

Veritas Storage Foundation™ and High Availability Solutions 6.0.5 Release Notes - HP-UX

6.0.5 Maintenance Release

Veritas Storage Foundation and High Availability Release Notes

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

Product version: 6.0.5

Document version: 6.0.5 Rev 0

Legal Notice

Copyright © 2014 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website 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 [supportolutions@symantec.com](mailto:supportsolutions@symantec.com)

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>

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 website.

<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 Veritas Storage Foundation and High Availability Solutions

This document includes the following topics:

- [Introduction](#)
- [List of products](#)
- [List of patches](#)
- [Changes introduced in 6.0.5](#)
- [System requirements](#)
- [Fixed issues](#)
- [Known Issues](#)
- [Software limitations](#)
- [Documentation errata](#)

Introduction

This document provides information about the products in Veritas Storage Foundation and High Availability Solutions 6.0.5 Maintenance Release (6.0.5 MR).

For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:

<http://www.symantec.com/docs/TECH164885>

The hardware compatibility list contains information about the supported hardware and is updated regularly. For the latest information on supported hardware visit:

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

Before installing or upgrading Veritas Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

For instructions to install or upgrade the product see the *Veritas Storage Foundation and High Availability Solutions 6.0.5 Installation Guide* at available on the Symantec website:

<http://sort.symantec.com/documents>

This Maintenance Release applies to the following releases of Storage Foundation and High Availability products:

- Storage Foundation and High Availability Solutions 6.0.1
- Storage Foundation and High Availability Solutions 6.0.3

This Maintenance Release is available as 6.0.5

List of products

Apply the patches for the following Veritas Storage Foundation and High Availability products:

- Veritas Dynamic Multi-Pathing (DMP)
- Veritas Volume Manager (VxVM)
- Veritas File System (VxFS)
- Veritas Storage Foundation (SF)
- Veritas Cluster Server (VCS)
- Veritas Storage Foundation and High Availability (SFHA)
- Veritas Storage Foundation Cluster File System and High Availability (SFCFSHA)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)

List of patches

This section lists the patches for 6.0.5.

Table 1-1 Patches for HP-UX

Patch ID	Description	Products Affected	Patch Size
PVCO_03982	VRTS 6.0.300.000 VRTSperl Command Patch	DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, VCS, VM	143M
PVCO_04004	VRTS 6.0.500.000 VRTSvxfs Command Patch	FS, SF, SF Oracle RAC, SFCFSHA, SFHA	67M
PVCO_04005	VRTS 6.0.500.000 VRTSvxvm Command Patch	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, VM	443M
PVCO_04007	VRTS 6.0.500.000 VRTSvcs Command Patch	SF Oracle RAC, SFCFSHA, SFHA, VCS	399M
PVCO_04008	VRTS 6.0.500.000 VRTSvcsag Command Patch	SF Oracle RAC, SFCFSHA, SFHA, VCS	10M
PVCO_04009	VRTS 6.0.500.000 VRTSvsea Command Patch	SF Oracle RAC, SFCFSHA, SFHA, VCS	4.2M
PVCO_04020	VRTS 6.0.500.000 VRTSdbac Command Patch	SF Oracle RAC	88K
PVCO_04023	VRTS 6.0.500.000 VRTSdbed Command Patch	SF, SF Oracle RAC, SFCFSHA, SFHA	76M
PVCO_04024	VRTS 6.0.500.000 VRTSsfcp601 Command Patch	DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, VCS, VM	9.9M
PVCO_04032	VRTS 6.0.500.000 VRTSscavf Command Patch	SF Oracle RAC, SFCFSHA	348K
PVCO_04033	VRTS 6.0.500.000 VRTSfsadv Command Patch	FS, SF, SF Oracle RAC, SFCFSHA, SFHA	8.8M
PVKL_04003	VRTS 6.0.500.000 VRTSvxfs Kernel Patch	FS, SF, SF Oracle RAC, SFCFSHA, SFHA	11M

Table 1-1 Patches for HP-UX (*continued*)

Patch ID	Description	Products Affected	Patch Size
PVKL_04006	VRTS 6.0.500.000 VRTSvxvm Kernel Patch	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, VM	7.3M
PVKL_04031	VRTS 6.0.500.000 VRTSodm Kernel Patch	SF, SF Oracle RAC, SFCFSHA, SFHA	7.6M
PVNE_04016	VRTS 6.0.500.000 VRTSvxfen Kernel Patch	SF Oracle RAC, SFCFSHA, SFHA, VCS	5.9M
PVNE_04017	VRTS 6.0.500.000 VRTSamf Kernel Patch	SF Oracle RAC, SFCFSHA, SFHA, VCS	7.6M
PVNE_04021	VRTS 6.0.500.000 VRTSgab Kernel Patch	SF Oracle RAC, SFCFSHA, SFHA, VCS	4.8M
PVNE_04022	VRTS 6.0.500.000 VRTSllt Kernel Patch	SF Oracle RAC, SFCFSHA, SFHA, VCS	4.8M
VRTSaslapm	Array Support Libraries and Array Policy Modules for Veritas Volume Manager		4.8M

Note: You can also view the list using the `installmr` command: `./installmr -listpatches`

Changes introduced in 6.0.5

This section lists the changes in 6.0.5.

Changes in documentation in 6.0.5

The following are the changes related to documentation introduced in this release:

SFHA Release Notes content now split into separate installation and release notes documents

Maintenance releases until version 6.0.5 included both release-specific and installation content in a single release notes document. Starting with 6.0.5, future maintenance releases will deliver the following documentation with the release:

Document	Description
Veritas Storage Foundation and High Availability Solutions Release Notes	This document will contain release-specific information such as system requirements, changes in the release, fixed issues in the release, known issues and limitations in the release.
Veritas Storage Foundation and High Availability Solutions Installation Guide	This document will contain instructions specific to installing, upgrading, or uninstalling the product.

Both documents will be available on the Symantec SORT web site at the time of release:

<https://sort.symantec.com/welcome/documentation>

Changes related to Storage Foundation and High Availability

This release supports HP Integrity Virtual Machines (IVM) 6.3.

Changes related to installing, upgrading and rolling back

The following changes are related to installing, upgrading and rolling back of the product in 6.0.5 release.

Using Install Bundles with `-base_path` option to install or upgrade to 6.0.5 in one execution.

In version 6.0.5, Symantec offers you a method to easily install or upgrade your systems directly to 6.0.5 in one step using Install Bundles with `-base_path` option. With this option, the installers have the ability to merge base releases like 6.0.1 with 6.0.5 which is a maintenance release, so that you can install or upgrade directly to 6.0.5 in one execution. You do not have to perform two install actions to install or upgrade systems to 6.0.5.

You can get base release from FileConnect that requires customer serial number. For 6.0.5, the base release version should be 6.0.1.

You can also download 6.0.5 from the SORT website.

When you want to install or upgrade to 6.0.5 using Install Bundles with `-base_path` option, the command must be executed from the 6.0.5 install script.

For example, enter the following command:

```
./installmr -base_path <path_to_base>
```

Enhancement on `VRTSaslapm` depot upgrade and rolling back

During the upgrade, if the version of `VRTSaslapm` depot is earlier than 6.0.5, then the installer upgrades the `VRTSaslapm` depot with the one in 6.0.5 release.

During the rolling back, `VRTSaslapm` depot does not revert to the earlier version that you roll back to. If you need to use `VRTSaslapm` depot of an earlier version, then uninstall the current `VRTSaslapm` depot and reinstall `VRTSaslapm` depot of a specific version after rolling back.

Local installer scripts' version suffix changed

The local installer scripts' name under `/opt/VRTS/install/` is changed from `[un]install<prod>601` to `[un]install<prod>605`. This script name change does not affect any functionality.

Changes related to Veritas Volume Manager

There are no changes related to Veritas Volume Manager in 6.0.5.

Changes related to Veritas File System

There are no changes related to Veritas File System in 6.0.5.

Changes related to Veritas Cluster Server

Veritas Cluster Server includes the following changes in 6.0.5:

New attribute `ClearClone` added to `DiskGroup` and `CVMVolDg` agents to support `-c` option to reset `clone_disk` flag during disk group import

In this release, Symantec has introduced boolean attribute `ClearClone` to `DiskGroup` and `CVMVolDg` agents. The default value of the `ClearClone` attribute is 0. If the value of `ClearClone` attribute is set to 1, then the disk group is imported with the

`-c` option. While importing the disk group, this option clears the clone and the `udid_mismatch` flags from the disks of the disk groups and also updates the UDID.

You can modify the `ClearClone` attribute using the following procedure.

To enable the `ClearClone` attribute

- 1 Enable the write access for VCS configuration.

```
#haconf -makerw
```

- 2 Set the value of `ClearClone` attribute to 1.

```
#hares -modify < resource_name > ClearClone 1
```

- 3 Save the VCS configuration.

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

To disable the `ClearClone` attribute

- 1 Enable the write access for VCS configuration.

```
#haconf -makerw
```

- 2 Set the value of `ClearClone` attribute to 0.

```
#hares -modify < resource_name > ClearClone 0
```

- 3 Save the VCS configuration.

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

New command for `hacli` in `vxfsnwap` utility

A new option `-p` is introduced to specify a protocol value that `vxfsnwap` utility can use to communicate with other nodes in the cluster. The supported values for the protocol can be `ssh`, `rsh`, or `hacli`.

Changes introduced in HPVirtualMachine agent.

In this release, Symantec has introduced a resource level attribute to HPVirtualMachine agent - `EnableForMigration`. When the `EnableForMigration` attribute value is set to 1, the virtual machine (VM) being monitored is enabled for online migration by setting the `not_runnable` (NR) flag to the VM when the resource is brought offline or when the resource is faulted and the clean entry point is called

online for the resource. The default value for `EnableForMigration` is 0. New `attr_changed` entry point is added to the `HPVirtualMachine` agent to handle `EnableForMigration` attribute value change.

Support for Oracle Single Instance 12c

In 6.0.5 release, Veritas Cluster Server supports Oracle Single Instance 12c.

Changes related to Veritas Storage Foundation for Oracle RAC in 6.0.5

Veritas Storage Foundation for Oracle RAC includes the following changes in 6.0.5:

This release introduces script-based installer support for configuring Highly Available IP (HAIP) addresses on SF Oracle RAC nodes running Oracle RAC 11.2.0.2 and later versions.

The Oracle Network Configuration menu now displays the following options:

- **Configure private IP addresses (HAIP Configuration)** - For Oracle RAC 11.2.0.2 and later
- **Configure private IP addresses (PrivNIC Configuration)** - For Oracle RAC prior to 11.2.0.2
- **Configure private IP addresses (MultiPrivNIC Configuration)** - For Oracle RAC prior to 11.2.0.2
- **Exit** - Exit SF Oracle RAC Configuration
- **Back** - Back to the previous menu

Oracle 12c support

In 6.0.5 release, Veritas Storage Foundation for Oracle RAC supports Oracle 12c.

Enabling health check monitoring in VCS agent for Oracle with SFHA 6.0.5

In Veritas Storage Foundation High Availability 6.0.5 release, Symantec has enabled the health check monitoring feature in VCS agent for Oracle. Please refer to the following tech note for more details:

<http://www.symantec.com/docs/TECH214888>

Changes related to Veritas Dynamic Multi-Pathing

There are no changes related to Veritas Dynamic Multi-Pathing in this release.

Changes related to Veritas Storage Foundation for databases (SFDB) tools

Veritas Storage Foundation for databases (SFDB) tools includes the following changes in 6.0.5:

Veritas Storage Foundation for databases (SFDB) tools support Oracle 12c release for Oracle databases.

Note: For Oracle 12c, the SFDB tools do not support the Multitenant database features, including the CDB and PDB databases.

System requirements

This section describes the system requirements for this release.

Supported HP-UX operating systems

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

[Table 1-2](#) shows supported operating system for this release.

Table 1-2 Supported operating systems

Operating system	Operating system version	Architecture
HP-UX 11i Version 3	HP-UX B.11.31.1103	PA-RISC/Itanium
	HP-UX B.11.31.1109	
	HP-UX B.11.31.1203	
	HP-UX B.11.31.1209	
	HP-UX B.11.31.1303	
	HP-UX B.11.31.1403	

For versions less than HP-UX 11.31 1403, ensure that the following drivers and packages use version 11.31.1311 or later.

- CommonIO,r=B.11.31.1311
- FibrChanl-01,r=B.11.31.1311
- FibrChanl-02,r=B.11.31.1311
- FibrChanl-03,r=B.11.31.1311

Note: SF Oracle RAC does not support HP-UX (PA) architectures.

Before installing or upgrading Veritas Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

For the latest information on supported hardware, visit:

<https://sort.symantec.com/documents>

Symantec recommends installing the latest HP-UX patches from HP.

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

Supported database software

For the latest information on supported Oracle database versions, see the following TechNote:

<http://www.symantec.com/docs/DOC5081>

Support for minor database versions is also documented in the afore-mentioned TechNote.

Additionally, see the following Oracle support site for information on patches that may be required by Oracle for each release.

<https://support.oracle.com>

Hardware compatibility list

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

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

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

Veritas Storage Foundation 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.

Installation and upgrades: issues fixed in 6.0.5

This section describes the incidents that are fixed in installation and upgrades in 6.0.5.

Table 1-3 Installation and upgrades fixed issues in 6.0.5

Fixed issues	Description
3131744	installmp CPI should not enable DMP native support ,it should ask customer if they want to enable it.
3243089	During a live upgrade the installation process takes more time than expected.
3295841	CPI patches 6.0.1_P4.pl and 6.0.3_P6.pl fails when ssh banner is enabled
3304955	The installer blocks upgrading the product to 6.0.1 if Veritas Cluster Server (VCS) is not configured.
3432524	For Oracle 12c, the installation of Clusterware's response file fails.
3472265	The installvcs script cannot set the heartbeat NIC name to be "L_101".
3436227	The VRTSicsco, VRTSobc33, and VRTSpx files are not removed after upgrading the product from version 50MP2 or 501RP3 to the version 6.0.1.
3448674	After upgrading from 5.0MP3RP5 to 6.0.5 using the base_path option, NFSRestart and NFS upper or lower resource cannot come online automatically.

Installation and upgrades: issues fixed in 6.0.3

This section describes the installation and upgrade issues fixed in 6.0.3.

Table 1-4 Installation and upgrades 6.0.3 fixed issues

Incident	Description
2967125	Eval injection vulnerability in the Digest module before 1.17 for Perl allows context-dependent attackers to execute arbitrary commands via the new constructor.

Installation and upgrades: issues fixed in 6.0.1

This section describes the incidents that are fixed related to installation and upgrades in this release.

Table 1-5 Fixed issues related to installation and upgrades

Incident	Description
2329580	Unable to stop some SFCFSHA processes.
2873102	Perl module error on completion of SFHA installation
2628469	Base501 packages doesn't get removed on fresh installation with 1109 Fusion.
2627076	Incorrect server names sometimes display if there is a clock synchronization issue.
2626311	Installer checks for VRTSfsadv if you specify -version.
2622987	sfmh discovery issue when you upgrade your Veritas product to 6.0.1
2593148	cssd agent configuration failed with CPI when have two priv IP's in setup.
2532432	CPI fails to remove older VRTSperl package after the upgrade.
2526709	DMP-OSN tunable value not get persistence after upgrade from 5.1SP1 to 6.0.
2088827	During product migration the installer overestimates disk space use.

Veritas Storage Foundation Cluster File System High Availability: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) in 6.0.5.

Table 1-6 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues

Fixed issues	Description
3259634	A Cluster File System (CFS) with blocks larger than 4GB may become corrupt.
3066116	The system panics due to NULL pointer dereference at vx_worklist_process()function.
3462694	The fsdedupadm(1M) command fails with error code 9 when it tries to mount checkpoints on a cluster.
3189562	Oracle daemons get hang with the vx_growfile() kernel function.
3214328	A mismatch is observed between the states for the Global Lock Manager (GLM) grant level and the Global Lock Manager (GLM) data in a Cluster File System (CFS) inode.
2495673	Mismatch of concurrent I/O related data in an inode is observed during communication between the nodes in a cluster.
3092114	The information output displayed by the "df -i" command may be inaccurate for cluster mounted file systems.
3224101	After you enable the optimization for updating the i_size across the cluster nodes lazily, the system panics.
2977035	A debug assert issue was encountered in vx_dircompact() function while running an internal noise test in the Cluster File System (CFS) environment
3312897	System can hang when the Cluster File System (CFS) primary node is disabled.
3226462	On a cluster mounted file-system with unequal CPUs, a node may panic while doing a lookup operation.
3228646	NFSv4 server panics in unlock path.
3369049	File system may hang with partitioned directory enabled (PD).
3274592	Internal noise test on cluster file system is unresponsive while executing the fsadm(1M) command
1949445	System is unresponsive when files were created on large directory.
2972183	The fsppadm(1M) enforce command takes a long time on the secondary nodes compared to the primary nodes.

Table 1-6 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3003679	When running the <code>fsppadm(1M)</code> command and removing a file with the named stream attributes (<code>natr</code>) at the same time, the file system does not respond.
3072036	Read operations from secondary node in CFS can sometimes fail with the <code>ENXIO</code> error code.
3364312	The <code>fsadm(1M)</code> command is unresponsive while processing the <code>VX_FSADM_REORGLK_MSG</code> message.
3359200	Internal test on Veritas File System (VxFS) <code>fsdedup(1M)</code> feature in cluster file system environment results in a hang.
2735912	The performance of tier relocation using the <code>fsppadm(1M)enforce</code> command degrades while migrating a large number of files.
3153919	The <code>fsadm (1M)</code> command may hang when the structural file set re-organization is in progress.
3332902	While shutting down, the system running the <code>fsclustadm(1M)</code> command panics.
3046983	Invalid CFS node number in <code>"__fsppadm_fclextract"</code> , causes the DST policy enforcement failure.
3364287	Debug assert may be hit in the <code>vx_real_unshare()</code> function in the cluster environment.
3364301	Assert failure because of improper handling of inode lock while truncating a reorg inode.
3444775	Internal noise testing on cluster file system results in a kernel panic in function <code>vx_fsadm_query()</code> with an error message.
3192985	Checkpoints quota usage on Cluster File System (CFS) can be negative.
3410837	The error message has an issue when the user uses the <code>cfsumount(1M)</code> to unmount a mount point which has a samba parent resource.
2756779	The code is modified to improve the fix for the read and write performance concerns on Cluster File System (CFS) when it runs applications that rely on the POSIX file-record using the <code>fcntl</code> lock.

Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues

Table 1-7 lists the Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) issues fixed in this release.

Table 1-7 Veritas Storage Foundation Cluster File System High Availability fixed issues

Incident	Description
2715175	<code>cfsumount</code> command runs slowly on large file system and for large file systems, reconfiguration process takes a longer time.
2750860	Performance write issue observed due to CFS (Cluster File System) fragmentation in CFS (Cluster File System) cluster.
2942776	CFS mount fails with the error ENXIO or EIO on volume <code>vset</code> device when the volumes in <code>vset</code> is not ready.
2923105	The upgrade <code>VRTSvxfs5.0MP4HFaf</code> hangs at <code>vxfs</code> (Veritas File System) preinstall scripts.
2923867	Internal test hits an assert <code>f:xted_set_msg_pri1:1</code> .
2841059	The file system gets marked for a full <code>fscck</code> operation and the attribute inode is marked as 'bad ondisk'.
2916691	Customer experiencing hangs when doing dedups.
2977697	<code>vx_idetach</code> generated kernel core dump while running filestore replication.
2906018	<code>vx_ireaderrors</code> after successful log replay and mount of the file system.

Veritas Storage Foundation Cluster File System High Availability: issues fixed in 6.0.1

This section describes the incidents that are fixed in Veritas Storage Foundation Cluster File System High Availability in this release.

Table 1-8 Veritas Storage Foundation Cluster File System High Availability fixed issues

Incident	Description
2867282	An ENOSPC error may return to the cluster file system application.

Table 1-8 Veritas Storage Foundation Cluster File System High Availability fixed issues (*continued*)

Incident	Description
2703747	CFS failover takes up to 20 minutes due to slow log replay.
2684573	The performance of the cfsumount(1M) command for the VRTScavf package is slow when some checkpoints are deleted.

Veritas Volume Manager: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas Volume Manager (VxVM) in 6.0.5.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues

Fixed issues	Description
1942051	IO hangs on a master node after disabling the secondary paths from slave node and rebooting the slave node.
2020017	Cluster node panics when mirrored volumes are configured in the cluster.
2054606	During the DMP driver unload operation the system panics.
2236443	Disk group import failure should be made fencing aware, in place of VxVM vxdmp V-5-0-0 i/o error message.
2308875	vxddladm(1M) list command options (hbas, ports, targets) don't display the correct values for the state attribute.
2398954	The system panics while performing I/O on a VxFS mounted instant snapshot with the Oracle Disk Manager (ODM) SmartSync enabled.
2599887	The DMP device paths that are marked as "Disabled" cannot be excluded from VxVM control.
2643506	vxconfigd dumps core when LUNs from the same enclosure are presented as different types, say A/P and A/P-F.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

Fixed issues	Description
2685230	In a Cluster Volume Replicator (CVR) environment, if the SRL is resized and the logowner is switched to and from the master node to the slave node, then there could be a SRL corruption that leads to the Rlink detach.
2735364	The "clone_disk" disk flag attribute is not cleared when a cloned disk group is removed by the "vxdg destroy dg-name" command.
2746907	The vxconfigd(1M) daemon can hang under the heavy I/O load on the master node during the reconfiguration.
2790864	For OTHER_DISKS enclosure, the vxdmpadm config reset CLLfails while trying to reset IO Policy value.
2804326	In the Veritas Volume Replicator (VVR) environment, secondary logging is seen ineffect even if Storage Replicator Log (SRL) size mismatch is seen across primary and secondary.
2812161	In a Veritas Volume Replicator (VVR) environment, after the Rlink is detached, the vxconfigd(1M) daemon on the secondary host may hang.
2825102	CVM reconfiguration and VxVM transaction code paths can simultaneously access volume device list resulting in data corruption.
2845383	The site gets detached if the plex detach operation is performed with the site-consistency set to off.
2860230	In a Cluster Volume Manager (CVM) environment, the shared disk remains as opaque after execution of vxdiskunsetup(1M) command on a master node.
2861011	The "vxdisk -g <dgname> resize diskname" command fails with an error for the Cross-platform Data Sharing(CDS) formatted disk.
2866299	The NEEDSYNC flag set on volumes in a Replicated Volume Group (RVG) not getting cleared after the vxrecover command is run.
2869514	In the clustered environment with large Logical unit number(LUN) configuration, the node join process takes long time.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
2882312	If an SRL fault occurs in the middle of an I/O load, and you immediately issue a read operation on data written during the SRL fault, the system returns old data.
2882412	The 'vxdisk destroy' command uninitialized a VxVM disk which belongs to a deported disk group.
2893530	With no VVR configuration, when system is rebooted, it panicked.
2898324	UMR errors reported by Purify tool in "vradmind migrate" command.
2907746	File Descriptor leaks are observed with the device-discovery command of VxVM.
2909668	In case of multiple sets of the cloned disks of the same source disk group, the import operation on the second set of the clone disk fails, if the first set of the clone disks were imported with "updateid".
2910367	When SRL on the secondary site disabled, secondary panics.
2916911	The vxconfigd(1M) daemon sends a VOL_DIO_READ request before the device is open. This may result in a scenario where the open operation fails but the disk read or write operations proceeds.
2921816	System panics while starting replication after disabling the DCM volumes.
2925746	In the cluster volume manager (CVM) environment, cluster-wide vxconfigd may hang during CVM reconfiguration.
2932214	"vxdisk resize" operation may cause the disk goes into "online invalid" state.
2933476	The vxdisk(1M) command resize fails with a generic error message. Failure messages need to be more informative.
2933688	When the 'Data corruption protection' check is activated by Dynamic Multi-Pathing (DMP), the device- discovery operation aborts, but the I/O to the affected devices continues, this results in data corruption.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
2938710	The vxassist(1M) command dumps core during the layout operation .
2950624	vradmind fails to operate on the new master when a node leaves the cluster.
2952403	Shared disk group fails to destroy if master has lost storage.
2952553	Refresh of a snapshot should not be allowed from a different source volume without force option.
2954455	During Dynamic Reconfiguration Operations in vxdiskadm, when a pattern is specified to match a range of LUNs for removal, the pattern is matched erroneously.
2957555	The vxconfigd(1M) daemon on the CVM master node hangs in the userland during the vxsnap(1M) restore operation.
2958983	Memory leak is observed during the remirror operations.
2959333	The Cross-platform Data Sharing (CDS) flag is not listed for disabled CDS disk groups.
2959733	Handling the device path reconfiguration in case the device paths are moved across LUNs or enclosures to prevent the vxconfigd(1M) daemon core dump.
2962010	The replication hangs when the Storage Replicator Log (SRL) is resized.
2966990	In a Veritas Volume Replicator (VVR) environment, the I/O hangs at the primary side after multiple cluster reconfigurations are triggered in parallel.
2969335	The node that leaves the cluster node while the instant operation is in progress, hangs in the kernel and cannot join back to the cluster node unless it is rebooted.
2969844	The device discovery failure should not cause the DMP database to be destroyed completely.
2972513	In CVM, PGR keys from shared data disks are not removed after stopping VCS.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
2973522	At cable connect on port1 of dual-port Fibre Channel Host Bus Adapters (FC HBA), paths via port2 are marked as SUSPECT.
2976130	Multithreading of the vxconfigd (1M) daemon for HP-UX 11i v3 causes the DMP database to be deleted as part of the device-discovery commands.
2979824	The vxdiskadm(1M) utility bug results in the exclusion of the unintended paths.
2980955	Disk group (dg) goes into disabled state if vxconfigd(1M) is restarted on new master after master switch.
2986596	The disk groups imported with mix of standard and clone logical unit numbers(LUNs) may lead to data corruption.
2992667	When new disks are added to the SAN framework of the Virtual Intelligent System (VIS) appliance and the Fibre Channel (FC) switcher is changed to the direct connection, the "vxdisk list" command does not show the newly added disks even after the "vxdisk scandisks" command is executed.
2993667	Veritas Volume Manager (VxVM) allows setting the Cross-platform Data Sharing (CDS) attribute for a disk group even when a disk is missing, because it experienced I/O errors.
2996142	Data is corrupted or lost if the mapping from disk access (DA) to Data Module (DM) of a disk is incorrect.
2996443	In a cluster volume replication (CVR) environment, log ownname mismatch configuration error is seen on Slave nodes after it brings down the master node.
2999871	The vxinstall(1M) command gets into a hung state when it is invoked through Secure Shell (SSH) remote execution.
3003991	The vxdbg adddisk command hangs when paths for all the disks in the disk group are disabled.
3006245	While executing a snapshot operation on a volume which has 'snappoints' configured, the system panics in frequently.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3010191	Previously excluded paths are not excluded after upgrade to VxVM 5.1SP1RP3.
3011405	Execution of "vxtune -o export" command fails and displays an error message.
3012929	The vxconfigbackup(1M) command gives errors when disk names are changed.
3015181	I/O hangs on both the nodes of the cluster when the disk array is disabled.
3022689	The vxbrk_rootmir(1M) utility succeeds with the following error message: " ioscan: /dev/rdisk/eva4k6k0_48s2: No such file or directory".
3031796	Snapshot reattach operation fails if any other snapshot of the primary volume is not accessible.
3038684	The restore daemon enables the paths of Business Continuance Volumes-Not Ready (BCV-NR) devices.
3041014	Beautify error messages seen during relayout operation.
3045033	"vxdg init" should not create a disk group on clone disk that was previously part of a disk group.
3047470	The device /dev/vx/esd is not recreated on reboot with the latest major number, if it is already present on the system.
3049633	In Veritas Volume Replicator (VVR) environment, the VxVMconfiguration daemon vxconfigd(1M) hangs on secondary node when all disk paths are disabled on secondary node.
3052770	The vradmin syncrvg operation with a volume set fails to synchronize the secondary RVG with the primary RVG.
3052879	Auto import of the cloned disk group fails after reboot even when source disk group is not present.
3053073	Dynamic Reconfiguration (DR) Tool doesn't pick thin LUNs in "online invalid" state for disk remove operation.
3060327	The vradmin repstatus(1M) shows "dcm contains 0 kbytes" during the Smart Autosync.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3065072	Data loss occurs during the import of a clone disk group, when some of the disks are missing and the import "useclonedev" and "updateid" options are specified.
3067452	If new LUNs are added in the cluster, and its naming scheme has the avid set option set to 'no', then DR (Dynamic Reconfiguration) Tool changes the mapping between dmpnode and disk record.
3067784	The grow and shrink operations by the vxresize(1M) utility may dump core in vfprintf() function.
3074579	The "vxddmpadm config show" CLI does not display the configuration file name which is present under the root(/) directory.
3076093	The patch upgrade script "installrp" can panic the system while doing a patch upgrade.
3084449	The shared flag sets during the import of private disk group because a shared disk group fails to clear due to minor number conflict error during the import abort operation.
3085519	Missing disks are permanently detached from the disk group because -o updateid and tagname options are used to import partial disks.
3088907	A node in a Cluster Volume Manager can panic while destroying a shared disk group.
3091916	The Small Computer System Interface (SCSI) I/O errors overflow the syslog.
3098559	Cluster File System (CFS) data corrupted due to cloned copy of logical unit numbers (LUNs) that is imported with volume asymmetry.
3099796	The vxevac command fails on volumes having Data Change Object (DCO) log. The error message "volume is not using the specified disk name" is displayed.
3101419	In CVR environment, I/Os to the data volumes in an RVG experience may temporary hang during the SRL overflow with the heavy I/O load.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3102114	A system crash during the 'vxsnap restore' operation can cause the vxconfigd(1M) daemon to dump core after the system reboots.
3111062	When diffsync is executed, vxrsync gets the following error in lossy networks: VxVM VVR vxrsync ERROR V-5-52-2074 Error opening socket between[HOST1] and [HOST2] -- [Connection timed out]
3114134	The Smart (sync) Autosync feature fails to work and instead replicates the entire volume size for larger sized volumes.
3120458	In cluster volume replication (CVR) in data change map (DCM) mode, cluster-widevxconfigd hang is seen when one of the nodes is stopped.
3121380	I/O of replicated volume group (RVG) hangs after one data volume is disabled.
3122828	Dynamic Reconfiguration (DR) tool lists the disks which are tagged with Logical Volume Manager (LVM), for removal or replacement.
3125631	Snapshot creation on volume sets may fail with the error: "vxsnap ERRORV-5-1-6433 Component volume has changed".
3127543	Non-labeled disks go into udid_mismatch after vxconfigd restart.
3130353	Continuous disable or enable path messages are seen on the console forEMC Not Ready (NR) devices.
3136272	The disk group import operation with the "-o noreonline" option takes additional import time.
3139300	Memory leaks are observed in the device discovery code path of VxVM.
3142315	Disk is misidentified as clone disk with udid_mismatch flag.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3144781	In the Veritas Volume Replicator (VVR) environment, execution of the vxlinkpause command causes a hang on the secondary node if the rlink disconnect is already in progress.
3146955	Remote disks (failed or missing disks) go into the "ONLINE INVALID LFAILED" or "ONLINE INVALID LMISSING" state after the disk loses global disk connectivity.
3152274	The dd command to SRDF-R2 (write disable) device hangs, which causes the vm command hangs for a long time. But no issues with the Operating System (OS) devices.
3162418	The vxconfigd(1M) command dumps core due to wrong check in ddl_find_cdevno() function.
3162987	The disk has a UDID_MISMATCH flag in the vxdisk list output.
3163549	vxconfigd(1M) hangs on master node if slave node joins the master having disks which are missing on master.
3163970	The "vxsnap -g disk group syncstart volume" command is unresponsive on the Veritas Volume Replicator (VVR) DR site.
3178029	When you synchronize a replicated volume group (RVG), the diff string is over 100%.
3178182	During a master take over task, shared disk group re-import operation fails due to false serial split brain (SSB) detection.
3188154	The vxconfigd(1M) daemon does not come up after enabling the native support and rebooting the host.
3199056	Veritas Volume Replicator (VVR) primary system panics in the vol_cmn_errfunction due to the VVR corrupted queue.
3222707	Dynamic Reconfiguration (DR) tool does not permit the removal of disks associated with a deported diskgroup(dg).

