attributes.

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 detect when bonded NICs come up (2604437)

When LLT is configured over a bonded NIC and that bonded NIC is DOWN with the `ifconfig` command, LLT marks the corresponding link down. When the bonded

NIC is UP again using the `ifconfig` command, LLT fails to detect this change and marks the link up.

Workaround: Close all the ports and restart LLT, then open the ports again.

LLT connections are not formed when a vlan is configured on a NIC (2484856)

LLT connections are not formed when a vlan is configured on a NIC that is already used to configure an LLT link.

Workaround: Do not specify the MAC address of a NIC in the `llttab` file while configuring LLT if you want to configure a vlan later. If you have already specified the MAC address of a NIC, then delete the MAC address from the `llttab` file, and update the file before you restart LLT.

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 RPM 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.

Veritas Cluster Server agents for Veritas Volume Replicator known issues in

The following are new additional Veritas Cluster Server agents for Veritas Volume Replicator known issues in release.

fdsetup cannot correctly parse disk names containing characters such as "-" (1949294)

The fdsetup cannot correctly parse disk names containing characters such as "-".

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).

Cluster Manager (Java Console) may display an error while loading templates (1433844)

You can access the Template View in the Cluster Manager from the Tools > Templates menu. If you have Storage Foundation configured in a VCS cluster setup, the following error may occur while the Cluster Manager loads the templates.

```
VCS ERROR V-16-10-65 Could not load :-
/etc/VRTSvcs/Templates/DB2udbGroup.tf
```

Workaround: Ignore the error.

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.

Upgrading VBS from VBS 6.0 to 6.0 RP1 deletes the /opt/VRTSvbs and /etc/VRTSvbs directories (2760273)

When VBS is upgraded from VBS 6.0 to 6.0 RP1 using the CPI VCS 6.0 RP1 installer, the /opt/VRTSvbs and /etc/VRTSvbs directories get deleted during the upgrade.

Workaround:

Perform the following steps to upgrade VBS from 6.0 to 6.0 RP1:

1. Verify VBS version on the cluster.

```
# vbsd -version
```

2. Check the state of VBS.

```
# vbssvc -display
```

3. Disable fault management from VOM.
4. Check the state of VBS again.

```
# vbssvc -display
```

5. Offline the VBS application resource.

```
# hares -offline vbsapp -sys system_name
```

6. Make the cluster writable.

```
# haconf -makerw
```

7. Delete the vbsapp resource.

```
# hares -delete vbsapp
```

8. Save the cluster configuration to `main.cf`.

```
# haconf -dump -makero
```

9. Upgrade using CPI to 6.0 RP1.
10. Check the version of VRTSvbs package.

```
# rpm -qa|grep vbs
```

11. Remove the VRTSvbs package.

```
# rpm -e VRTSvbs
```

12 Install 6.0 RP1 VRTSvbs package on all cluster nodes from the rpms directory location where the bits for VCS 6.0 RP1 were downloaded or saved.

```
# rpm -ivh rpms/VRTSvbs-6.0.001.000-RP1_os_type.rpm
```

13 Refresh all the nodes from VOM and check the version of the rpm installed.

```
# rpm -qa|grep vbs
```

Note: Refer to the latest *Veritas Operations Manager Administrator's Guide* for steps on how to refresh nodes from VOM GUI.

14 Enable fault management from VOM GUI for the VBS.

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 `halogin`command, VCS stores encrypted authentication information in the user’s home directory.

NFS resource might not come online while configuring CNFS share (2488685)

If SELinux is configured as `enforcing` or `permissive`, NFS resource might not come online and go into `FAULTED` state while configuring CNFS share `cfsnfssg` service group.

Sample output:

```
# hastatus -sum

-- SYSTEM STATE
-- System               State                Frozen

A  swlx14               RUNNING              0

-- GROUP STATE
-- Group           System   Probed     AutoDisabled    State

B  cfsnfssg        swlx14   Y          N               OFFLINE|FAULTED
B  cfsnfssg_dummy  swlx14   Y          N               OFFLINE
B  cvm             swlx14   Y          N               ONLINE
B  vip1            swlx14   Y          N               OFFLINE

-- RESOURCES FAILED
-- Group           Type                 Resource             System

D  cfsnfssg        NFS                  nfs                  swlx14
```

Workaround

To resolve this issue you need to add the Ethernet port into the trusted list for SELinux.

- In the System Setup->Firewall configuration, select customize.
- In the Trusted device, select the Ethernet port.