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3225660	The Dynamic Reconfiguration (DR) tool does not list thin provisioned LUNs during a LUN removal operation.
3238397	Dynamic Reconfiguration (DR) Tool's Remove LUNs option does not restart the vxattachd daemon.
3239521	When you do the PowerPath pre-check, the DynamicReconfiguration (DR) tool displays the following error message: 'Unable to runcommand [/sbin/powermt display]' and exits.
3240858	The /etc/vx/vxesd/.udev_lock file may have different permissions at different instances.
3243355	The vxres_lvmroot(1M) utility which restores the Logical Volume Manager (LVM) root disk from the VxVM root disk fails.
3244217	Cannot reset the clone_disk flag during vxdg import.
3247040	vxdisk scandisks enables the PowerPath (PP) enclosure which was disabled previously.
3248281	When the "vxdisk scandisks" or "vxdctl enable" commands are run consecutively the "VxVM vxdisk ERROR V-5-1-0 Device discovery failed." error is encountered.
3254311	System panics when reattaching site to a site-consistent diskgroup having volumelarger than 1.05 TB
3259732	In a CVR environment, rebooting the primary slave followed by connect-disconnectin loop causes rlink to detach.
3261485	The vxcdsconvert(1M) utility failed with the error "Unable to initialize the disk as a CDS disk".
3261601	System panics when dmp_destroy_dmpnode() attempts to free an already free virtual address.
3271595	Veritas Volume Manager (VxVM) should prevent the disk reclaim flag from getting turned off, when there are pending reclaims on the disk.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3271985	In Cluster Volume Replication (CVR), with synchronous replication, aborting a slave node from the Cluster Volume Manager (CVM) cluster makes the slave node panic.
3279932	The vxdisksetup and vxdiskunsetup utilities were failing on disk which is part of a deported disk group (DG), even if "-f" option is specified.
3280830	Multiple vxresize operations on a layered volume fail with error message "There are other recovery activities. Cannot grow volume"
3287880	In a clustered environment, if a node doesn't have storage connectivity to clone disks, then the vxconfigd on the node may dump core during the clone disk group import.
3289202	Handle KMSG_EPURGE error in CVM disk connectivity protocols.
3323548	In the Cluster Volume Replicator (CVR) environment, a cluster-wide vxconfigd hang occurs on primary when you start the cache object.
3368361	When site consistency is configured within a private disk group and CVM is up, the reattach operation of a detached site fails.
3373142	Updates to vxassist and vxedit man pages for behavioral changes after 6.0.
3385753	Replication to the Disaster Recovery (DR) site hangs eventhough Replication links (Rlinks) are in the connected state.
3399131	For Point Patch (PP) enclosure, both DA_TPD and DA_COEXIST_TPD flags are set.
3400504	Upon disabling the host side Host Bus Adapter (HBA) port, extended attributes of some devices are not seen anymore.
3408320	Thin reclamation fails for EMC 5875 arrays.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3409016	During cold install and recovery of archives with HPUX ignite-UX, the error message "VxVM vxvol WARNING V-5-1-16864 Unable to open reclaim_log file" is displayed.
3409612	The value of reclaim_on_delete_start_time cannot be set to values outside the range: 22:00-03:59
3415188	I/O hangs during replication in Veritas Volume Replicator (VVR).
3416622	The hot-relocation feature fails for a corrupted disk in the CVM environment.
3417044	System becomes unresponsive while creating a VVR TCP connection.
3424798	Veritas Volume Manager (VxVM) mirror attach operations(e.g., plex attach, vxassist mirror, and third-mirror break-off snapshot resynchronization) may take longer time under heavy application I/O load.
3435225	In a given CVR setup, rebooting the master node causes one of the slaves to panic.
3438271	The vxconfigd may hang while adding new LUNs using DR (Dynamic Reconfiguration) tool.
3461383	The vxlink(1M) command fails when the "vxlink -g <DGNAME> -a att <RLINK>" command is executed.
3250450	In the presence of a linked volume, running the vxdisk(1M) command with the -o thin, fssize list option causes the system to panic.
3250369	Execution of vxdisk scandisks command causes endless I/O error messages in syslog.
3249264	Veritas Volume Manager (VxVM) thin disk reclamation functionality causes disk label loss, private region corruption and data corruption.
3237503	System hangs after creating space-optimized snapshot with large size cache volume.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3236773	Multiple error messages of the same format are displayed during setting or getting the failover mode for EMC Asymmetric Logical Unit Access (ALUA disk array).
3235350	I/O on grown region of a volume leads to system panic if the volume has instant snapshot.
3230148	Clustered Volume Manager (CVM) hangs during split brain testing.
3218013	Dynamic Reconfiguration (DR) Tool does not delete the stale OS (Operating System device handles).
3199398	Output of the command "vxmpadm pgrreg" depends on the order of DMP node list where the terminal output depends on the last LUN (DMP node).
3194358	The continuous messages displayed in the syslog file with EMC not-ready (NR) LUNs.
3194305	In the Veritas Volume Replicator (VVR) environment, replication status goes in a paused state.
3182350	VxVM volume creation or size increase hangs.
3182175	The vxdisk -o thin,fssize list command can report incorrect File System usage data.
3158320	VxVM (Veritas Volume Manager) command "vxdisk -px REPLICATED list (disk)" displays wrong output.
3156295	When DMP native support is enabled for Oracle Automatic Storage Management (ASM devices, the permission and ownership of /dev/raw/raw# devices goes wrong after reboot.
3146715	Rlinks' do not connect with NAT configurations on Little Endian Architecture.
3139983	Failed I/Os from SCSI are retried only on very few paths to a LUN instead of utilizing all the available paths, and may result in DMP sending I/O failures to the application bounded by the recovery option tunable.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3125711	When the secondary node is restarted and the reclaim operation is going on the primary node, the system panics.
3119102	Support LDOM Live Migration with fencing enabled.
3107741	The vxrvg snapdestroy command fails with the "Transaction aborted waiting for io drain" error message.
3090667	The system panics or hangs while executing the "vxdisk -o thin,fssize list" command as part of Veritas Operations Manager (VOM) Storage Foundation (SF) discovery.
3086627	The "vxdisk -o thin,fssize list" command fails with error message V-5-1-16282.
3081410	Dynamic Reconfiguration (DR) tool fails to pick up any disk for LUNs removal operation.
3063378	VM commands are slow when Read Only disks are presented.
3056311	For release < 5.1 SP1, allow disk initialization with CDS format using raw geometry.
3021970	A secondary node panics due to NULL pointer dereference when the system frees an interlock.
3019684	I/O hang is observed when SRL is about to overflow after the logowner switches from slave to master.
2994976	System panics during mirror break-off snapshot creation or plex detach operation in vol_mv_pldet_callback() function.
2959325	The vxconfigd(1M) daemon dumps core while performing the disk group move operation.
2957645	When the vxconfigd daemon/command is restarted, the terminal gets flooded with error messages.
2857044	System crash on voldco_getalloffset when trying to resize filesystem.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
2824977	The Command Line Interface (CLI) "vxddmpadm setattr enclosure <enclname> failovermode" which is meant for Asymmetric Logical Unit Access ALUA type of arrays fails with an error on certain arrays without providing an appropriate reason for the failure.
2665425	The vxdisk -px "attribute" list(1M) Command Line Interface (CLI) does not support some basic VxVM attributes.
2152830	A diskgroup (DG) import fails with a non-descriptive error message when multiple copies (clones) of the same device exist and the original devices are either offline or not available.
2091520	The ability to move the configdb placement from one disk to another using "vxdisk set <disk> keepmeta=[always skip default]" command.
1783763	In a Veritas Volume Replicator (VVR) environment, the vxconfigd(1M) daemon may hang during a configuration change operation.
3377383	The vxconfigd crashes when a disk under Dynamic Multi-pathing (DMP) reports device failure.
3325371	Panic occurs in the vol_multistepsio_read_source() function when snapshots are used.
3373208	DMP wrongly sends the SCSI PR OUT command with APTPL bit value as A0A to arrays.
3327842	In the Cluster Volume Replication (CVR) environment, with IO load on Primary and replication going on, if the user runs the vradm resizevol(1M) command on Primary, often these operations terminate with error message "vradm ERROR Lost connection to host".
3301470	All cluster volume replication (CVR) nodes panic repeatedly due to null pointer dereference in vxio.
3300418	VxVM volume operations on shared volumes cause unnecessary read I/Os.
3283525	The vxconfigd(1M) daemon hangs due to Data Change Object (DCO) corruption after volume resize.

Table 1-9 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3325122	In a Clustered Volume Replicator (CVR) environment, when you create stripe-mirror volumes with logtype=dcm, creation may fail.
3312162	Data corruption may occur on the Secondary Symantec Volume Replicator (VVR) Disaster Recovery (DR) Site.
3326964	VxVM hangs in Clustered Volume Manager (CVM) environments in the presence of FMR operations.
3332796	Getting message: VxVM vxiasm INFO V-5-1-0 seeking block #... while initializing disk that is not ASM disk.
3353211	<p>A. After EMC Symmetrix BCV (Business Continuance Volume) device switches to read-write mode, continuous vxdump (Veritas Dynamic Multi Pathing) error messages flood syslog.</p> <p>B. DMP metanode/path under DMP metanode gets disabled unexpectedly.</p>
3218013	Dynamic Reconfiguration (DR) Tool does not delete the stale OS (Operating System) device handles.
3423316	The vxconfigd(1M) daemon observes a core dump while executing the vxdisk(1M) scandisks command.
3279039	HP Ignite integration for VxVM 6.0.1
3326964	VxVM hangs in Clustered Volume Manager (CVM) environments in the presence of FMR operations.
1783763	In a Veritas Volume Replicator (VVR) environment, the vxconfigd(1M) daemon may hang during a configuration change operation.
2994976	System panics during mirror break-off snapshot creation or plex detach operation in vol_mv_pldet_callback() function.
2857044	System crashes while resizing a volume with Data Change Object (DCO) version 30.

Veritas Volume Manager: Issues fixed in 6.0.3

Table 1-10 lists the Veritas Volume Manager issues fixed in this release.

Table 1-10 Veritas Volume Manager fixed issues

Incident	Description
2858853	After master switch, <code>vxconfigd</code> dumps core on old master.
2779580	Secondary node gives configuration error 'no Primary RVG' after reboot of master node on primary site.
2866059	Improve error messages that are hit during <code>vxdisk</code> resize operation.
2149922	Record the diskgroup import and deport events in syslog.
2851403	System panic is seen while unloading "vxio" module. This happens whenever VxVM uses SmartMove feature and the "vxportal" module gets reloaded.
2930569	The LUNs in 'error' state in output of ' <code>vxdisk list</code> ' cannot be removed through DR (Dynamic Reconfiguration) Tool.
2916094	Enhancements have been made to the Dynamic Reconfiguration Tool (DR Tool) to create a separate log file every time DR Tool is started, display a message if a command takes longer time, and not to list the devices controlled by TPD (Third Party Driver) in 'Remove Luns' option of DR Tool.
1859018	Link <link-name> link detached from volume <volume-name> warnings are displayed when a linked-breakoff snapshot is created.
2715129	<code>Vxconfigd</code> hangs during Master takeover in a CVM (Clustered Volume Manager) environment.
2753954	When a cable is disconnected from one port of a dual-port FC HBA, the paths via another port are marked as SUSPECT PATH.
2886402	When re-configuring devices, <code>vxconfigd</code> hang is observed.
2878876	<code>vxconfigd</code> dumps core in <code>vol_cbr_dolog()</code> due to race between two threads processing requests from the same client.
1973983	<code>vxunreloc</code> fails when DCO (Data Change Object) plex is in disabled state.
2801962	Growing a volume takes significantly large time when the volume has version 20 DCO (Data Change Object) attached to it.
2886333	<code>vxdg(1M) join</code> command allowed mixing clone and non-clone disk group. Subsequent import of new joined disk group fails.
2000585	<code>vxrecover</code> does not start remaining volumes, if, one of the volumes is removed during <code>vxrecover</code> command run.

Table 1-10 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2834046	NFS migration failed due to device reminoring.
2567618	VRTSexplorer core dumps in <code>checkhbaapi/print_target_map_entry</code> .
1765916	VxVM socket files do not have proper write protection.
1903700	Removing mirror using <code>vxassist</code> does not work.
2919627	Dynamic Reconfiguration tool should be enhanced to remove LUNs feasibly in bulk.
2892983	<code>vxvol</code> dumps core if new links are added while the operation is in progress.
1982965	<code>vx dg import DGNAME <da-name..></code> fails when "da-name" used as an input to <code>vx dg</code> command is based on naming scheme which is different from the prevailing naming scheme on the host.
2899173	<code>vxconfigd</code> hangs after executing <code>vradmin stoprep</code> command.
2910043	Frequent swapin/swapout of pages seen due to higher order memory requests.
2876256	<code>vx disk -f -g <dg1> set <da_name> mediatype=ssd</code> command fails with new naming scheme.
2859470	EMC SRDF (Symmetrix Remote Data Facility) R2 disk with EFI label is not recognized by VxVM (Veritas Volume Manager) and its shown in error state.
2836798	In VxVM, resizing simple EFI disk fails and causes system panic/hang.
2919720	<code>vxconfigd</code> dumps core in <code>rec_lock1_5()</code> function.
2940446	I/O can hang on volume with space optimized snapshot if the underlying cache object is of very large size. It can also lead to data corruption in cache-object.
1725593	The <code>vx dmpadm listctlr</code> command has to be enhanced to print the count of device paths seen through the controller.
2970368	Enhancing handling of SRDF-R2 WD devices in DMP.
2510928	The extended attributes reported by <code>vx disk -e list</code> for the EMC SRDF LUNs are reported as <code>tdev mirror</code> , instead of <code>tdev srdf-r1</code> .

Table 1-10 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2942609	Message displayed when user quits from Dynamic Reconfiguration Operations is shown as error message.
2911040	Restore from a cascaded snapshot leaves the volume in unusable state if any cascaded snapshot is in detached state.
2744004	<code>vxconfigd</code> hangs on the VVR (Veritas Volume Replicator) secondary node during VVR (Veritas Volume Replicator) configuration.
2919318	During CVM (Cluster Volume Manager) node join and shared disk group operation, the I/O fencing key value of data disk are different and abnormal in a VCS cluster with I/O fencing.
2833498	<code>vxconfigd</code> daemon hangs in <code>vol_ktrans_commit()</code> while reclaim operation is in progress on volumes having instant snapshots.
2857827	During early boot, recovery of linked volume resize fails due to <code>/usr</code> not mounted.
2815517	<code>vx dg adddisk</code> allows mixing of clone and non-clone disks in a DiskGroup.
2915063	System panics during detaching plex of volume in CVM (Cluster Volume Manager) environment.
2837717	<code>vx disk(1M) resize</code> command fails if 'da name' is specified.
2826125	VxVM (Veritas Volume Manager) script daemon is terminated abnormally on its invocation.
2919714	On a THIN LUN, <code>vx evac</code> returns 0 without migrating unmounted VxFS (Veritas FileSystem) volumes.
2692012	When moving subdisks, using <code>vx evac</code> command fails with a generic message which does not convey exactly why the operation failed.
2851085	DMP does not detect implicit LUN ownership changes for some of the DMP nodes.
2898547	<code>vradmin</code> dumps core on Veritas Volume Replicator secondary Site , when Logowner Service Group on VVR (Veritas Volume Replicator) Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes.
2798673	System panics in <code>voldco_alloc_layout()</code> function while creating volume with instant DCO (Dynamic Change Object).

Table 1-10 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2933138	System panics in <code>voldco_update_itemq_chunk()</code> function due to accessing invalid buffer.
2619600	Live migration of virtual machine having SFHA/SFCFSHA stack with data disks fencing enabled, causes service groups configured on virtual machine to fault.
2149922	Record the diskgroup import and deport events in <code>syslog</code> .
1901838	Incorrect setting of "No license" flag leads to DMP (Dynamic Multi-Pathing) database inconsistency.
2273190	Incorrect setting of UNDISCOVERED flag can lead to database inconsistency.
2898547	<code>vradmind</code> on VVR (Veritas Volume Replicator) Secondary Site dumps core, when Logowner Service Group on VVR (Veritas Volume Replicator) Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes.
2962262	Uninstallation of DMP (Dynamic Multi - Pathing) fails in presence of other multipathing solutions.
2851085	DMP (Dynamic Multi - Pathing) does not detect implicit LUN ownership changes for some of the DMP nodes.
2948172	Execution of <code>vxdisk -o thin, fssize list</code> command can result in panic.
2935771	In VVR (Veritas Volume Replicator) environment, RLINK diconnects after switching a master on the primary.
3002770	Accessing NULL pointer in <code>dmp_aa_recv_inquiry()</code> causes system panic.
2869594	Master node panics due to corruption, if, space optimized snapshots are refreshed and master node is selected using <code>vxclustadm setmaster</code> .
2965910	When <code>-o ordered</code> is used, <code>vxassist</code> handles non-disk parameters in a differenn way. This scenario may result in invalid comparison, leading to a core dump.
2398416	<code>vxassist</code> dumps core while creating volume when attribute <code>wantmirror=ctlr</code> is added to the <code>/etc/default/vxassist</code> file.

Table 1-10 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2851403	System panic is seen while unloading <code>vxio</code> module. This happens whenever VxVM uses SmartMove feature and the <code>vxportal</code> module gets reloaded (For example, during VxFS package upgrade).

Veritas Volume Manager: issues fixed in 6.0.1

This section describes the incidents that are fixed in Veritas Volume Manager in this release. This list includes Veritas Volume Replicator fixed issues.

This section describes the incidents that are fixed in Veritas Volume Manager in this release. This list includes Veritas Volume Replicator and Cluster Volume Manager fixed issues.

Table 1-11 Veritas Volume Manager fixed issues

Incident	Description
2838059	VVR Secondary panic in <code>vol_rv_update_expected_pos</code> .
2832784	ESX panicked after applying a template file from GUI.
2826958	The <code>pwn</code> number is not displayed in the output of command <code>vxdmadm list dmpnode dmpnodename=dmpnode name</code> .
2818840	Enhance the <code>vxdmpraw</code> utility to support permission and "root:non-system" ownership to be set and make it persistent.
2794625	Unable to configure ASM to use DMP native block device path.
2792748	CVM join fails because of wrong <code>fd</code> close.
2792242	I/O hang after performing zone remove/add operations.
2774406	The <code>svol_flush_srl_to_dv_start</code> fails to start.
2771452	IO hung because of hung port deletion.
2763206	The <code>vxdisk rm</code> command core dumps when list of disknames is very long.
2756059	Panic in <code>voldco_or_dr1_to_pvm</code> when volume started at boot.
2754819	Live deadlock seen during disk group rebuild when the disk group contains cache object.
2751278	The <code>vxconfigd</code> daemon hung on all cluster nodes during <code>vxsnap</code> operation.

Table 1-11 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2743926	DMP <code>restored</code> daemon fails to restart during system boot.
2741240	The <code>vx dg join</code> transaction failed and did not rollback to the <code>sourcedg</code> .
2739709	Disk group rebuild related issues.
2739601	VVR: <code>repstatus</code> output occasionally reports abnormal timestamp.
2737420	The <code>vxconfigd</code> daemon dumps core while onlining of the disk.
2729501	Exclude path not working properly and can cause system hang while coming up after enabling native support.
2710579	Do not write backup labels for CDS disk - irrespective of disk size.
2710147	Node panics in <code>dmp_pr_do_reg</code> during key registration with fencing enabled.
2703858	Site failure (storage and all nodes including master node) led to 'configuration daemon not accessible' error on all the sites.
2700792	SEGV in <code>vxconfigd</code> daemon during CVM startup.
2700486	The <code>vradmin</code> daemon core dumps when Primary and Secondary have the same hostname and an active Stats session exists on Primary.
2700086	EMC BCV (NR) established devices are resulting in multiple DMP events messages (paths being disabled/enabled).
2698860	The <code>vxassist mirror</code> command failed for thin LUN because <code>statvfs</code> failed.
2689845	After upgrade, some VxVM disks changed to error status and the disk group import failed.
2688747	Logowner local sequential I/Os starved with heavy I/O load on logclient.
2688308	Do not disable other disk groups when a re-import of a disk group fails during master take-over.
2680482	Empty <code>vx.*</code> directories are left in the <code>/tmp</code> directory.
2680343	Node panic during <code>cur pri</code> path update in cluster while running I/O shipping.
2679917	Corrupt space optimized snapshot after a refresh with CVM master switching.
2677016	The <code>vxesd</code> daemon dumps core with SIGILL.
2675538	The <code>vx disk resize</code> command may cause data corruption.

Table 1-11 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2664825	Disk group import fails when disk contains no valid UDID tag on config copy and config copy is disabled.
2656803	Race between <code>vxnetd start</code> and <code>stop</code> operations causes panic.
2652485	Inactive snapshot LUNs cause trespassing.
2648176	Performance difference on Master versus Slave during recovery with Data Change Object (DCO).
2645196	Campus Cluster + Hot Relocation: When a disk failure is detected, the associated disks for that site are detached and ALL disks as marked as RLOC.
2643634	Message enhancement for a mixed (non-cloned and cloned) disk group import.
2627126	Lots of I/Os and paths are stuck in <code>dmp_delayq</code> and <code>dmp_path_delayq</code> respectively. DMP daemon did not wake up to process them.
2626199	The <code>vxdumpadm list dmpnode</code> printing incorrect path type.
2620555	I/O hang due to SRL overflow & CVM reconfig.
2617277	Need man pages for the <code>vxautoanalysis</code> and <code>vxautoconvert</code> commands.
2580393	Removal of SAN storage cable on any node brings Oracle Application Groups down on all nodes.
2566174	Null pointer dereference in <code>volcvm_msg_rel_gslock()</code> .
2564092	Automate the LUN provisioning (addition) / removal steps using <code>vxdiskadm</code> .
2553729	Status of the EMC Clariion disk changed to "online clone_disk" after upgrade.
2495346	The <code>vxvmconvert</code> utility is broken to convert LVM to VxVM:hpdisk for larger configurations.
2495338	Disks with <code>hpdisk</code> format can't be initialized with private region offset other than 128.
2441283	The <code>vxsnap addmir</code> command sometimes fails under heavy I/O load.
2427894	Opaque disk support for VIS appliance.
2249445	Develop a tool to get the disk-related attributes like geometry, label, media capacity, partition info etc.
2240056	The <code>vxdg move</code> transaction not completing and backups fail.

Table 1-11 Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2227678	The second rlink gets detached and does not connect back when overflowed in a multiple-secondaries environment.
1675482	The <code>vx dg list dgname</code> command gives error 'state=new failed'.
1190117	<code>vx disk -f init</code> can overwrite some of the public region contents.

Dynamic Multi-Pathing: issues fixed in 6.0.1

This section describes the incidents that are fixed for Dynamic Multi-Pathing in this release.

Table 1-12 Veritas Dynamic Multi-Pathing fixed issues

Incident	Description
2826958	<code>pwwn no</code> is not displayed in the output of command " <code>vx dmpadm list dmpnode dmpnodename=</code> ".
2818840	Enhance the <code>vx dmp raw</code> utility to support permission and <code>root:non-system</code> ownership be set and make it persistent.
2792242	I/O hang after performing zone remove/add operations.
2743926	DMP restored fails to restart during system boot in 6.0.
2729501	exclude path not working properly and can cause system hang while coming up after enabling native support.
2700086	EMC BCV (NR) established devices are resulting in multiple dmp events messages (paths being disabled/enabled).
2677016	<code>vx esd</code> dumps core with SIGILL.
2652485	Inactive snapshot luns cause trespassing.
2626199	<code>vx dmpadm list dmpnode</code> printing incorrect path-type.
2564092	[VxVM][Usability]Automate the lun provisioning (addition) / removal steps using <code>vx diskadm /or new VxVM CLI</code> command.
2556467	DMP-ASM: disable all paths and reboot host cause <code>/etc/vx/.vx dmp raw dev</code> records losing.

Veritas File System: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas File System (VxFS) in 6.0.5

Table 1-13 Veritas File System 6.0.5 fixed issues

Fixed issues	Description
2059611	The system panics due to a NULL pointer dereference while flushing bitmaps to the disk.
2439261	When the vx_fiostats_tunable value is changed from zero to non-zero, the system panics.
2444146	The Oracle Disk Manager read returns EINTR while running unspecified Oracle jobs.
2646933	VxFS takes long time to process the large sequential writes.
2667658	The 'fscdsconv endian' conversion operation fails because of a macro overflow.
2834192	You are unable to mount the file system after the full fsck(1M) utility is run.
2839871	On a system with DELICACHE enabled, several file system operations may hang.
2850738	The system may hang in the low memory condition.
2908391	It takes a long time to remove checkpoints from the VxFS file system, when there are a large number of files present.
2926684	In rare cases, the system may panic while performing a logged write.
2963763	When the thin_friendly_alloc() and deliache_enable() functionality is enabled, VxFS may enter a deadlock.
2964018	On a high end machine with about 125 CPUs, the operations using the lstat64(2) function, may hang.
2966277	Systems with high file system activity like read/write/open/lookup may panic the system.
2970219	Panic in fcache_as_map+0x70 due to null v_vmdata.
2977828	The file system is marked bad after an inode table overflow error.
2982157	During internal testing, the f:vx_trancommit:4 debug asset was hit when the available transaction space is lesser than required.

Table 1-13 Veritas File System 6.0.5 fixed issues (*continued*)

Fixed issues	Description
2999493	The file system check validation fails with an error message after a successful full fsck operation during internal testing.
2999560	The 'fsvoladm'(1M) command fails to clear the 'metadataok' flag on a volume.
3010444	On a NFS filesystem cksum(1m) fails with the "cksum: read error on : Bad address" error.
3031869	"vxfsstat -b" does not print correct information on maximumbuffer size.
3031901	The 'vxtunefs(1M)' command accepts the garbage value for the 'max_buf_dat_size' tunable.
3042485	During internal stress testing, the f:vx_purge_nattr:1 assert fails.
3049408	When the system is under the file-cache pressure, the find(1) command takes time to operate.
3079215	Oracle RAC Database creation failed with the Ora-00600[ksfd_odmio1] error when Veritas ODM links.
3096834	Intermittent vx_disable messages are displayed in the system log.
3099638	The vxfs_ifree_timelag(5) tunable when tuned, displays incorrect minimum value.
3101418	The current time returned by the operating system (Oracle error code ORA-01513)during Oracle startup is invalid.
3121933	The pwrite(2) function fails with the EOPNOTSUPP error.
3149174	Veritas Oracle Disk Manager (ODM) clone shutdown fails with the ORA-03113: end-of-file on communication channel error.
3150368	vx_writesuper() function causes the system to panic in evfsevol_strategy().
3157624	The fcntl() system call when used for file share reservations(F_SHARE command) can cause a memory leak in Cluster File System (CFS).
3164418	Internal stress test on locally mounted VxFS filesytem results in data corruption in no space on device scenario while doing spilt on Zero Fill-On-Demand (ZFOD) extent
3194635	The internal stress test on a locally mounted file system exited with an error message.

Table 1-13 Veritas File System 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3197901	Prevent duplicate symbol in VxFS libvxfpriv.a and vxfspriv.so
3214816	With the DELICACHE feature enabled, frequent creation and deletion of the inodes of a user may result in corruption of the user quota file.
3233284	FSCK binary hangs while checking Reference Count Table (RCT).
3252983	On a high-end system greater than or equal to 48 CPUs, some file system operations may hang.
3253210	File system hangs when it reaches the space limitation.
3261462	File system with size greater than 16TB corrupts with vx_mapbad messages in the system log.
3291635	Internal testing found debug assert vx_freeze_block_threads_all:7c on locally mounted file systems while processing preambles for transactions.
3297840	A metadata corruption is found during the file removal process.
3298041	With the delayed allocation feature enabled on a locally mounted file system, observable performance degradation might be experienced when writing to a file and extending the file size.
3308673	A fragmented file system is disabled when delayed allocations feature is enabled.
3310755	Internal testing hits a debug assert vx_rcq_badrecord:9:corruptfs.
3312030	The default quota support on Veritas File System (VxFS) version 6.0.4 and later is changed to 32 bit.
3331045	Kernel Oops in unlock code of map while referring freedmlink due to a race with iodone routine for delayed writes.
3331093	Issue with MountAgent Process for vxfs. While doing repeated switchover on HP-UX, MountAgent got stuck.
3331109	The full fsck does not repair the corrupted reference count queue (RCQ) record.
3331419	System panic because of kernel stack overflow.
3335272	The mkfs (make file system) command dumps core when the log size provided is not aligned.

Table 1-13 Veritas File System 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3349634	Assert failure if tried to write on file snapped allocated HOLE.
3364282	The fsck(1M) command fails to correct inode list file
3364290	The kernel may panic in Veritas File System (VxFS) when it is internally working on reference count queue (RCQ) record.
3364306	Stack overflow seen in extent allocation code path.
3394803	A panic is observed in VxFS routine vx_upgrade7() function while running the vxupgrade command(1M).
3436699	An assert failure occurs because of a race condition between clone mount thread and directory removal thread while pushing data on clone.

Veritas File System: Issues fixed in 6.0.3

Table 1-14 lists the Veritas File System issues fixed in this release.

Table 1-14 Veritas File System fixed issues

Incident	Description
2895743	Accessing named attributes for some files seems to be slow.
2881211	File ACLs not preserved in checkpoints properly if file has hardlink.
2756779	Write and read performance concerns on CFS when running applications that rely on posix file-record locking (fcntl).
2858683	Reserve extent attributes changed after vxrestore, only for files greater than 8192 bytes.
2806466	A reclaim operation on a file system mounted on a Logical Volume Manager (LVM) volume using the fsadm(1M) command with the 'R' option may panic the system.
2624262	Panic hit in vx_bc_do_brelse() function while executing dedup functionality.
2616622	The performance of the mmap() function is slow when the file system block size is 8 KB and the page size is 4 KB.
2857751	The internal testing hits the assert f:vx_cbdnlc_enter:1a.

Table 1-14 Veritas File System fixed issues (*continued*)

Incident	Description
2730759	The sequential read performance is poor because of the read-ahead issues.
2850730	LM conformance hits an assert <code>f:vx_do_getpage:6b,3</code> and panics.
2417858	VxFS quotas do not support 64 bit limits.
2857629	File system corruption can occur requiring a full <code>fscck</code> of the system.
2590918	Delay in freeing unshared extents upon primary switch over.
2885592	<code>vxdump</code> to the <code>vxcompress</code> file system is aborted.

Veritas File System: issues fixed in 6.0.1

This section describes the incidents that are fixed in Veritas File System in this release.

Table 1-15 Veritas File System fixed issues

Incident	Description
2781322	VxFS returning error 61493 (VX_EFCLNOSPC) on CFS.
2764861	Uncompress by <code>vxcompress</code> ignores quota limitation.
2753944	The file creation threads can hang.
2735912	The performance of tier relocation using <code>fspadm enforce</code> is poor when moving a large amount of files.
2715186	System panic spinlock: locker forgot to unlock.
2712392	Threads hung in VxFS.
2709869	System panic with redzone violation when <code>vx_free()</code> tried to free <code>fiostat</code> .
2696067	When a <code>getaccess()</code> command is issued on a file which inherits the default Access Control List (ACL) entries from the parent, it shows incorrect group object permissions.
2670022	Duplicate file names can be seen in a directory.
2655788	Using cross-platform data sharing to convert a file system that has more than 32k <code>nlinks</code> does not update the <code>vx_maxlink</code> and <code>maxlink_enable</code> tunables.

Table 1-15 Veritas File System fixed issues (*continued*)

Incident	Description
2651922	Is -l command on local VxFS file system is running slow and high CPU usage is seen.
2599590	Expanding or shrinking a DLV5 file system using the fsadm(1M)command causes a system panic.
2597347	fsck should not coredump when only one of the device record has been corrupted and the replica is intact.
2566875	The write(2) operation exceeding the quota limit fails with an EDQUOT error (Disc quota exceeded) before the user quota limit is reached.
2559450	Command fsck_vxfs(1m) may core-dump with SEGV_ACCERR error.
2555198	sendfile() does not create DMAPI events for Hierarchical Storage Management(HSM) on VxFS.
2536130	fscdsconv fails to convert FS between specific platforms if FCL is enabled.
2272072	GAB panics the box because VCS engine HAD did not respond. The lobolt wraps around.
2183320	VxFS mmap performance degradation on HP-UX 11.31.
2086902	Spinlock held too long on vxfs spinlock, and there is high contention for it.

Veritas Cluster Server: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas Cluster Server (VCS) in 6.0.5

Table 1-16 Veritas Cluster Server 6.0.5 fixed issues

Fixed issues	Description
1919203	Add Health check monitoring for Oracle Agent.
2848020	When IP is unplumbed or network cable is pulled, the SambaShare agent fails to detect the fault.
3028760	The NFSRestart resource does not start the NFS processes such as statd and lockd, during the online or offline operation.
3042450	A parent service group which is frozen and configured with online local hard dependency is brought offline when its child service group faults.

Table 1-16 Veritas Cluster Server 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3079893	The value of LastSuccess attribute of the service group equals the GlobalCounter value of the cluster if the resource faults while you online the service group.Hence the service group fails to come online.
3090229	The Asynchronous Monitoring Framework (AMF) driver panics the node when the vxconfigd daemon is unresponsive.
3090710	The High Availability Daemon (HAD) starts and stops before the VxFEN driver configuration completes.
3097342	The Asynchronous Monitoring Framework (AMF) driver causes a panic in the node when AMF is being stopped.
3101761	The vcsauthserver process dumps core due to issues in VxAT library.
3104071	The service group online propagate operation fails without giving proper error message.
3106493	If for some reason, kernel components of the Veritas Cluster Server (VCS) software stack are stopped and restarted in quick succession, then during a restart, the cluster communication may fail.
3112608	Resource fails to become online after switch operations fails for a service group.
3117829	A very high memory load on a system may cause a panic in the Cluster File System (CFS) node.
3125918	The Asynchronous Monitoring Framework (AMF) driver causes a panic in the node when the vxconfigd process is unresponsive.
3137377	High Availability Daemon (HAD) dumps core due to failure in the allocation memory.
3140359	Global Atomic Broadcast (GAB) fails to start when the gabconfig -c and gabconfig -cx commands are executed simultaneously on the same system.
3145047	The Asynchronous Monitoring Framework (AMF) driver causes a panic in the node after VXFS driver is unloaded.
3153987	In the Application agent, the clean operation is reported successful even when the CleanProgram returns a non-zero value.
3154104	For Application agent, an error message is logged when the StartProgram or StopProgram returns a non-zero value. This gives incorrect implication of the program failure.

Table 1-16 Veritas Cluster Server 6.0.5 fixed issues (*continued*)

Fixed issues	Description
3207190	After Integrity Virtual Machine (VM) is migrated, the HPVirtualMachine resource may fail to perform the failover operation for VM.
3207663	Incorrect user privileges are set in case of incorrect use of the '-group' option in command "hauser -addprive".
3211834	CurrentLimits attribute value is not updated correctly when a service group faults.
3233895	Error message does not specify the source directory for the missing detailed monitoring script of the Db2udb agent.
3240209	During the Oracle online operation, due to an incorrect pattern match, the Oracle agent unnecessarily tries to back up the database.
3246141	The vxfenswap(1M) utility does not work in the clusters where rsh/ssh logins to other nodes are disabled.
3259682	If vxconfigd daemon hangs, then the registration thread of IMFD daemon trying to get disk group status from vxconfigd daemon also hangs. Therefore, the amfregister command waiting for IMFD daemon gets stuck.
3302143	Mount agent fails to bring the NFS file system online, when the mount options are configured.
3318764	Unexpected deletion of temporary files causes the VCS agents to report an incorrect state.
3341320	The "Cannot delete event (rid %d) in reaper" error message is repeatedly logged in the Syslog file.
3347536	The Application agent may dump a core while registering the resource with Asynchronous Monitoring Framework (AMF).
3362108	The system panics if LLT receives a corrupt packet from the network.
3409593	The ASMDG agent shows offline before volumes are released and service group fail-over will be delayed because volumes won't stop.

Veritas Cluster Server 6.0.3 fixed issues

[Table 1-17](#) lists the Veritas Cluster Server issues fixed in this release.

Table 1-17 Veritas Cluster Server fixed issues

Incident	Description
2736627	The remote cluster remains in "INIT" state and the ICMP heartbeat status is "UNKNOWN".
2737653	If you override the OnlineTimeout attribute value for the RVGPrimary resource, the agent does not consider it.
2848009	Asynchronous Monitoring Framework (AMF) panics the system when an agent exits.
2861253	In the vxfen driver debug log message, the jeopardy membership status is printed as garbage.
2937673	The AMF driver panics the system when the amfstat utility is executed.
2941155	Veritas Cluster Server (VCS) does not report the group as offline on a failed cluster when the cluster failure is declared in the global cluster configuration.
2964772	If the NFSRestart resource is taken offline, the NFSRestart agent may unexpectedly stop the NFS processes in a local container.
3013940	<p>In non-MPP mode, when no virtual host is configured in the db2nodes.cfg file, the db2start command fails with the following message:</p> <pre>A communication error occurred during START or STOP DATABASE MANAGER processing. SQL1032N No start database manager command was issued. SQLSTATE=57019.</pre>
3013962	In DB2 10.1, the monitor fails to detect the online DB2 instance.

Veritas Cluster Server: Issues fixed in 6.0.1

This section describes Veritas Cluster Server fixed issues in 6.0.1.

LLT, GAB, and I/O fencing fixed issues in 6.0.1

[Table 1-18](#) lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-18 LLT, GAB, and I/O fencing fixed issues

Incident	Description
2845244	<p><code>vxfen</code> startup script gives error <code>grep: can't open /etc/vxfen.d/data/cp_uid_db</code>.</p> <p>The error comes because <code>vxfen</code> startup script tries to read a file that might not be present. This error is typically seen when starting <code>vxfen</code> for the very first time after installation.</p>
2554167	Setting <code>peerinact</code> value to 0 in the <code>/etc/llttab</code> file floods the system log file with large number of log messages.

Bundled agents fixed issues in 6.0.1

Table 1-19 lists the fixed issues for bundled agents.