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.

cvm_clus resource goes into faulted state after the resource is manually panicked and rebooted in a 32 node cluster (2278894)

The `cvm_clus` resource goes into faulted state after the resource is manually panicked and rebooted in a 32 node cluster.

Workaround: There is no workaround for this issue.

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.

tail -f run on a cluster file system file only works correctly on the local node (2613030)

When using the `tail -f` command to monitor a file on a cluster file system, changes to the file made on remote nodes are not detected. This is due to the `tail` command now utilizing `inotify`. Symantec is currently unable to support `inotify` with a cluster file system due to GPL restrictions.

Workaround: To revert to the old behavior, you can specify the `---disable-inotify` option with the `tail` command.

"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.

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 -t 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 -t vxfs /dev/vx/dsk/dg/data_volume
```

Clone command errors in a Data Guard environment using the MEMORY_TARGET feature for Oracle 11g (1824713)

The `dbed_vmclonedb` command displays errors when attempting to take a clone on a STANDBY database in a dataguard environment when you are using the MEMORY_TARGET feature for Oracle 11g.

When you attempt to take a clone of a STANDBY database, the `dbed_vmclonedb` displays the following error messages:

```
Retrieving snapshot information ...                              Done
Importing snapshot diskgroups ...                                Done
Mounting snapshot volumes ...                                    Done
Preparing parameter file for clone database ...                  Done
Mounting clone database ...
ORA-00845: MEMORY_TARGET not supported on this system


SFDB vxsfadm ERROR V-81-0612 Script
/opt/VRTSdbed/applications/oracle/flashsnap/pre_preclone.pl failed.
```

This is Oracle 11g-specific issue known regarding the MEMORY_TARGET feature, and the issue has existed since the Oracle 11gr1 release. The MEMORY_TARGET feature requires the `/dev/shm` file system to be mounted and to have at least 1,660,944,384 bytes of available space. The issue occurs if the `/dev/shm`file system is not mounted or if the file system is mounted but has available space that is less than the required minimum size.

Workaround

To avoid the issue, remount the `/dev/shm` file system with sufficient available space.

To remount the /dev/shm file system with sufficient available space

1. Shut down the database.
2. Unmount the `/dev/shm` file system:

   ```
   # umount /dev/shm
   ```

3. Mount the `/dev/shm` file system with the following options:

   ```
   # mount -t tmpfs shmfs -o size=4096m /dev/shm
   ```

4. Start the database.

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.

The following message may be displayed on running the `strace` command:

```
# /usr/bin/strace -ftt -p pid_of_ohasd.bin
14:05:33.527288 open("/var/tmp/.oracle/npohasd",
O_WRONLY <unfinished ...>
```

For possible causes and workarounds, see the Oracle Metalink document: 1069182.1

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.

Veritas Storage Foundation for Sybase ASE CE known issues

This section describes the known issues in this release of Veritas Storage Foundation for Sybase ASE CE.

Stale .vxfendargs file lets hashadow restart vxfend in Sybase mode (2554886)

When I/O fencing is configured in customized mode, vxfend, the user mode daemon of I/O fencing, creates the `/opt/VRTSvcs/lock/.vxfendargs` file. VCS uses this file to restart the vxfend daemon when it gets killed. However, VCS does not use this file when I/O fencing is configured in Sybase mode. This file is not removed from the system when I/O fencing is unconfigured.

If user configures I/O fencing in Sybase mode and an old `/opt/VRTSvcs/lock/.vxfendargs` file is present in the system from an earlier

configuration of I/O fencing in customized mode, then VCS attempts to restart the vxfend daemon every time it is killed. This interferes with the functioning of I/O fencing in the Sybase mode.

Workaround: Before you configure I/O fencing in Sybase mode, delete the `/opt/VRTSvcs/lock/.vxfendargs` file if it is present in the system.

Sybase Agent Monitor times out (1592996)

Problem: The Sybase Agent Monitor has issues with timing out, in cases where qrmutil reports delay.

The Sybase Agent monitor times out, if qrmutil fails to report the status to the agent within the defined MonitorTimeout for the agent.

Solution: If any of the following configuration parameters for Sybase Database is increased, it will require a change in its MonitorTimeout value:

- quorum heartbeat interval (in seconds)
- Number of retries

If the above two parameters are changed, Symantec recommends that the MonitorTimeout be set to a greater value than the following: ((number of retries + 1) * (quorum heartbeat interval)) + 5.

Installer warning (151550)

Problem: During configuration of Sybase instance under VCS control, if the quorum device is on CFS and is not mounted, the following warning message appears on the installer screen:

```
Error: CPI WARNING V-9-40-4563 The quorum file /qrmmnt/qfile
cannot be accessed. This may be due to a file system not being mounted.
```

The above warning may be safely ignored.

Unexpected node reboot while probing a Sybase resource in transition (1593605)

Problem: A node may reboot unexpectedly if the Sybase resource is probed while the resource is still in transition from an online to offline state.

Normally the monitor entry point for Sybase agent completes with 5-10 seconds. The monitor script for the Sybase agent uses the qrmutil binary provided by Sybase. During a monitor, if this utility takes longer time to respond, the monitor entry point will also execute for longer duration before returning status.

Resolution: During the transition time interval between online and offline, do not issue a probe for the Sybase resource, otherwise the node may reboot.

Unexpected node reboot when invalid attribute is given (2567507)

Problem: A node may reboot unexpectedly if the Home, Version, or Server attributes are modified to invalid values while the Sybase resources are online in VCS.

Resolution: Avoid setting invalid values for the Home, Version, or Server attributes while the Sybase resources are online in VCS, to avoid panic of the node.

Deporting issues with shared disk groups

If you manually deport a shared disk group, the CVMVolDg agent does not automatically reimport it as a shared disk group. You must manually reimport it as a shared disk group.

Resources not brought online after had process is killed and restarted

If the `had` process is killed either manually or by some other means, the `hashadow` process restarts it. During `had` restart time, the resources remain in the same state unless their states are changed manually, but after restarting the `had`, process VCS may not bring some resources online.

To work around this issue

- Manually restart the resources that were not started during `had` restart.

AutoFailOver = 0 attribute absent in the sample files at `/etc/VRTSagents/ha/conf/Sybase` (2615341)

Problem: AutoFailOver = 0 attribute is not present in the sample files at `/etc/VRTSagents/ha/conf/Sybase`.

Resolution: If you copy the main.cf file from the `/etc/VRTSagents/ha/conf/Sybase` location, add the AutoFailOver = 0 attribute to the binmnt and sybasece service groups.

Symantec VirtualStore known issues

This section describes the known issues in this release of Symantec VirtualStore.

The cluster node may panic (2524087)

On SLES 10 SP4, the cluster node may panic while iSCSI initiator access the LUN from the target.

Workaround

There is no workaround at this time.

VirtualStore machine clones created while the VirtualStore cluster reboots will probably not start (2164664)

In some cases when you clone while rebooting the SVS nodes, you may receive several of the following error messages:

```
clone vms could not start X server
```

Workaround

Delete all the clones that got created while the node crashed and redo the cloning operation.

Cloning may not work (2348628)

If you cannot clone and you are using the VMware vAPP and OVF templates, then you must disable the vApp.

Workaround

To disable the vAPP

1. In VI Client, right-click on the virtual machine > **Edit Settings** > **Options** > **vApp Options**.
2. Click **Disable**.

Need intelligent NDMP/NBU backups for virtual machines (2378396)

When using NDMP or the NBU client to backup a virtual machine, the space consumed by the backup is equivalent to the size of the disks in the virtual machine, even though not all of the disk space in the virtual machine is used.

If a VMDK (Virtual Machine Disk) file is 10GB in size, but only consumes 1GB of disk space, an backup done by NDMP or the NBU client generates 10GB of backup data, even though the original VMDK file contains 9GB of unassigned disk space.

Workaround

Use VMware-specific backup applications (such as NetBackup for VMware) to create space-efficient backups.

The Symantec Quick Clone Virtual Machine Wizard may not behave as expected when multiple instances are open (2309702)

The wizard may not behave as expected, if you invoke multiple parallel session of the wizard from a single vSphere Client at the same time.

For example, if you do the following:

- Right-click wingoldvm1 and invoke the wizard.
- Then soon after, right-click slesgoldvm1 and invoke the wizard.

This causes you to have two instances of the wizard running from the same vSphere Client and can cause unexpected behavior.

Workaround

To resolve this issue:

- Close both instances of the wizard.
- Reopen a new instance of the wizard.

Virtual machines created by the Symantec Quick Clone Virtual Machine Wizard might not boot correctly if during the process the FileStore cluster node, the ESX Server, or the vCenter Server reboots (2164664, 2374229)

In some cases when you clone using the wizard, and one of the following servers crashes or reboots while the clone process is in progress, the clones might not get created correctly:

- FileStore nodes
- ESX host on which the clones are being created
- vCenter Server

Even if the clones appear in the vCenter inventory as created, the clones GuestOS might not be able to boot.

Workaround

Delete all of the clones that were created when the servers crashed or were rebooted, and redo the wizard operation.

Error message does not always display when you select an incorrect cluster to clone (2372713)

In cases where multiple FileStore clusters are registered with the same Virtual Center, the Symantec Quick Clone Virtual Machine Wizard might not provide a

warning that you selected an incorrect cluster to clone a golden image. This could happen if all of the FileStore clusters are exporting the same file system path, such as `/mnt`. Instead of an advanced warning that you selected the wrong cluster, you instead see an error on the final page of the wizard when the wizard attempts to clone the disks (vmdks) of the golden image. The error that displays is similar to the following example:

```
/mnt/goldvm/goldvm.vmdk no such file or directory...
```

Workaround

There is no workaround for this issue.

Software limitations

There are no software limitations in this release.

List of RPMs

Table 1-9 lists the Red Hat Enterprise Linux 5 RPMs that are updated in this release.

Table 1-9 RPMs for Red Hat Enterprise Linux 5

Name	Version	Products Affected	Patch Size in Bytes
VRTSamf	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	885160
VRTScavf	6.0.001.000	SF Oracle RAC, SFCFSHA, SVS	195047
VRTSfsadv	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS	4698367
VRTSfssdk	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS	353878
VRTSlvmconv	6.0.001.000	SF, SF Oracle RAC, SFCFSHA, SFHA, SVS, VM	72099
VRTSob	3.4.530	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS, VM	32434618
VRTSsfcpi60	6.0.001.000	DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS, VM	717793

Table 1-9 RPMs for Red Hat Enterprise Linux 5 *(continued)*

Name	Version	Products Affected	Patch Size in Bytes
VRTSsvs	6.0.001.000	SVS	65670991
VRTSvbs	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	17899295
VRTSvcs	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SVS, VCS	63611799
VRTSvxfs	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS	10711114
VRTSvxvm	6.0.001.000	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, SVS, VM	31133582

Table 1-10 lists the Red Hat Enterprise Linux 6 RPMs that are updated in this release.

Table 1-10 RPMs for Red Hat Enterprise Linux 6

Name	Version	Products Affected	Patch Size in Bytes
VRTSamf	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	613728
VRTScavf	6.0.001.000	SF Oracle RAC, SFCFSHA, SFSYBASECE, SVS	176312
VRTSfsadv	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	3807188
VRTSfssdk	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	293240
VRTSlvmconv	6.0.001.000	SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	65476
VRTSob	3.4.530	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	32434618

Table 1-10 RPMs for Red Hat Enterprise Linux 6 *(continued)*

Name	Version	Products Affected	Patch Size in Bytes
VRTSsfcpi60	6.0.001.000	DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS, VM	717793
VRTSsvs	6.0.001.000	SVS	65670991
VRTSvbs	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	17899295
VRTSvcs	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	47936932
VRTSvxfs	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	6671816
VRTSvxvm	6.0.001.000	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	20353816

Table 1-11 lists the SUSE Linux Enterprise Server 10 RPMs that are updated in this release.

Table 1-11 RPMs for SUSE Linux Enterprise Server 10

Name	Version	Products Affected	Patch Size in Bytes
VRTSamf	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	4770504
VRTScavf	6.0.001.000	SF Oracle RAC, SFCFSHA, SFSYBASECE, SVS	163498
VRTSfsadv	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	4867100
VRTSfssdk	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	308671

Table 1-11 RPMs for SUSE Linux Enterprise Server 10 *(continued)*

Name	Version	Products Affected	Patch Size in Bytes
VRTSlvmconv	6.0.001.000	SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	62769
VRTSob	3.4.530	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	32434618
VRTSsfcpi60	6.0.001.000	DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS, VM	717793
VRTSsvs	6.0.001.000	SVS	65670991
VRTSvbs	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	17899295
VRTSvcs	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	58655172
VRTSvxfs	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	13033103
VRTSvxvm	6.0.001.000	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	29649352

Table 1-12 lists the SUSE Linux Enterprise Server 11 RPMs that are updated in this release.

Table 1-12 RPMs for SUSE Linux Enterprise Server 11

Name	Version	Products Affected	Patch Size in Bytes
VRTSamf	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	1311939
VRTScavf	6.0.001.000	SF Oracle RAC, SFCFSHA, SFSYBASECE, SVS	168680

Table 1-12 RPMs for SUSE Linux Enterprise Server 11 *(continued)*

Name	Version	Products Affected	Patch Size in Bytes
VRTSfsadv	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	3452239
VRTSfssdk	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	269595
VRTSlvmconv	6.0.001.000	SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	64040
VRTSob	3.4.530	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	32434618
VRTSsfcpi60	6.0.001.000	DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS, VM	717793
VRTSsvs	6.0.001.000	SVS	65670991
VRTSvbs	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	17899295
VRTSvcs	6.0.001.000	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VCS	47493461
VRTSvxfs	6.0.001.000	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS	6015813
VRTSvxvm	6.0.001.000	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, SVS, VM	20408607

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.

Chapter 2

Installing the products for the first time

This chapter includes the following topics:

Installing the Veritas software using the script-based installer

This section describes how to install a 6.0 RP1 Veritas Storage Foundation and High Availability Solutions product for the first time on a host. 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 108.

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 107.

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 103.

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.

Chapter 3

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 107.

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), Veritas Storage Foundation for Oracle RAC (SFRAC) or Symantec VirtualStore (SVS) 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 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 101.

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 107.

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.

Note: You need to make sure that IPv4RouteOptions attribute is configured for MultiNICA resources, otherwise network connection may be interrupted.

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 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
```

where *node1* and *node2* are nodes which are to be upgraded.

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 SFCFS cluster

1 Make sure you have downloaded the latest software required for the upgrade.

2 Log in as superuser.

3 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
```

4 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
```

Note: If file system is CFS mounted then use `cfsumount` command.

---

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 on master node:

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

Verify that no volumes remain open:

```
# vxprint -Aht -e v_open
```

8 If required, apply the OS kernel patches.

9 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
```

where *node1* and *node2* are nodes which are to be upgraded.

10 After all nodes in the cluster are upgraded, the processes will be restarted automatically. Should there be any problem, the `installrp` script will ask you to reboot the system. Then the application failover capability will be available.

11 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.

12 Restart all the volumes by entering the following command on the master node, for each disk group:

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

13 If you stopped any RVGs in step 5, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

14 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

15 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

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 an 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/VRTSvcs/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
```

5 Enter the following command to freeze HA service group operations on each node:

```
# hasys -freeze -persistent nodename
```

6 Make the configuration read-only:

```
# haconf -dump -makero
```

7 If Oracle Clusterware is not controlled by VCS, enter the following command on each node of the cluster to stop Oracle Clusterware:

```
# $CRS_HOME/bin/crsctl stop crs
```

8 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.

If the applications are under VCS control:

```
# hagrp -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

9 For Oracle RAC 10g and Oracle RAC 11g:

Stop all Oracle RAC resources.

- If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

  ```
  # hagrp -offline group_name -any
  ```

- If the database instances are not managed by VCS, then run the following on one node:

  ```
  $ srvctl stop database -d db_name
  ```

10 - 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 -maker
  ```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

  ```
  $ srvctl modify database -d db-name -y manual
  ```

11 Unmount the VxFS file system, which is not under VCS control.

```
# mount |grep vxfs

# fuser -m /mount_point

# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

12 Stop VCS.

```
# hastop -all
```

13 If required, apply the OS kernel patches.

See *Oracle's* documentation for the procedures.

14 From the directory that contains the extracted and untarred 6.0 RP1 rolling patch binaries, change to the directory that contains the installrp script. If ssh key authentication is configured then enter:

```
# ./installrp  node1 node2
```

If ssh is not configured then enter:

```
# ./installrp -rsh node1 node2
```

where `node1` and `node2` are nodes which are to be upgraded.

15 Follow the instructions from the installer. If there is some module load/unload issue, reboot all of the nodes of the cluster.