Table 1-19 Bundled agents fixed issues

Incident	Description
2850904	Concurrency violation and data corruption of a Volume resource may occur, if storage connectivity is lost or all paths under VxDMP are disabled and <code>PanicSystemOnDGLoss</code> is set to 0.
2728802	If the 'httpd' binary or the 'ab' binary is not present at the location that you specified in the 'httpdDir' attribute, the Apache agent cannot perform detail monitoring or start the HTTP server.
2850905	IMF registration for Mount resource for file systems type other than VxFS and NFS should be blocked.
2850916	Mount resource does not get registered with IMF if the attributes <code>BlockDevice</code> and/or <code>MountPoint</code> have a trailing slash in their values.
2850917	LVMLogicalVolume agent does not support zero byte logical volume and the resource goes into a faulted state.
2822920	DNSAgent goes to UNKNOWN state if the Top Level Domain (TLD) is more than 4 characters in length.
2846389	In releases prior to VCS 6.0.1, the upper bound value of <code>FaultTolerance</code> attribute of the <code>CoordPoint</code> agent was the one less than the number of coordination points. If the majority number of coordination points fault, the entire cluster panicked under network partition scenario. Therefore, the upper bound value of the <code>FaultTolerance</code> attribute of <code>CoordPoint</code> agent had to be set to less than the majority of the coordination points. Subsequent to VCS 6.0.1, the <code>FaultTolerance</code> attribute of <code>CoordPoint</code> agent is less than the majority of coordination points.

VCS engine fixed issues in 6.0.1

Table 1-20 lists the fixed issues for VCS engine.

Table 1-20 VCS engine fixed issues

Incident	Description
2879413	You may see two instances of CmdServer running on a node. One of these using IPv4 and the other IPv6.
2832754	When a Global Cluster Option (GCO) is configured across clusters having duplicate system names, command-line utility <code>hagrp</code> gives incorrect output with the "-clear", "-flush", "-state" options.
2741299	CmdSlave gets stuck in a tight loop when it gets an EBADF on a file descriptor(fd). The CmdSlave process keeps retrying on the FD and eventually dumps core.
2850906	If a group is auto-enabled, the engine clears the Start attribute even if the resource is online.
2692173	Engine does not check whether remote parent is online when <code>-nopre</code> option is selected.
2684818	If the following attributes are specified before SystemList attribute in main.cf, then the value got rejected when HAD started: <ul style="list-style-type: none"> ■ PreOnline ■ ContainerInfo ■ TriggersEnabled ■ SystemZones
2696056	Memory leak occurs in the engine when <code>haclus -status <cluster></code> command is run.
2746802	When failover group is probed, VCS engine clears the MigrateQ and TargetCount.
2746816	The syslog call used in <code>gab_heartbeat_alarm_handler</code> and <code>gabsim_heartbeat_alarm_handler</code> functions is not async signal safe.

Enterprise agents fixed issues in 6.0.1

Table 1-21 lists the fixed issues for enterprise agents.

Table 1-21 Enterprise agents fixed issues

Incident	Description
1985093	Ensure 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.
2831044	Sybase agent script entry points must handle large process command line.

Agent framework fixed issues in 6.0.1

[Table 1-22](#) lists the fixed issues for agent framework.

Table 1-22 Agent framework fixed issues

Incident	Description
2660011	Resource moves to FAULTED state even if value of ManageFaults attribute is set to NONE at service group level. This will cause service group to fault if the resource is Critical.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 6.0.5

This section describes the incidents fixed in Veritas Storage Foundation for Oracle RAC in 6.0.5.

Table 1-23 Veritas Storage Foundation for Oracle RAC 6.0.5 fixed issues

Fixed issues	Description
3090447	The CRSResource agent does not support the C shell (csh) environment.

Veritas Storage Foundation for Oracle RAC 6.0.3 fixed issues

There are no issues fixed in SF Oracle RAC 6.0.3.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 6.0.1

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 6.0.1.

Issues fixed in 6.0.1

[Table 1-24](#) lists the issues fixed in 6.0.1.

Table 1-24 Issues fixed in 6.0.1

Incident	Description
2585899	<p>The SF Oracle RAC 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:</p> <pre>The node sys1 doesn't seem to be part of the cluster, or CVM is not running on the node sys1.</pre>
2329580	<p>If you install and start SFHA, but later configure SFHA using <code>installvcs</code>, 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.</p>
2392741	<p>Policy-managed Oracle RAC databases fail to come online on some of the nodes in the server pool.</p> <p>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.</p>
2749412	<p>Setting the <code>UseVirtualIP</code> attribute to 1 overwrites the IP address of the virtual interface on some nodes in the cluster.</p>
2628469	<p>Installation fails to remove the following 5.0.1 depots: <code>Base-VXFS</code>, <code>Base-VxFS-501</code>, <code>Base-VxTools-501</code>, <code>Base-VxVM-501</code></p>
2622786	<p>Minimum installation of SFHA does not install <code>VRTSvcsea</code> depot.</p> <p>The installer does not install the <code>VRTSvcsea</code> depot when you perform a minimum installation of SFHA by selecting the option Install minimum required depots.</p>
2757032	<p>The <code>PrivNIC/MultiPrivNIC</code> agents fail to match the exact IP address configured in the agent configuration with the IP address configured on the system. As a result, the agent detects the wrong interface as the active interface resulting in a resource fault.</p>
2580393	<p>Removal of SAN cable from any node in a global cluster setup takes application service groups offline on all nodes.</p> <p>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.</p>

Table 1-24 Issues fixed in 6.0.1 (*continued*)

Incident	Description
2734745	The PrivNIC resource faults after the UseVirtualIP attribute is set to 1.
2740150	The SF Oracle RAC installer fails to set the value of the CSSD resource attribute OfflineWaitLimit to 3.
2746948	Some drivers fail to add to the system. Sometimes during bootup, some of the drivers fail to add in the system because of add_drv/rem_drv race between our modules which are independent of each other.
2532432	Presence of multiple VRTSperl versions after operating system upgrade causes product upgrade to fail. When you upgrade the operating system from HP-UX 11i Version 2 to HP-UX 11i Version 3 September 2011 or later, the <code>swinstall</code> command fails to remove the lower version of VRTSperl depot before installing the higher version. As a result, multiple versions of VRTSperl depot exist on the system after the operating system upgrade causing the product upgrade to fail.

Veritas Storage Foundation for databases (SFDB) tools: Issues fixed in 6.0.5

This section describes the incidents that are fixed in Veritas Storage Foundation for databases (SFDB) tools in 6.0.5

Table 1-25 Veritas Storage Foundation for databases (SFDB) tools 6.0.5 fixed issues

Fixed issues	Description
2715323	The DBED operations may not work with the non-standard Oracle database character sets like ZHS16GBK.
3237852	Oracle 12c database is not supported. SYMPTOM: Oracle 12c database is not supported.
3290416	Some DBED operations may fail with the following error message: "ORA-01406: fetched column value was truncated".
3211388	While cloning a Veritas Database Edition (DBED) instant checkpoint, if you enable the Block Change Tracking feature, the error message ORA-00600 is displayed.

Storage Foundation for Databases (SFDB) tools: Issues fixed in 6.0.3

[Table 1-26](#) lists the Veritas Storage Foundation for Databases (SFDB) tools fixed issues in this release.

Table 1-26 Veritas Storage Foundation for Databases (SFDB) tools fixed issues

Incident	Description
3030663	dbed_vmclonedb does not read pfile supplied by -p 'pfile_modification_file' option.

Veritas Storage Foundation for databases (SFDB) tools: Issues fixed in 6.0.1

This section describes Veritas Storage Foundation for databases (SFDB) tools fixed issues in 6.0.1.

Storage Foundation for Databases (SFDB) tools: issues fixed in 6.0.1

[Table 1-27](#) describes the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-27 SFDB tools fixed issues

Incident	Description
2585643	<p>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:</p> <pre>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.</pre> <p>The error messages are unclear.</p>
2703881 (2534422)	<p>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:</p> <pre>SFAE Error:0642: Storage for diskgroup oradatadg is not splittable.</pre>

Table 1-27 SFDB tools fixed issues (*continued*)

Incident	Description
2582694 (2580318)	After you have done FlashSnap cloning using a snapplan, any further attempts to create a clone from the same snapplan using the <code>dbed_vmclonedb</code> continue to use the original clone SID, rather than the new SID specified using the <code>new_sid</code> 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.
2579929	<p>The <code>sfae_auth_op -o auth_user</code> command, used for authorizing users, fails with the following error message:</p> <pre>SFDB vxsfadm ERROR V-81-0384 Unable to store credentials for <username></pre> <p>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.</p>

Known Issues

This section lists the known issues in this release.

- [Issues related to installation](#)
- [Issues related to any OS or supported technology](#)
- [Veritas Storage Foundation known issues](#)
- [Veritas Cluster Server known issues](#)
- [Veritas Storage Foundation for Oracle RAC known issues](#)
- [Veritas Storage Foundation Cluster File System High Availability \(SFCFSHA\) known issues](#)
- [Veritas Storage Foundation for Databases \(SFDB\) tools known issues](#)

Issues related to installation

This section describes known issues related to installation in this release.

Warning messages may be seen during script-based installation (2615447)

When you install SFHA using the script-based installer, you may see the following warning message:

```
interpreter "/opt/VRTSperl/bin/perl" not found
```

Workaround: You must install perl to resolve the issue.

To install perl

- 1 Exit the installer.
- 2 Install the `VRTSperl` depot from the product media manually:

```
# cd /dvd_path/depot
# /usr/sbin/swinstall -x enforce_dependencies=false
-x autoreboot=false -s `pwd` VRTSperl
```

- 3 Start the installer again.

Errors recorded in the swremove logs of VRTSgab during VCS upgrade from 4.1 to 5.0.1 (1719136)

When VCS is upgraded from 4.1 to 5.0.1 on HP-UX 11i v3 using the Veritas product installer, the installer reports errors for GAB and errors are recorded in the swremove logs related to VRTSgab.

Workaround: You can safely ignore these error messages.

VCS agents dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the update-ux command (1630968)

On PA-RISC architecture, the VCS agents (Oracle, Netlsnr, Sybase, SybaseBk, MultiNICB, and so on) may dump core after the operating system is upgraded from HP-UX 11i v2 to HP-UX 11i v3 using the `update-ux` command.

This is because on HP-UX PA-RISC systems, the default thread stack size is limited to 64k. When the agent requires more than 64k stack memory, it may dump core due to SIGBUS error.

Workaround: Before running the `update-ux` command, edit the `/opt/VRTSvcs/bin/vcsenv` file to append following lines to it:

```
PLATFORM=`uname -s`
ARCHITECTURE=`uname -m`
if [ "${PLATFORM}" = "HP-UX" ] && [ "${ARCHITECTURE}" = "9000/800" ];
then
    PTHREAD_DEFAULT_STACK_SIZE=524288
    export PTHREAD_DEFAULT_STACK_SIZE
fi
```

After a locale change restart the vxconfig daemon (2417547)

You need to restart the vxconfig daemon after 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 manual 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 manual 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 not 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:

```
# vxdctl -c mode
```

- 2 On the CVM master node, upgrade the CVM protocol:

```
# vxdctl upgrade
```

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.

Error message seen in swagent.log after removing the 6.0.1 VRTS depots (2324553)

After removing the 6.0.1 VRTS depots and before rebooting the system, you sometimes see the following message in the `swagent.log` file:

```
vxfs mount: V-3-21272: mount option(s) incompatible with file system  
/dev/vg00/lvol1
```

This message appears because the VRTS depots are removed and the kernel is not yet loaded.

Workaround: Reboot the system.

Stopping the Web installer causes Device Busy error messages (2633924)

If you start the Web installer, and then perform an operation (such as prechecking, configuring, or uninstalling), you may get an error message saying the device is busy.

Workaround: Do one of the following:

- Kill the start.pl process.
- Start the webinstaller again. On the first Web page you see that the session is still active. Either take over this session and finish it or terminate it directly.

Erroneous resstatechange trigger warning

You may encounter the following warning when you restart resources:

```
CPI WARNING V-9-40-4317 The installer has detected that resstatechange  
trigger is configured by setting TriggerResStateChange attributes.
```

Workaround: In future releases, the resstatechange trigger will not be invoked when a resource is restarted. Instead, the resrestart trigger will be invoked if you set the TriggerResRestart attribute. The resrestart trigger is available in the current release. Refer to the VCS documentation for details.

Adding a node to a cluster fails if you did not set up passwordless ssh or remsh

Adding a node to a cluster fails if you did not set up passwordless `ssh` or `remsh` prior to running the `./installsfcfsha<version> -addnode` command.

Workaround: Set up passwordless `ssh` or `remsh`, and then run the `./installsfcfsha<version> -addnode` command.

Where `<version>` is the current release version.

The VRTSacclib depot is deprecated (2032052)

The VRTSacclib depot is deprecated. For installation, uninstallation, and upgrades, note the following:

- Fresh installs: Do not install VRTSacclib.
- Upgrade: Ignore VRTSacclib.
- Uninstall: Ignore VRTSacclib.

Partial upgrade causes “installmr -version” to detect wrong product version [3438634]

If any packages of the product are upgraded to higher version by installer or manually, the `./installmr -version` command detects the product as a 6.0.5 product.

Workaround: No workaround.

NetBackup 6.5 or older version is installed on a VxFS file system (2056282)

If you have NetBackup 6.5 or older version installed on a VxFS file system and before upgrading to Veritas Storage Foundation (SF) 6.0.1, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (`/usr/opensv`), then while upgrading to SF 6.0.1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure depots `VRTSspbx`, `VRTSat`, and `VRTSicisco`. This causes NetBackup to stop working.

Workaround: Before you unmount the VxFS file system that hosts NetBackup, copy the `/usr/opensv/netbackup/bin/version` file and `/usr/opensv/netbackup/version` file to the `/tmp` directory. If you have clustered NetBackup installed, you must also copy the `/usr/opensv/netbackup/bin/cluster/NBU_RSP` file to the `/tmp` directory. After you unmount the NetBackup file system, manually copy these two version files from `/tmp` to their original directories. If you have clustered NetBackup installed, you must also copy the `/usr/opensv/netbackup/bin/cluster/NBU_RSP` file from `/tmp` to its original directory.

If the `version` files' directories do not exist, create the directories:

```
# mkdir -p /usr/opensv/netbackup/bin
```

Run the installer to finish the upgrade process. After upgrade process completes, remove the two version files and their directories.

If your system is already affected by this issue, then you must manually install the `VRTSpxb`, `VRTSat`, and `VRTSicisco` depots after the upgrade process completes.

Issues related to any OS or supported technology

This section describes known issues related to any Operating System (OS) or supported technology in this release.

Veritas Storage Foundation known issues

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

- [Veritas File System known issues](#)
- [Veritas Volume Manager known issues](#)
- [Veritas Dynamic Multi-pathing known issues](#)

Veritas File System known issues

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

Mismatched summary and bitmap data for filesystem allocation unit may cause full fsck flag set on file system [3451284]

If extent allocation happens during write operations, mismatched summary and bitmap data for filesystem allocation unit may cause full fsck flag set on file system.

Workaround:

To correct the summary and bitmap data for filesystem allocation unit, run the `fsck` command on the file system. Enter:

```
# fsck -o full -y device_name
```

Enabling delayed allocation on a small file system sometimes disables the file system (2389318)

When you enable delayed allocation on a small file system, such as around 100 MB, the file system can get disabled. In this case, the following error message displays in the system console log:

```
msg 001: V-2-1: vx_nospace - file_system file system full  
(size block extent)
```

Workaround: Use the `vxtunefs` command to turn off delayed allocation for the file system.

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.

Deleting a large number of files at the same time drastically increases CPU usage (2129455)

When you delete a large number of files at the same time, the CPU usage drastically increases beyond what you should expect.