```
# /sbin/shutdown -r now
```

16 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.

17 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.

18 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

19 Enter the following command on each node to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent nodename
```

20 Make the configuration read-only:

```
# haconf -dump -makero
```

21 Enter the following command on each node to take service groups online:

```
# hagrp -online service_group -sys nodename
```

22 Restart all the volumes by entering the following command for each disk group:

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

23 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

24 If Oracle Clusterware is not controlled by VCS, enter the following command on each node to start Oracle Clusterware.

```
# $CRS_HOME/bin/crsctl start crs
```

25 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
```

26 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
```

27 ■ 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_groupname 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
```

28 For upgrade scenarios that involve Oracle RAC 9i, start `gsd` as the Oracle user:

```
$ ORACLE_HOME/bin/gsdctl start
```

29 Upgrade Oracle RAC.

Note: Oracle RAC 11g Release 1 Clusterware is not supported. Make sure that you install Oracle RAC 11g Release 2 Grid Infrastructure in order to use the Oracle RAC 11g Release 1 database. All database versions starting from Oracle 10g Release 2 and later are supported.

For instructions, see the chapter *Upgrading Oracle RAC* in *Veritas Storage Foundation™ for Oracle® RAC Installation and Configuration Guide*.

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.

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

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

6 Unmount all Storage Checkpoints and file systems:

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

7 If 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 Navigate to the folder that contains the installation program. Run the `installrp` script:

```
# ./installrp nodename
```

12 If necessary, reinstate any missing mount points in the `/etc/fstab` 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
```

Upgrading to 6.0 RP1 on a system that has encapsulated boot disk

You can use this procedure to upgrade to 6.0 RP1 on a system that has encapsulated boot disk.

Note: Upgrading with encapsulated boot disk from 6.0 to 6.0 RP1 requires multiple reboots.

To upgrade to 6.0 RP1 on a system that has encapsulated boot disk

1. Manually unmount file systems and stop open volumes.
2. If required, manually break the mirror.
3. Upgrade to 6.0 RP1 using `installrp` command.
4. After upgrading, reboot the system to have the new VM drivers take effect.
5. If the mirrors of boot disk are split manually in step 2, re-join the mirrors manually when systems reboot after upgrading to 6.0 RP1.

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.

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 as subcluster1and nodeB as subcluster2.

To perform the rolling upgrade on kernel packages: phase 1

1. On the first sub-cluster, start the installer for the rolling upgrade with the `-rollingupgrade_phase1` option.

   ```
   # ./installrp -rollingupgrade_phase1 nodeA
   ```

2. Note that if the boot-disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
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 and starts all the relevant processes and brings all the service groups online.
5. If the boot disk is encapsulated, reboot the first sub-cluster's system. Otherwise go to step 6.

6 After rolling upgrade 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 rolling upgrade phase 1 by itself on nodeB.

If you choose `n` or `q`, you need to complete step 1 to step 4 on nodeB.

7 After rolling upgrade phase 1 of the cluster, the following message displays:

```
Would you like to perform rolling upgrade phase 2 on the cluster?
[y,n,q] (y)
```

- If you choose `y`, it continues to run rolling upgrade phase 2 of the cluster by itself. You don't need to run phase 2: See "Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSHA : phase 2" on page 121.
  After phase 2 upgrade, 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 use the following steps to finish rolling upgrade phase 2 of the cluster: See "Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSHA : phase 2" on page 121.

Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSHA : phase 2

In this phase installer installs all non-kernel packages 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, patch versions, product versions, and completes prechecks. It upgrades non-kernel patches. It also 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 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. Failing to perform this step results in the database attempting to come online after the upgrade; the attempt fails due to the presence of old libraries on the system.

```
# 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:

```
$ srvctl modify database -d db_name -y manual
```

4 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.

5 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
```

6 Switch over all failover service groups to the other nodes in the cluster:

```
# hagrp -switch grp_name -to node_name
```

7 Take all the VCS service groups offline:

```
# hagrp -offline grp_name -sys node_name
```

8 Unmount all the VxFS file system which is not under VCS control.

```
# mount |grep vxfs

# fuser -m /mount_point

# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

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.

---

9 On the sub-cluster, start the installer for the rolling upgrade with the `-rollingupgrade_phase1` option.

```
# ./installrp -rollingupgrade_phase1 nodeA nodeB
```

10 Note that if the boot-disk is encapsulated, you do not need to perform an unencapsulation for upgrades.

11 The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.

12 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.

13 If the boot disk is encapsulated, the installer strongly recommends a reboot of the nodes. Reboot the nodes as prompted by the installer.

Note: Before you reboot the nodes, ensure that the boot device is set to the disk containing the upgraded version of the product.