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/rdisk/dg1/voll -
blocks are currently in use.
VxVM vxresize ERROR V-5-1-7514 Problem running fsadm command for volume
voll, in diskgroup dg1
```

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

Debug kernel panics with `spin_deadlock_failure` panic string while enabling auxiliary swap space (2521695)

The debug kernel panics with a `spin_deadlock_failure` panic string while enabling auxiliary swap space. The following example is of the relevant part of the stack trace:

```
spinlock+0x50
vx_inactive+0x140
vx_vn_inactive+0x30
vn_rele_inactive+0x1e0
vx_dnlc_getpathname+0x12b0
```

Not all partitioned directory entries display after exporting a VxFS file system over an HP-UX NFS server (2623412)

After you export a VxFS file system over an HP-UX NFS server, the file system might not list all of the entries in partitioned directories if accessed by NFS clients. This issue is specific to HP-UX NFS servers and VxFS disk layout Version 8 and later.

Workaround: There is no workaround for this issue.

Possible assertion failure in `vx_freeze_block_threads_all()` (2244932)

There is a possible assertion failure in the `vx_freeze_block_threads_all()` call when the `pdir_threshold` tunable is set to 1.

Workaround: There is no workaround for this issue.

fsppadm operations issued on multi-volume file system fail if there are other mounted file systems with a disk layout Version less than 6 (2909203)

The `fsppadm` command checks all mounted file systems, and if it finds any file systems with a disk layout Version that is less than 6, then it exits with the following error message:

```
# fsppadm assign /dst_vset /tmp/pol_test.xml
```

```
UX:vxfs fsppadm: ERROR: V-3-26510: Low level Volume enumeration failure
on / with message Function not implemented
```

This error occurs because the `fsppadm` command functionality is not supported on a disk layout Version that is less than 6.

Workaround: There is no workaround for this issue.

Veritas Dynamic Multi-pathing known issues

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

Enclosure name limitation when using HP-UX LVM `vgcreate` command on DMP device [1980759]

For HP-UX LVM on a DMP device, you cannot use the `vgcreate` command if the enclosure-based name of the DMP device contains the 's' character. This is a limitation of the `vgcreate` utility on HP-UX LVM.

Workaround: Before you run the `vgcreate` command, rename the enclosure to replace the 's' with some other character in the name of the enclosure. To rename the enclosure, use the following command:

```
# vxdmppadm setattr enclosure  
enclr_name name=new_enclr_name
```

Path name character limit when converting LVM volumes over DMP to VxVM volumes over DMP (2035399)

The HP-UX `lvdisplay` utility truncates physical volume path names to 22 characters. If a path name is truncated, utilities such as `vxvmconvert` or `vxautoconvert` that depend on the `lvdisplay` output may not function properly. If you intend to use the `vxvmconvert` utility or the `vxautoconvert` utility to convert LVM over DMP to VxVM over DMP, Symantec recommends that you reduce the length of the enclosure name to at most 8 characters before enabling native stack support.

DMP path discovery behavior when a device is removed from PowerPath control (2144891)

To remove a device from PowerPath control, you use the `powermt unmanage` command. When you remove a device from PowerPath control, DMP requires two device discovery cycles to discover the attributes of the paths of the device correctly.

Workaround:

Issue the following command to start the device discovery:

```
# vxdisk scandisks
```

After the discovery completes, issue the command again to start a second device discovery cycle.

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 6.0.5 (2082414)

The Veritas Volume Manager (VxVM) 6.0.5 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 6.0.5, 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 6.0.5. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-28 shows the Hitachi arrays that have new array names.

Table 1-28 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 6.0.5. 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

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.

Enclosure name limitation when using HP-UX LVM pvcreate command on DMP device (1980759)

For HP-UX LVM on a DMP device, you cannot use the `pvcreate` command if the enclosure-based name of the DMP device contains the 's' character. This is a limitation of the `pvcreate` utility on HP-UX LVM.

Workaround:

Rename the enclosure to replace the 's' with some other character in the name of the enclosure before you run the `pvcreate` command. To rename the enclosure, use the following command:

```
# vxdlmpadm setattr enclosure enclr_name name=new_enclr_name
```

After disconnecting and reconnecting the Fibre Channel, DMP is unable to present the device tree (2509636)

On some HP-UX 11i version 3 systems, after disconnecting and reconnecting the Fibre Channel, DMP is unable to present the device tree.

Workaround:

Restart the `vxconfigd` daemon with the following command:

```
# vxconfigd -k
```

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.

The pvcreate command failed with DMP devices on IA machines (2482178)

When `dmp_native_support` is enabled on HP-UX 11i version 3 systems, you must run the `pvcreate` command on `/dev/disk/<disk#>` before creating a LVM volume group on the corresponding DMP device.

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.

After LUNs remapped using different target ID, DMP reports error with device discovery (2526605)

After LUNs are re-mapped using different target IDs, device discovery fails with the following error message:

```
VxVM vxdisk ERROR V-5-1-16007 Data Corruption Protection Activated -  
User Corrective Action Needed To recover, first ensure that the OS  
device tree is up to date (requires OS specific commands).
```

Workaround:

To recover from this issue

- 1 Use Operating System (OS) commands to ensure that the OS device tree is up to date.
- 2 Remove the specified devices from VxVM control:

```
# vxdisk rm devicename
```

- 3 Restart device discovery.

```
# vxdisk scandisks
```

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:

```
# vxddmpadm settune dmp_native_support=on
```

Some DMP commands do not display proper output [3438692]

Some DMP command outputs may have blank entries for devices that are claimed under Tachyon or Qlogic CNA driver. For example:

```
# vxddmpadm getctrl all
```

LNAME	PNAME	VENDOR	CTLR-ID
c49	0/2/1/0.0x5006016939a01884	-	-
c50	0/2/1/0.0x5006016239a01884	-	-
c53	0/3/1/0.0x5006016939a01884	-	-
c54	0/3/1/0.0x5006016239a01884	-	-


```
# vxddladm list hbas
```

HBA-ID	DRIVER	FIRMWARE	DISCOVERY	STATE	ADDRESS
c43	-	-	-	Online	-
c38	-	-	-	Online	-
c50	-	-	-	Online	-
c49	-	-	-	Online	-
c9	-	-	-	Online	-
c21	-	-	-	Online	-

Tachyon driver (shipped as part of FibrChanl-00 product) is out of support since December 2012, and Qlogic CAN driver (shipped as part of FibrChanl-04 product) is also not supported. Hence the multi-threaded SNIA libraries that are needed to get Host bus adapter (HBA) attributes for the mentioned drivers are not available.

If Tachyon or Qlogic CNA HBAs are used on the system, consequently Event Source Daemon (`vxesd`) fails to get additional attributes for these HBAs. Hence some fields are shown as empty in the output of DMP commands.

Workaround: Users of Tachyon or Qlogic CNA drivers should migrate to supported drivers or configurations to get the needed functionality.

Veritas Volume Manager known issues

This section describes Veritas Volume Manager known issues in this release.

Dynamic LUN expansion is not supported for EFI disks in simple or sliced formats (2836798)

Dynamic LUN expansion is not supported for EFI (Extensible Firmware Interface) disks in simple or sliced formats. The recommended format is the Cross-platform Data Sharing (CDS) disk format.

Workaround:

Convert the disk format to CDS using the `vxcdsconvert` utility.

On HP-UX 11.31, vxdiskadm option 22-2 Dynamic Reconfiguration (DR) operation 'Remove Luns' might fail with error (2957766)

When the user tries to remove LUNs from the system using `vxdiskadm` option 22-2 Dynamic Reconfiguration operation 'Remove Luns', the device removal operation fails and reports the following error message:

```
ERROR: Please make sure to remove Luns from Array
```

This is due to the Dynamic Reconfiguration Tool not being able to find devices that are not part of the legacy HP-UX I/O device tree but are seen only in the agile I/O device tree.

Workaround:**Perform the following steps:**

- 1 Remove the device with no hardware (NO_HW in output of 'ioscan -fNC disk') using `rmsf(1M)`.
- 2 Run `ioscan(1M)`.
- 3 Run `vxdisk scandisks`.

Veritas Volume Manager Event Source daemon (vxesd) requires certain drivers and packages to use version 11.31.1311 or later (3473334)

Systems running the HP-UX operating system must ensure that the following drivers and packages use version 11.31.1311 or later to take advantage of the `vxesd` daemon.

```
CommonIO, r=B.11.31.1311  
FibrChanl-01, r=B.11.31.1311  
FibrChanl-02, r=B.11.31.1311  
FibrChanl-03, r=B.11.31.1311
```

The `vxesd` daemon stops working with versions lower than 11.31.1311.

The `vxesd` daemon does not depend upon the version of the following drivers or packages:

```
FibrChan1-00 B.11.31.1003 FibreChannel;HW=A6795A,A5158A
FibrChan1-04 B.11.31.1303 FCQ;HW=P3P
```

The SCSI registration keys are not removed even if you stop VCS engine for the second time (3037620)

If you stop VCS engine for the first time, the SCSI registration keys can be removed. But if you stop VCS engine for the second time, the keys are not removed.

Workaround:

There is no workaround for this issue.

In some cases with large LUN setup, the storage disappears after DMP device scan (2828328)

This issue is typically seen on a large LUN setup. In some cases, the storage disappears after the DMP device scan. The DMP device scan is generated with the `vxdisk scandisks` command or the `vxctl enable` command. Even if the OS command `ioscan` can discover devices, VxVM/DMP cannot.

Workaround:

Restarting the `vxconfigd` daemon on the affected node may resolve the issue. If that does not work, you must reboot the system.

Dynamic LUN expansion is not supported for EFI disks in simple or sliced formats (2836798)

Dynamic LUN expansion is not supported for EFI (Extensible Firmware Interface) disks in simple or sliced formats. It may lead to corruption. The recommended format is the Cross-platform Data Sharing (CDS) disk format.

Workaround:

Convert the disk format to CDS using the `vxcdsconvert` utility.

System may not boot from a VxVM root disk on a thin LUN (2753626)

The system may fail to boot from a VxVM root disk on a thin LUN. This is an intermittent issue seen only with thin LUNs. The boot process aborts with the following error:

```
System Console is on the Built-In Serial Interface
AF_INET socket/streams output daemon running, pid 52
afinet_prelink: module installed
Starting the STREAMS daemons-phase 1
NOTICE: reading the krs value is failed rc 2
Swap device table: (start & size given in 512-byte blocks)
entry 0 - major is 2, minor is 0x1; start = 0, size = 6242304
```

```
Starting vxconfigd in boot mode (pre_init_rc).
pre_init_rc[86]: 81 Illegal instruction
Error returned from vxconfigd -m boot, halting
ERROR: The configuration could not be locked. It may be in use by
another process.
Calling function e000000001a98660 for Shutdown State 1 type 0x1
```

Workaround:

In most cases, rebooting the system resolves the issue.

The vxdmp and other drivers have the incorrect release version (2878024)

The `vxdmp` and other drivers have the incorrect release version. The version displays as 50.0, as shown in the following output:

```
kcmodule -v vxdmp
Module                vxdmp    (50.0)
Description            VxVM DMP Subsystem
Timestamp              Wed Aug  1 10:17:12 2012 [50195688]
State                  static (best state)
State at Next Boot    static (best state)
Capable                static unused
Depends On             interface HPUX_11_31_PERF:1.0
```

vxdg split or join operations can fail for disks with a disk media name greater than or equal to 27 characters (2063387)

If a disk's media name is greater than or equal to 27 characters, certain operations, such as diskgroup split or join, can fail with the following error:

```
VxVM vxdg ERROR : vxdg move/join dg1
                  dg2 failed subdisk_name : Record
already exists in disk group
```

VxVM uses disk media names to create subdisk names. If multiple subdisks are under the same disk, then the serial number, starting from 1, is generated and appended to the subdisk name so as to identify the given subdisk under the physical disk. The maximum length of the subdisk name is 31 characters. If the disk media name is long, then the name is truncated to make room for serial numbers. Therefore, two diskgroups can end up having same subdisk names due to this truncation logic, despite having unique disk media names across diskgroups. In such scenarios, the diskgroup split or join operation fails.

Workaround:

To avoid such problems, Symantec recommends that disk media name length should be less than 27 characters.

The vxrecover command does not handle RAID5 volumes correctly (2715124)

The `vxrecover` command calls the recovery process for the top-level volume, which internally takes care of recovering its subvolumes. The `vxrecover` command does not handle RAID5 volumes correctly. The recovery process fails to recover the subvolumes, which remain in the NEEDSYNC state.

Workaround:

Manually recover the RAID5 volumes using the `vxvol` utility, as follows:

```
# vxvol -g diskgroup resync volume
```

After initializing a disk for native LVM, the first instance of vxdisk list fails with a 'get_contents' error and errant flags are displayed (2074640)

After you initialize a disk that is under the operating system's native LVM control and not under Veritas Volume Manager (VxVM) control by using the `pvcreeate path_to_physical_disk` command, the first time that you run the `vxdisk list disk_name` command results in a VxVM error message related to `get_contents`, and the `flags` field is incorrectly populated. However, in the next instantiation of the same command, VxVM does not produce an error and the flags are correctly populated with the LVM tag.

Workaround:

Issue the `vxdisk list disk_name` command a second time.

vxconfig fails to allocate memory until the daemon is restarted (2112448)

Veritas Volume Manager (VxVM) utilities may fail with the following error message:

```
Memory allocation failure
```

This error implies that there is insufficient memory for the `vxconfigd` daemon. A program's data segment size is enforced by the operating system tunable `maxdsiz`. The default value of `maxdsiz` is 1 GB. With this default `maxdsiz` value, the `vxconfigd` daemon can allocate a maximum of 1 GB of memory.

Workaround:

You might need to increase the operating system `maxdsiz` tunable's value appropriately to increase the data storage segment for the programs.

See the `maxdsiz(5)` manual page for more information.

After increasing the value, you must stop and restart the `vxconfigd` daemon. Depending on the `maxdsiz` tunable value, `vxconfigd` can allocate a maximum up to 2 GB of memory on PA machines, and 4 GB of memory on IA machines.

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.

vxdisksetup fails on a LUN that is larger than 1 TB and has the cdsdisk format if the system is using Tachyon HBAs (2146340)

The `vxdisksetup` command fails to initialize a LUN that is larger than 1 TB and has the `cdsdisk` format if the system is using Tachyon HBAs. The `vxdisksetup` command displays the following error:

```
VxVM vxdisk ERROR V-5-1-5433 Device disk_name: init failed:  
Disk is not useable, bad format
```

Workaround:

There is no workaround for this issue.

The "vx dg listclone" command output may not list all the disks with "clone_disk" or "udid_mismatch" flag set (2354560)

In Cluster Volume Manager environment, "vx dg 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 "vx disk scandisks" or "vx disk -o all dgs list" followed by "vx dg listclone" to get all the disks containing "clone_disk" or "udid_mismatch" flag on respective host.

Known Issue related to EFI disk initialization (2585433)

For disks initialized with EFI format using `idisk`, DA record becomes invisible from "vx disk list" output after executing "vx disk scandisks".

Workaround:

For devices to be correctly seen with slices in "vx disk list" output, VxVM needs to flush the cached open and reopen the disk device. Further, VxVM needs to search for this new EFI format on the disk and generate new DA record.

To recover from this issue

- ◆ To achieve this functionality run following VxVM commands:

```
# vx disk rm <DANAME>  
# vx dtl cache flush  
# vx disk scandisks
```

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 cache object (DCO) volumes. That is, the command output may show less %dirty than actual.

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

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

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

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 `vxdatapadm disable` command. In this case, CVM may place the disks into the `lfailed` state. When you restore connectivity with the `vxdatapadm 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
```

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.

Workaround:

To work around this issue

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

During online migration from LVM to VxVM volumes, LVM sometimes incorrectly reports the remapped LVM device paths as valid LVM volumes

Problem: In a migrated or committed configuration, only the renamed LVM names of the form <lvname>_vxlv are valid LVM volumes. The original LVM names, in turn, point to target VxVM volumes. However, LVM sometimes incorrectly reports these original LVM device paths pointing to VxVM volumes, as valid LVM volumes.

Do not assume these as LVM volumes or do any operations on them, as it would disrupt the application's access to the target VxVM volumes.

Diskgroup import of BCV luns using -o updateid and -o useclonedev options is not supported if the diskgroup has mirrored volumes with DCO or has snapshots. (2831658)

VxVM uses guid stored in configuration to uniquely identify all objects. The DCO volume stores the guid of mirrors and snapshots. If the diskgroup is imported with -o updateid and -o useclonedev, it changes the guid of objects in VxVM configuration database and the guids stored in DCO volume are not updated. So the operations involving DCO will not be able to find objects with the stored guid and this could lead to failure of certain operations involving DCO or could lead to unexpected behaviour.

Workaround:

No workaround available.

LVM to VxVM convert operation fail to show the expected time [3419565]

The following error message is displayed when you convert LVM volume to VxVM with the `vxvmconvert` command that the expected time fails to display:

```
The expected time for convert is: syntax error on line 1,\  
hrs syntax error on line 1, mins syntax error on line 1, secs.
```

Workaround: No action is required. The message does no harm to the convert operation.

Failed to roll back the LVM that is created on DMP device [3421103]

If you create an LVM volume on DMP devices and convert the volume to VxVM, you can't roll back the volume from VxVM to LVM.

Workaround: No workaround.

Vradmin verifydata reports differences in case of mirrored or layered volumes with SmartMove enabled [3426434]

When the SmartMove utility is enabled, mirrored or layered volumes plexes are not fully synced. The `vradmin verifydata` command compares the checksum block wise, and reports differences on mirrored or layered volumes. The following error message is displayed:

```
VxVM VVR vxrsync INFO V-5-52-10190 \  
Verification of the remote volumes found differences.
```

Workaround: No workaround. Since it does not relate any data corruption, it is safe to ignore the message. You may want to use file checksum to verify whether the volumes are same.

Duplicate disk access (da) entries on vxdisk list (2705055)

If there is a change in the naming scheme and some disks in the disk group are not accessible, then duplicate disk access (da) entries will be visible on the same node.

Workaround: Perform the following steps to resolve the issue:

1. Remove the duplicate disk entry.

```
# vxdisk rm duplicate_da_name
```

2. Verify the disk.

```
# vxdisk scandisks
```

vradmin functionality may not work after a master switch operation (2163712)

In certain situations, if you switch the master role, `vradmin` functionality may not work. The following message displays:

```
VxVM VVR vxrlink ERROR V-5-1-15861 Command is not supported for  
command shipping. Operation must be executed on master
```

Workaround:

To restore vradmin functionality after a master switch operation

1 Restart `vradmind` on all cluster nodes. Enter the following:

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

2 Re-enter the command that failed.

Veritas Cluster Server known issues

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

- [Issues related to installing and upgrading VCS](#)
- [Operational issues for VCS](#)
- [Issues related to the VCS engine](#)
- [Issues related to the agent framework](#)
- [Issues related to global clusters](#)
- [LLT known issues](#)
- [I/O fencing known issues](#)
- [Issues related to Intelligent Monitoring Framework \(IMF\)](#)
- [Issues related to the Cluster Manager \(Java Console\)](#)

Operational issues for 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

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 [2653688]

When HAD is restarted by `hshadow` 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 parent dependency 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 [2490331]

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: Terminate 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 [2653695]

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 [2653695]

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.

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.

Elevated TargetCount prevents the online of a service group with `hagrp -online -sys` command [2871892]

When you initiate an offline of a service group and before the offline is complete, if you initiate a forced flush, the offline of the service group which was initiated earlier is treated as a fault. As start bits of the resources are already cleared, service group goes to OFFLINE|FAULTED state but TargetCount remains elevated.

Workaround: No workaround.

Auto failover does not happen in case of two successive primary and secondary cluster failures [2858187]

In case of three clusters (clus1, clus2, clus3) in a GCO with steward not configured, if clus1 loses connection with clus2, it sends the inquiry to clus3 to check the state of clus2 one of the following condition persists:

1. If it is able to confirm that clus2 is down, it will mark clus2 as FAULTED.
2. If it is not able to send the inquiry to clus3, it will assume that a network disconnect might have happened and mark clus2 as UNKNOWN

In second case, automatic failover does not take place even if the ClusterFailoverPolicy is set to Auto. You need to manually failover the global service groups.

Workaround: Configure steward at a geographically distinct location from the clusters to which the above stated condition is applicable.

GCO clusters remain in INIT state [2848006]

GCO clusters remain in INIT state after configuring GCO due to :

- Trust between two clusters is not properly set if clusters are secure.
- Firewall is not correctly configured to allow WAC port (14155).

Workaround: Make sure that above two conditions are rectified. Refer to *Veritas Cluster Server Administrator's Guide* for information on setting up Trust relationships between two clusters.

The `ha` commands may fail for non-root user if cluster is secure [2847998]

The `ha` commands fail to work if you first use a non-root user without a home directory and then create a home directory for the same user.

Workaround

- 1 Delete `/var/VRTSat/profile/<user_name>`.
- 2 Delete `/home/user_name/.VRTSat`.

- 3 Delete `/var/VRTSat_lhc/<cred_file>` file which same non-root user owns.
- 4 Run `ha` command with same non-root user (this will pass).

Older ClusterAddress remains plumbed on the node while modifying ClusterAddress [2858188]

If you execute `gcoconfig` to modify `ClusterAddress` when `ClusterService` group is online, the older `ClusterAddress` remains plumbed on the node.

Workaround: Un-plumb the older `ClusterAddress` from the node manually or offline `ClusterService` group by executing the following command before running `gcoconfig`:

```
hagrpl -offline -force ClusterService -any
```

or

```
hagrpl -offline -force ClusterService -sys <sys_name>
```

The ha commands may fail for non-root user if cluster is secure [2847998]

The `ha` commands fail to work if you first use a non-root user without a home directory and then create a home directory for the same user.

Workaround

- 1 Delete `/var/VRTSat/profile/<user_name>`,
- 2 Delete `/home/user_name/.VRTSat`.
- 3 Delete `/var/VRTSat_lhc/<cred_file>` file which same non-root user owns.
- 4 Run `ha` command with same non-root user (this will pass).

Issues related to the bundled agents

Process and ProcessOnOnly agent rejects attribute values with white spaces [2303513]

Process and ProcessOnOnly agent does not accept Arguments attribute values that are separated by multiple whitespaces. The Arguments attribute specifies the set of arguments for a process. If a script controls the process, the script is passed as an argument. You must separate multiple arguments by using a single whitespace. A string cannot accommodate more than one space between arguments, or allow leading or trailing whitespace characters. This attribute must not exceed 80 characters.

Workaround: You should use only single whitespace to separate the argument attribute values. Make sure you avoid multiple whitespaces between the argument attribute values or trailing whitespace characters.

Application agent cannot handle a case with user as root, envfile set and shell as csh [2490296]

Application agent does not handle a case when the user is root, envfile is set, and shell is csh. The application agent uses the `system` command to execute the `Start/Stop/Monitor/Clean` Programs for the root user. This executes `Start/Stop/Monitor/Clean` Programs in `sh` shell, due to which there is an error when root user has csh shell and EnvFile is written accordingly.

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

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

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

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

NIC agent may report incorrect interface state due to less traffic [2556355]

When PingOptimize is set to 1 and no NetworkHosts is specified, NIC agent depends on packet count to report the health of the interface. If the traffic on the interface is not sufficient enough, NIC agent may report incorrect state of the interface.

Workaround: Any of the following workaround must resolve the issue:

- Setting PingOptimize = 0. This makes NIC agent ping the broadcast address whenever there is no traffic on the interface.
- Setting valid NetworkHosts value. This makes NIC agent to ping NetworkHosts to check health of status.

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.

Resource gets faulted with zero byte logical volume of LVM agent [2393787]

LVM Agent does not support zero byte logical volume and the resource goes into faulted state.

When you configure the resource and try to bring it online, the resource goes into faulted state.

Workaround: No workaround.

CoordPoint agent remains in faulted state [2852872]

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

Workaround: After HAD has stopped, reconfigure fencing.

Process resource fails to come online if call to exec are present in the profile of the root user [2611530]

Process agent fails to bring the resource online if there are calls to exec in the shell profile file of the root user.

Workaround: Make sure there no calls to exec in the shell profile of the root user.

CoordPoint agent reports faulted if agile disk naming scheme is used [2871893]

If LevelTwo monitoring is enabled, i.e., if LevelTwoMonitorFreq attribute value is greater than zero, CoordPoint agent reports faulted if agile disk naming scheme is used on HP-UX platform.

Workaround: Use a disk naming scheme other than agile naming scheme.

Probing LVMLogicalVolume resource makes it OFFLINE when LVMLogicalVolume attribute is "" and volume group is not activated [2858185]

If volume group is not activated and the LogicalVolume attribute is set to "" (empty string), a probe of LVMLogicalVolume resource shows the resource state as OFFLINE instead of UNKNOWN.

Workaround: Activate the volume group before configuring LVMLogicalVolume resource under VCS control.

LVMVolumeGroup resource remains ONLINE if VolumeGroup is deactivated outside VCS [2858165]

If LVMVolumeGroup resource is deactivated outside of VCS, clean entry point cannot clean the resource. As a result, the resource state remains ONLINE and clean is called repeatedly.

Workaround: Do not deactivate LVMVolumeGroup from outside SFHA.

SambaShare agent clean entry point fails when access to configuration file on shared storage is lost [2858183]

When the Samba server configuration file is on shared storage and access to the shared storage is lost, SambaShare agent clean entry point fails.

Workaround: No workaround.

Concurrency violation in the service group [2870982]

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 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 disk group is not deported manually, then volume may get started if disk group is not deported during the service group failover. Thus, volume resource shows its state as online on both the nodes and thus causes concurrency violation. This may lead to data corruption.

Workaround: Ensure that the disk group is deported soon after storage connectivity is restored. Always configure volume resource whenever disk group resource is configured and set the attribute PanicSystemOnDGLoss to 1 or 2 as required.

Service group with LVMLogicalVolume resource does not failover if the underlying storage is not available (2916108)

If service group with LVMLogicalVolume resource is ONLINE and the underlying storage is disconnected or is unavailable, then it does not failover. This is because the clean entry point is not able to clean the resource, causing the service group to remain ONLINE.

Workaround: Make sure the underlying storage is always available.

NFS cluster I/O fails when storage is disabled [2555662]

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.

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.

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 with configuration of resource values (1718043)

If you configure a resource that has more than 425 values in its **ArgListValues**, the agent managing that resource logs a message such as:

```
VCS WARNING V-16-2-13806 Thread(1437547408) ArgListValues overflow;  
Cannot append values more than upper limit of (425).
```

Normally, the number of values in **ArgListValues** for a resource must not exceed 425. However, in case of a keylist, association or vector type of attribute appears in the ArgList for a resource-type. Since these attributes can take multiple values, there is a chance for the resource values in **ArgListValues** to exceed 425.

Delayed response to VCS commands observed on nodes with several resources and system has high CPU usage or high swap usage [3432749]

You may experience a delay of several minutes in the VCS response to commands if you configure large number of resources for monitoring on a VCS node and if the CPU usage is close to 100 percent or swap usage is very high.

Some of the commands are mentioned below:

- # hares -online
- # hares -offline
- # hagrps -online
- # hagrps -offline
- # hares -switch

The delay occurs as the related VCS agent does not get enough CPU bandwidth to process your command. The agent may also be busy processing large number of pending internal commands (such as periodic monitoring of each resource).

Workaround: Change the values of some VCS agent type attributes which are facing the issue and restore the original attribute values after the system returns to the normal CPU load.

- 1 Back up the original values of attributes such as MonitorInterval, OfflineMonitorInterval, and MonitorFreq of IMF attribute.
- 2 If the agent does not support Intelligent Monitoring Framework (IMF), increase the value of MonitorInterval and OfflineMonitorInterval attributes.

```
# haconf -makerw
# hatype -modify <TypeName> MonitorInterval <value>
# hatype -modify <TypeName> OfflineMonitorInterval <value>
# haconf -dump -makero
```

Where <TypeName> is the name of the agent with which you are facing delays and <value> is any numerical value appropriate for your environment.

- 3 If the agent supports IMF, increase the value of MonitorFreq attribute of IMF.

```
# haconf -makerw
# hatype -modify <TypeName> IMF -update MonitorFreq <value>
# haconf -dump -makero
```

Where <value> is any numerical value appropriate for your environment.

- 4 Wait for several minutes to ensure that VCS has executed all pending commands, and then execute any new VCS command.

- 5 If the delay persists, repeat step 2 or 3 as appropriate.
- 6 If the CPU usage returns to normal limits, revert the attribute changes to the backed up values to avoid the delay in detecting the resource fault.

Issues related to global clusters

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.

LLT known issues

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

On reboot of cluster nodes that are connected via a single switch, a race condition may cause one of the llt links to not come up (2848001)

If cluster nodes are connected via a single switch and nodes are rebooted multiple times then sometimes a race condition may cause one of the links to be down. Run the `lltstat -nvv` command to know the link that is down.

Workaround: Restart LLT on the rebooted node.

Cannot use CPI response files to add nodes to a cluster that is using LLT over UDP (2869763)

When you run the `addnode -responsefile` command, if the cluster is using LLT over UDP, then the `/etc/llttab` file generated on new nodes is not correct. So, the procedure fails and you cannot add nodes to a cluster using CPI response files.

Workaround: None

GAB known issues

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 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.

Workaround: There is no workaround for this issue.

GAB can panic due to had not responding (2166263)

GAB can panic due to `had` not responding. This is caused by threads becoming stuck in the `vx_event_wait()` call and the `vx_rwsleep_rec_lock_em()` call.

Workaround: There is no workaround for this issue.

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.

Workaround: There is no workaround for this issue.

I/O fencing known issues

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

Installer is unable to split a cluster that is registered with one or more CP servers (2110148)

Splitting a cluster that uses server-based fencing is currently not supported.

You can split a cluster into two and reconfigure Veritas Storage Foundation HA on the two clusters using the installer. For example, you can split a cluster `clus1` into `clus1A` and `clus1B`.

However, if you use the installer to reconfigure the Veritas Storage Foundation HA, the installer retains the same cluster UUID of `clus1` in both `clus1A` and `clus1B`. If both `clus1A` and `clus1B` use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: There is no workaround for this issue.

CoordPoint agent does not report the addition of new disks to a Coordinator disk group [2727672]

The LevelTwo monitoring of the CoordPoint agent does not report a fault even if the constituent of a coordinator disk group changes due to addition of new disks in the coordinator disk group

Workaround: There is no workaround for this issue.

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 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).

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@sys1,  
domaintype vx; not allowing action
```

The `vxfsend` 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 cluster, ensure that the cluster details such as cluster name, UUID, nodes, and privileges are added to the CP server.

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.

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 information to the CP server to resolve this issue. Alternatively, you can use installer as the installer adds cluster information to the CP server during configuration.

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.

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.

ping_cps and server_security fail to communicate with the secure CPS Server if the CPS variables are not exported manually (2791763)

The `ping_cps` and `server_security` commands fail to communicate with the secure CPS Server from a client if the CPS variables are not exported manually.

Workaround: Set and export the following variables manually on the client cluster:

```
# CPS_DOMAINTYPE="vx"
# export CPS_DOMAINTYPE
# EAT_HOME_DIR="/opt/VRTScps"
# export EAT_HOME_DIR
# CPS_HOME="/opt/VRTScps"
# export CPS_HOME
# CPS_USERNAME="CPSADM@VCS_SERVICES"
# export CPS_USERNAME
```

Hostname and username are case sensitive in CP server (2846392)

The hostname and username on the CP server are case sensitive. The hostname and username used by fencing to communicate with CP server must be in same case as present in CP server database, else fencing fails to start.

Workaround: Make sure that the same case is used in the hostname and username on the CP server.

Cannot run the vxfentsthdw utility directly from the install media if VRTSvxfen depot is not installed on the system (2858190)

If VRTSvxfen depot is not installed on the system, then certain script files that are needed for the vxfentsthdw utility to function are not available. So, without the VRTSvxfen depot installed on the system you cannot run the utility from the install media.

Workaround: Install VRTSvxfen depot, then run the utility from either the install media or from the `/opt/VRTSvcsvxfen/bin/` location.

After upgrading coordination point server in secure mode the cpsadm command may fail with error - Bus error (core dumped) (2846727)

After upgrading the coordination point server from SFHA 5.0 to the next version on the client system, if you do not remove the VRTSat depot that were installed on the system, the `cpsadm` command fails. The command fails because it loads old security libraries present on the system. The `cpsadm` command is also run on the coordination point server to add or upgrade client clusters. The command also fails on the server because it loads old security libraries present on the system.

Workaround: Perform the following steps on all the nodes on the coordination point server:

1 Rename `cpsadm` to `cpsadmbin`

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

2 Create the `/opt/VRTScps/bin/cpsadm` file with the following details.

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

3 Give executable permissions to the new file.

```
# chmod 775 /opt/VRTScps/bin/cpsadm
```

During a race scenario, the larger subcluster of a cluster can lose to a smaller subcluster, which may cause the large subcluster to panic (2858189)

It may happen that during a split-brain scenario, GAB and `vxfen` modules may take more time to confirm memberships of nodes on a larger subcluster than the time taken to for the same action on a smaller subcluster. So, GAB and `vxfen` modules on the larger subcluster may lose the race to confirm new node memberships. Hence, the larger subcluster may panic.

CoordPoint agent goes into faulted state if you change the disk naming scheme

If `LevelTwo` monitoring is enabled (`LevelTwoMonitorFreq` attribute value set to a value greater than zero), the `coordpoint` resource goes into faulted state if you use agile disk naming scheme.

Workaround: Use a disk naming scheme other than the agile naming scheme.

Fencing command, `vxfenadm`, does not print the registration keys in character format (2760308)

The `vxfenadm` command does not print character format of keys with leading NULL bytes. This behavior happens because the `vxfenadm` command prints entire registration key as a string and if there is a leading NULL byte in the string key the character format of the key is not printed.

Workaround: None

Server-based fencing may fail to start after reinstalling the stack (2802682)

Server-based fencing may fail to start if you use the existing configuration files after reinstalling the stack.

Workaround:

After reinstalling the stack, add the client cluster information on the coordination point server because the client cluster information is removed when the stack is uninstalled. For more details, see the Setting up server-based I/O Fencing manually section in the Veritas Storage Foundation and High Availability Solutions Installation Guide. Alternatively, you can manually modify the `/etc/vxfenmode` file and the `main.cf` file to start fencing in disable mode and then configure fencing.

Common product installer cannot setup trust between a client system on release version 5.1SP1 and a server on release version 6.0 or later (2824472)

The issue exists because the 5.1SP1 release version does not support separate directories for truststores. But, release version 6.0 and later support separate directories for truststores. So, because of this mismatch in support for truststores, you cannot set up trust between client systems and servers.

Workaround: Set up trust manually between the coordination point server and client systems using the `cpsat` or `vcsat` command. Now, the servers and client systems can communicate in secure mode.

Fencing may show the RFSM state as replaying for some nodes in the cluster (2555191)

Fencing based on coordination point clients in Campus cluster environment may show the RFSM state as replaying for some nodes in the cluster.

Workaround:

Restart fencing on the node that shows RFSM state as replaying.

vxfen module does not come up after phased upgrade from release version 4.1MP1 to 6.0.1 (2846209)

With HP-UX 11iv3, after upgrade, vxfen module does not allow raw disks to be specified as coordinator disks. So, even if you set the `vxfen_disk_policy` attribute to **raw** in the `/etc/vxfenmode` file fencing does not come up.

Workaround: Set the `vxfen_disk_policy` to `dmp` in the `/etc/vxfenmode` file.

Installation of Oracle 12c Clusterware using Oracle response file fails (3319554)

The installation of Oracle 12c Clusterware using Oracle response file fails with the following error:

There are issues using the DISPLAY value you provided. Either the DISPLAY variable has not been set properly or there are display connectivity problems.

This is because Oracle 12c Clusterware response file does not require the DISPLAY environment variable whereas the SF Oracle RAC installer requires it.

Workaround: Before starting the SF Oracle RAC installer, export the DISPLAY environment variable as follows:

```
$ Export DISPLAY=10.200.58.255:4
```

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.

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

When you configure fencing in customized mode and do not 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 at least one CP server. The default port value is 14250.

The cpsadm command fails after upgrading CP server to 6.0 or above in secure mode (2846727)

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTS at depot 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 procedure on all of the nodes of the CP server.

To resolve this issue

- 1 Rename `cpsadm` to `cpsadmbin`:

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

- 2 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 "$@"
```

- 3 Change the permissions of the new file to 775:

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

After you run the `vxfsnwap` utility the `CoordPoint` agent may fault (3462738)

After you run the `vxfsnwap` utility, if the value of the `FaultTolerance` attribute of the `CoordPoint` agent is more than the majority (more than 50%) of the coordination points then the `Coordination Point` agent faults.

Workaround: Manually set the value of the `FaultTolerance` attribute of `CoordPoint` agent to be less than the majority (more than 50%) of the coordination points.

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.

Perl errors seen while using haimfconfig command

Perl errors seen while using `haimfconfig` command:

```
Perl 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`.

Workaround: 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"
```

IMF does not provide notification for a registered disk group if it is imported using a different name [2730774]

If a disk group resource is registered with the AMF and the disk group is then imported using a different name, AMF does not recognize the renamed disk group and hence does not provide notification to DiskGroup agent. Therefore, the DiskGroup agent keeps reporting the disk group resource as offline.

Workaround: Make sure that while importing a disk group, the disk group name matches the the one registered with the AMF.

Direct execution of `linkamf` displays syntax error [2858163]

Bash cannot interpret Perl when executed directly.

Workaround: Run `linkamf` as follows:

```
# /opt/VRTSperl/bin/perl /opt/VRTSamf/imf/linkamf <destination-directory>
```

Error messages displayed during reboot cycles [2847950]

During some reboot cycles, the following message might get logged in the engine log:

```
AMF libvxamf ERROR V-292-2-149 Cannot unregister event: no rid -1 found  
AMF libvxamf ERROR V-292-2-306 Unable to unregister all events (errno:405)
```

This does not have any effect on the functionality of IMF.