```
# eeprom
```

---

14 After rolling upgrade phase 1 is completed on nodeA and nodeB, the following message displays:

```
It is recommended to perform rolling upgrade phase 1 on the
systems nodeC and nodeD in the next step.
```

```
Would you like to perform rolling upgrade phase 1 on the systems?
[y,n,q] (y)
```

- If you choose `y`, first complete step 4 to step 8 on the remaining subcluster. Then it continues to run rolling upgrade phase 1 on nodeC and nodeD by itself.
- If you choose `n` or `q`, go to step 15.

15 After rolling upgrade phase 1 of the cluster, the following message displays:

```
Would you like to perform rolling upgrade phase 2 on the cluster?
[y,n,q] (y)
```

- If you choose `y`, it continues to run rolling upgrade phase 2 of the cluster 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 127. After rolling upgrade phase 2, complete step 14 to step 20 (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 14 to step 20 and run rolling upgrade phase 2: See "Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2" on page 127.

16 Manually mount the VxFS and CFS file systems that are not managed by VCS.

17 Bring the Oracle database service group online.

- If VCS manages the Oracle database:

  ```
  # hagrp -online oracle_group -sys node_name
  ```

- 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 4 to 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
```

Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2

In this phase, the 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, package 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.

Upgrading the operating system

This section describes how to upgrade the operating system on a Storage Foundation node where you plan to upgrade to 6.0 RP1. This section includes the following topics:

Upgrading RHEL 6

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 6.0 RP1.

To upgrade to a later version of RHEL 6

1. Stop Storage Foundation.
2. Upgrade to the latest update version of RHEL 6.
3. Upgrade to 6.0 RP1.

   See "Upgrading to 6.0 RP1" on page 108.
4. Start Storage Foundation.

Upgrading RHEL 5

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 6.0 RP1.

To upgrade to a later version of RHEL 5

1. Stop Storage Foundation.
2. Upgrade to the latest update version of RHEL 5.
3. Upgrade to 6.0 RP1.

   See "Upgrading to 6.0 RP1" on page 108.
4. Start Storage Foundation.

Upgrading RHEL 5

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 6.0 RP1.

To upgrade to a later version of RHEL 5

1. Stop Storage Foundation.
2. Upgrade to the latest update version of RHEL 5.
3. Upgrade to 6.0 RP1.

   See "Upgrading to 6.0 RP1" on page 108.
4. Start Storage Foundation.

Upgrading OEL 6

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 6.0 RP1.

To upgrade to a later version of OEL 6

1. Stop Storage Foundation.
2. Upgrade to the latest update version of OEL 6.
3. Upgrade to 6.0 RP1.

   See "Upgrading to 6.0 RP1" on page 108.
4. Start Storage Foundation.

Upgrading OEL 5

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 6.0 RP1.

To upgrade to a later version of OEL 5

1. Stop Storage Foundation.
2. Upgrade to the latest update version of OEL 5.
3. Upgrade to 6.0 RP1.

   See "Upgrading to 6.0 RP1" on page 108.
4. Start Storage Foundation.

Upgrading SLES 10

This section describes how to upgrade the OS on a Storage Foundation node and upgrade to 6.0 RP1.

To upgrade to a later version of SLES 10

1. Stop Storage Foundation.
2. Upgrade to the latest update version of SLES 10.
3. Upgrade to 6.0 RP1.

   See "Upgrading to 6.0 RP1" on page 108.
4. Start Storage Foundation.

Verifying software versions

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

```
# rpm -qa | egrep VRTS
```

Chapter 4

Removing Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

About removing Veritas Storage Foundation and High Availability Solutions 6.0 RP1

Symantec recommends that you follow the steps in the following sections to remove all the installed Veritas software, and then perform a complete reinstallation of the previous release.

For extensive uninstallation and reinstallation procedures, refer to the appropriate product's Installation Guide.

Uninstalling Veritas Storage Foundation Cluster File System High Availability 6.0 RP1

This section provides procedures for uninstalling Veritas Storage Foundation Cluster File System High Availability (SFCFSHA). You must complete the preparatory tasks before you uninstall SFCFS.

Preparing to uninstall Veritas Storage Foundation Cluster File System High Availability

The following procedure prepares your system for uninstalling Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

To prepare to uninstall Veritas Storage Foundation Cluster File System High Availability

1. Log in as the root user on any node in the cluster.
2. Verify that the following directories are set in your `PATH` environment variable:

   ```
   /opt/VRTS/bin
   /opt/VRTSvcs/bin
   ```

3. Back up the following configuration files:

   ```
   # mv /etc/llttab /etc/llttab.`date +%m-%d-%y_%H-%M-%S`
   # mv /etc/llthosts /etc/llthosts.`date +%m-%d-%y_%H-%M-%S`
   # mv /etc/gabtab /etc/gabtab.`date +%m-%d-%y_%H-%M-%S`
   # mv /etc/vxfenmode /etc/vxfenmode.`date +%m-%d-%y_%H-%M-%S`
   ```

4. Determine if each node's root disk is under VxVM control and proceed as follows.

   - Check if each node's root disk is under VxVM control:

     ```
     # df -v /
     ```

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

   - Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes mirrootvol-01, and mirswapvol-01 that are configured on a disk other than the root disk:

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

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

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices.

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

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

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
```

6 Check if any VxFS file systems or Storage Checkpoints are mounted:

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

7 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint1
# umount /filesystem1
```

If file system is mounted in a cluster, then use `cfsumount` command.

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

```
# vxvol -g dg1 stopall
```

To verify that no volumes are open:

```
# vxprint -Aht -e v_open
```

9 Stop VCS:

```
# hastop -all
```

Uninstalling Veritas Storage Foundation Cluster File System High Availability

The following procedure uninstalls Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

To uninstall Veritas Storage Foundation Cluster File System High Availability

1 Log in as the root user on any node in the cluster.

2 Navigate to the directory that contains the uninstallation program:

```
# cd /opt/VRTS/install
```

3 Run the uninstallation program while specifying each node in the cluster:

```
# ./uninstallsfcfsha node1 node2
```

4 Confirm the uninstallation:

```
Are you sure you want to uninstall SFCFS [y,n,q] (y)
```

The installer stops the Veritas Storage Foundation Cluster File System High Availability processes and uninstalls the packages.

5 Reboot the nodes:

```
# /sbin/shutdown -r now
```

After uninstalling the Veritas Storage Foundation Cluster File System High Availability, refer to the *Veritas Storage Foundation Cluster File System High Availability 6.0 Installation Guide* to reinstall the 6.0 software.

Uninstalling Veritas Storage Foundation for Oracle RAC

The following procedure uninstalls Veritas Storage Foundation for Oracle RAC (SFRAC).

Note: This procedure will remove the complete SFRAC stack from all nodes.

To uninstall Veritas Storage Foundation for Oracle RAC

1 On each node, take the Oracle resources in the VCS configuration file (`main.cf`) offline.

```
# hagrp -offline oracle_group -sys node_name
```

If the database is not managed by VCS, stop the Oracle database as follows:

```
$ srvctl stop database -d db_name
```

2 If Oracle Clusterware is not under VCS Control, then enter the following command on each node of the cluster to stop Oracle Clusterware.

- For 10gR2:

```
# /etc/init.d/init.crs stop
```

- For 11gR2:

```
# /etc/init.d/ohasd stop
```

3 Stop the applications that use CVM or CFS that are not under VCS control

- Using native application commands, stop the applications that use CVM or CFS on all nodes.
- Verify that no processes use the CFS mount point:

```
# fuser -c mount_point
```

4 Unmount CFS file systems that are not under VCS control

- Determine the file systems that need to be unmounted by checking the output of mount command.

```
# mount -v | grep vxfs | grep cluster
```

- Unmount each file system that is not controlled by VCS on each node:

  ```
  # umount mount_point
  ```

5 Stop VCS to take the service groups on all nodes offline

  On any one node execute following command to stop VCS:

  ```
  # hastop -all
  ```

6 Stopping the applications that use VxVM or VxFS that are not under VCS control

- Using native application commands, stop the applications that use VxVM or VxFS.
- Verify that no processes use the VxFS mount point:

  ```
  # fuser -c mount_point
  ```

7 Unmounting VxFS file systems that are not under VCS control.

- Determine the file systems that need to be unmounted by checking the output of mount command.

  ```
  # mount -v | grep vxfs
  ```

  Unmount each file system that is not controlled by VCS on each node:

  ```
  # umount mount_point
  ```

8 Remove SF for Oracle RAC.

- On any one node, navigate to the directory that contains the uninstallsfrac program:

  ```
  # cd /opt/VRTS/install
  ```

- Start the uninstallsfrac program:

  ```
  # ./uninstallsfrac
  ```

9 After uninstalling the SFRAC, refer to the *Veritas Storage Foundation for Oracle RAC 6.0 Installation and Configuration Guide* document to reinstall the SFRAC 6.0 software.