Workaround: No workaround.

Error message displayed when ProPCV prevents a process from coming ONLINE to prevent concurrency violation does not have I18N support [2848011]

The following message is seen when ProPCV prevents a process from coming ONLINE to prevent concurrency violation. The message is displayed in English and does not have I18N support.

```
Concurrency Violation detected by VCS AMF.  
Process <process-details> will be prevented from startup.
```

Workaround: No Workaround.

Error message seen during system shutdown [2954309]

During some system shutdowns, you might see the following message in the syslog.

```
Stopping AMF...  
AMF amfconfig ERROR V-292-2-405 AMF_UNCONFIG failed, return value = -1
```

The system continues to proceed with the shutdown.

Workaround: No workaround.

AMF displays StartProgram name multiple times on the console without a VCS error code or logs [2872064]

When VCS AMF prevents a process from starting, it displays a message on the console and in syslog. The message contains the signature of the process that was prevented from starting. In some cases, this signature might not match the signature visible in the PS output. For example, the name of the shell script that was prevented from executing will be printed twice.

Workaround: No workaround.

Terminating the `imfd` daemon orphans the `vxnotify` process [2728787]

If you terminate `imfd` daemon using the `kill -9` command, the `vxnotify` process created by `imfd` does not exit automatically but gets orphaned. However, if you stop `imfd` daemon with the `amfconfig -D` command, the corresponding `vxnotify` process is terminated.

Workaround: The correct way to stop any daemon is to gracefully stop it with the appropriate command (which is `amfconfig -D` command in this case), or to terminate the daemon using Session-ID. Session-ID is the -PID (negative PID) of the daemon.

For example:

```
# kill -9 -27824
```

Stopping the daemon gracefully stops all the child processes spawned by the daemon. However, using `kill -9 pid` to terminate a daemon is not a recommended option to stop a daemon, and subsequently you must kill other child processes of the daemon manually.

Core dump observed when `amfconfig` is run with set and reset commands simultaneously [2871890]

When you run `amfconfig -S -R` on a node, a command core dump is observed, instead of displaying the correct usage of the command. However, this core dump has no effect on the AMF functionality on that node. You need to use the correct command syntax instead.

Workaround: Use the correct commands:

```
# amfconfig -S <options>
# amfconfig -R <options>
```

Issues related to the Cluster Manager (Java Console)

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

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

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

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

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

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

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

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

During online migration from LVM to VxVM volumes, LVM sometimes incorrectly reports the remapped LVM device paths as valid LVM volumes

Problem: In a migrated or committed configuration, only the renamed LVM names of the form <lvolname>_vxlv are valid LVM volumes. The original LVM names, in turn, point to target VxVM volumes. However, LVM sometimes incorrectly reports these original LVM device paths pointing to VxVM volumes, as valid LVM volumes.

Do not assume these as LVM volumes or do any operations on them, as it would disrupt the application's access to the target VxVM volumes.

Network File System (NFS) client reports I/O error because of network split brain (3257399)

When network split brain occurs, the failing node may take some time to panic. As a result, the service group on the failover node may fail to come online as some of the resources, such as IP resources, are still online on the failing node. The disk group on the failing node may also get disabled, but IP resources on the same node stays online. As the result, I/O error occurs.

Workaround:

Configure the pre-online trigger for the service groups containing DiskGroup resources with reservation on each system in the service group:

- 1 Copy the preonline_ipc trigger from

```
/opt/VRTSvcs/bin/sample_triggers/VRTSvcs to
```

```
/opt/VRTSvcs/bin/triggers/preonline/filename as T0preonline_ipc:
```

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

- 2 Enable T0preonline_ipc, the pre-online trigger for the service group:

```
# hagrps -modify group_name TriggersEnabled PREONLINE -sys node_name
```

Issues related to installing and upgrading VCS

Manual upgrade of VRTSvlic depot loses keyless product levels [2737124]

If you upgrade the VRTSvlic depot 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 depot.

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` depot

```
# swremove VRTSvlic
```

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

```
swinstall -s 'pwd' VRTSvlic
```

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

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

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

After upgrading from 5.0.1 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]
```

### **Installer does not detect the duplicate cluster ID in an already configured SF Oracle RAC cluster [2368898]**

When you run the installer using `installsfprac -configure` command and if you choose to check the cluster ID, the installer correctly checks if the cluster ID is in use by any other setup. However, if you perform the same check on an already configured SF Oracle RAC cluster, it is unable to detect it.

**Workaround:** No workaround.

### **VxSS may go to a faulted state After stack and OS upgrade [2564568]**

During the upgrade using the installer, if cluster is in secure mode and the upgrade is from 1123 to 1131, the installer may send a the following warning message.

```
Warning: /opt/VRTSat/bin/vxatd is not running
on <system name>. Will be unable to setup trust with shared
broker, however secure upgrade can still proceed.
```

**Workaround:** You can ignore this warning and proceed with the upgrade.

### **Manual install of VRTSvcs depot using /usr/sbin/swinstall -s 'pwd' VRTSvcs may fail [2399744]**

Manual installation of VRTSvcs depot using `/usr/sbin/swinstall -s `pwd` VRTSvcs` might fail on freshly installed HP-UX machine.

AVXFS is an HP-owned integration product and has dependency on VRTSvlic 3.02.24.0. This causes `/usr/sbin/swinstall` to not select VRTSvcs dependent products from the depot.

**Workaround:** The workaround for this issue is to remove the AONLINEJFS, OnlineJFS01, and AVXFS depots manually before stack installation. Moreover, before installing VCS on freshly installed HP-UX machine, uninstall older VRTSvlic 3.02.24.0 depot (if installed).

### **During rolling upgrade swverify command displays errors [2439492]**

While performing Rolling upgrade from VCS 5.1SP1 to VCS 6.0 or later, after phase 1 of the Rolling Upgrade process, `swverify` command throws the following errors for all VCS related man pages:

```
"/opt/VRTS/man/man1m/hastatus.1m" missing.
```

This does not cause any product functionality issues. The man pages reappear after phase 2 of the Rolling Upgrade is complete.

Workaround: Not required, as the man pages reappear after Rolling Upgrade phase 2. In you wish to retain all the man pages even after phase 1, copy the `/opt/VRTS/man` directory in a safe location before starting the Rolling Upgrade procedure.

### Errors seen during verification of VRTSamf depot (2599242)

If VRTSvcsea depot is manually upgraded using `swinstall` command from VCS 5.1SP1RP1 to VCS 6.0, permissions for `/opt/VRTSamf/imf/imf_register` file are affected. Due to this, errors are seen during verification of VRTSamf.

Workaround: First uninstall the VRTSvcsea depot from VCS 5.1SP1RP1 using `swremove` command and then proceed to install VRTSvcsea depot from VCS 6.0 release.

## VCS 5.0.1 Rolling Patch 1 known issues

The VCS issues in this release are as follows:

- The Oracle agent with 11g Release 2 does not support Health check monitoring using the MonitorOption attribute. If the database is 11g Release 2, the MonitorOption attribute for the Oracle agent should be set to 0. The Oracle agent with 11g Release 2 database does not support the Intentional Offline feature. [1975007]
- The ASMInst agent does not support pfile or spfile for the ASM Instance on the ASM diskgroups in 11g Release 2. Symantec recommends that you store the file on the local file system. [1975010]
- If you try to enable debug logs for the DB2 agent, the logs are not written to the `engine_A.log` file. [1954752]

Workaround: Download and install the GNU Awk software from the GNU Web site. Then, create a soft link to the default awk binary on the cluster nodes as follows:

```
ln -s /usr/local/bin/gawk /bin/awk
```

- The VRTSperl patch takes more than 10 minutes to install on an HP Integrity system node:  
On an HP Integrity system node, installing the VRTSperl patch takes more than 10 minutes and requires that VCS is offline during this period. The installation time may vary based on the configuration of the machine on which the VRTSperl patch is being installed.

## Veritas Storage Foundation for Oracle RAC known issues

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

- [Oracle RAC issues](#)
- [SFHA issues](#)

### Oracle RAC issues

This section lists the known issues in Oracle RAC.

#### **During installation or system startup, Oracle Grid Infrastructure may fail to start [1933542]**

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/npoohasd",
O_WRONLY <unfinished ...>
```

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

#### **Oracle VIP Configuration Assistant fails with an error message (1182220)**

During Oracle RAC 10g Release 2 installation, the VIP Configuration Assistant may fail with the following error message:

```
The given interface(s), "" is not public.
Public interfaces should be used to configure virtual IPs.
```

This message appears only when the VIP is not from the regular public IP range (for example, 200.).

Workaround: Invoke the `vipca` utility manually as the superuser.

```
export DISPLAY=nebula:0.0
$CRS_HOME/bin/vipca
```

#### **Oracle Cluster Verification utility displays a warning message**

During the final stage of Oracle RAC 10g Release 2 installation, you may receive a warning message with the Oracle Cluster Verification utility.

For example:

Utility

```
=====
OUI-25031: Some of the configuration assistants failed. It is
strongly recommended that you retry the configuration
assistants at this time. Not successfully running any "
Recommended" assistants means your system will not be correctly
configured.
```

1. Check the Details panel on the Configuration Assistant Screen to see the errors resulting in the failures.
2. Fix the errors causing these failures.
3. Select the failed assistants and click the 'Retry' button to retry them.

```
=====
```

**Workaround:** You may safely ignore this message if the cluster is operating satisfactorily.

## SFHA issues

This section lists the known issues in SFHA for this release.

### Installation known issues

This section describes the known issues during installation and upgrade.

#### **PrivNIC and MultiPrivNIC agents not supported with Oracle RAC 11.2.0.2 and later versions**

The PrivNIC and MultiPrivNIC agents are not supported with Oracle RAC 11.2.0.2 and later versions.

For more information, see the following Technote:

<http://www.symantec.com/business/support/index?page=content&id=TECH145261>

#### **File system check daemon fails to restart after abnormal termination (2689195)**

The file system check daemon (`vxfsckd`) fails to update the `vxfsckd-pid` file with the new process ID (pid) of the `vxfsckd` process after abnormal termination. As a result, the CFSfsckd agent fails to detect the status of the `vxfsckd` daemon.

**Workaround:** Perform the following steps to resolve the issue on the node where the `vxfsckd` resource faults:

1. Log into the node as the root user.

2. Kill all `vxfscsd` processes:

```
kill -9 `ps -ef|grep vxfscsd|awk '{print $2}'`
```

3. Remove the `vxfscsd-pid` file:

```
rm /var/adm/cfs/vxfscsd-pid
```

4. Bring the `vxfscsd` resource online:

```
hares -online vxfscsd_resname -sys node_name
```

### Startup or shutdown failure messages reported for LLT, GAB, VXFEN, and VCSMM (1666327) Startup or shutdown failure messages reported for LLT, GAB, and VXFEN (1666327)

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

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

### 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.

### 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.

### On nodes with heavy load, the CSSD resource may fault [3404403]

The CSSD agent checks the status of Oracle Clusterware using the Oracle Clusterware command `crsctl check crs`. On nodes with heavy load, the command does not complete within the period that the `MonitorTimeout` defines. After the 4 (default value of the `FaultOnMonitorTimeout` attribute) successive monitor timeouts, the CSSD resource goes to the FAULT state.

**Workaround:** Set the value of the `FaultOnMonitorTimeouts` attribute to 0 and use the `AlertOnMonitorTimeouts` attribute.

- 1 Change the permission on the VCS configuration file to read-write mode. Enter:

```
haconf -makerw
```

- 2 Set the `AlertOnMonitorTimeouts` attribute value to 4 for the CSSD resource. Enter:

```
hatype -display Application | grep AlertOnMonitorTimeouts
Application AlertOnMonitorTimeouts 0
```

```
hares -override cssd_resname AlertOnMonitorTimeouts
hatype -modify Application AlertOnMonitorTimeouts 4
```

- 3 Set the `FaultOnMonitorTimeouts` attribute value to 0 for the CSSD resource. Enter:

```
hatype -display Application | grep FaultOnMonitorTimeouts
Application FaultOnMonitorTimeouts 4
hares -override cssd_resname FaultOnMonitorTimeouts
hatype -modify Application FaultOnMonitorTimeouts 0
```

- 4 Verify the `AlertOnMonitorTimeouts` and `FaultOnMonitorTimeouts` settings. Enter:

```
hatype -display Application |
egrep "AlertOnMonitorTimeouts|FaultOnMonitorTimeouts"
Application AlertOnMonitorTimeouts 4
Application FaultOnMonitorTimeouts 0
```

- 5 Change the permission on the VCS configuration file to read-only mode. Enter:

```
haconf -dump -makero
```

## Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) known issues

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

### Internal stress test on cluster file system hits debug assert in GLM [3364309]

In Group Lock Manager (GLM), the code to handle last revoke for a lock may cause deadlock, which is caught upfront by debug assert.

**Workaround:** There's no workaround for the issue.

### The mount command may hang when there are large number of inodes with extops and a small `vxfs_ninode`, or a full `fsck` cannot fix the link count table corruptions (2689326)

You might encounter one of the following issues:

- If there are large number of inodes having extended operations (`extops`), then the number of inodes used by the `mount` command reaches the maximum number of inodes that can be created in `core`. As a result, the `mount` command will not get any new inodes, which causes the `mount` command to run slowly and sometimes hang.

**Workaround:** Increase the value of `vxfs_ninode`.

- The link count table (LCT) file can get damaged such that the flag is set, but the attribute inode is already freed. In this case, the `mount` command tries to free an inode that has been already freed thereby marking the file system for a full structural file system check.

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

## An ENOSPC error may return to the cluster file system application (2867282)

In some cases, when a large number of exclusion zones are set by commands such as `fsadm`, an ENOSPC error may return to the cluster file system application when delegations with free extents are not available.

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

## CFS commands might hang when run by non-root (2403263, 3038283)

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.

## 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.

## Panic due to null pointer de-reference in `vx_bmap_lookup()` (2582232, 3038285)

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.

## MountAgent process can get stuck from repeatedly switching a service group from one node to another (2170318)

The `MountAgent` process can get stuck from repeatedly switching a service group from one node to another. This occurs because the `MountAgent` process is waiting for notification, but the notification is unregistered.

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

## The `svsdatastore(1M)` command may set the return value to zero even in cases of error. (3313498)

The `svsdatastore(1M)` command may set the return value to zero even in cases of error.

For example:

```
#svsdatastore add invalid disk name

Error: V-35-585: Disk invaliddisk does not exists

echo $?

0
```

### Workaround:

There is no workaround for this issue.

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

This section describes Veritas Storage Foundation for Databases (SFDB) tools known issues in this release.

### Relinking ODM after upgrading from 5.0.x

The `VRTSodm` library path has changed from `/opt/VRTSodm/lib/libodm.sl` to `/opt/VRTSodm/lib/libodm.so`.

After upgrading to from 5.0.x you must update the ODM link for your database to the new `VRTSodm` library path `/opt/VRTSodm/lib/libodm.so`.

### SFDB commands do not work in IPV6 environment (2619958)

In IPV6 environment, SFDB commands do not work for SFHA. There is no workaround at this point of time.

### 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.

## 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.

## 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.

## 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 of 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 5.0MP2 to 6.0.5 (2003131)

While upgrading from 5.0 MP2 to 6.0.5 the following error message could be seen when running `sfua_rept_migrate`:

```
/opt/VRTSdbed/migrate/sfua_rept_migrate
Mounting SFUA Sybase ASA repository.
SFORA sfua_rept_migrate ERROR V-81-8903 Could not start repository
database.
/usr/lib/dld.sl: Can't find path for shared library: libcur_colr.1
/usr/lib/dld.sl: No such file or directory
sh: 3845 Abort(coredump)
Symantec DBMS 3.0.85.0 vxdbms_start_db utility
ASA failed. Sybase ASA error code: [134].
Sybase ASA Error text: {{{}}}
```

```
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

## Workaround

### To upgrade without an existing SFDB repository set up

- 1 Verify X/Open curses is installed on the system.
- 2 Create the following link: `ln -s /usr/lib/libxcurses.1 /usr/lib/libcur_colr.1`
- 3 Run:

```
sfua_rept_migrate
```

## Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0.5 (2184482)

When upgrading from SFHA version 5.0 or 5.0.1 to SFHA 6.0.5 the S\*vxdms3 startup script is renamed to NO\_S\*vxdms3. The S\*vxdms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S\*vxdms3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdms3 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\*vxdms3 to S\*vxdms3.

## Upgrading in an HP Serviceguard environment (2116452)

When upgrading SFDB to 5.1SP1 from the previous release in an HP Serviceguard environment, first verify that the `cmviewcl` command can be executed by a non-root user. This permission change must be done before executing SFDB upgrade commands.

## Clone command fails if PFILE entries have their values spread across multiple lines (2844247)

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`.

## Frequent occurrence of SFDB remote or privileged command error (2869262)

If you installed a single instance database and try to run SFDB-related commands, then an error similar to the following might occur:

```
$ /opt/VRTSdbed/bin/dbed_update
```

```
No repository found for database faildb, creating new one.
```

```
SFDB vxsfadm ERROR V-81-0450 A remote or privileged command could not be executed on host1
```

Reason: This can be caused by the host being unreachable or the vxdbd daemon not running on that host.

Action: Verify that the host swpa04 is reachable. If it is, verify that the vxdbd daemon is running using the `/opt/VRTS/bin/vxdbdctrl status` command, and start it using the `/opt/VRTS/bin/vxdbdctrl start` command if it is not running.

There is no workaround at this point of time.

## Data population fails after datafile corruption, rollback, and restore of offline Storage Checkpoint (2869259)

Sometimes when a datafile gets corrupted below its reservation size, the rollback may not pass and the file may not be rolled back correctly.

There is no workaround at this point of time.

## Offline mode Storage Checkpoint or FlashSnap does not confirm the offline status of the database in CFS environment, leading to clone failure (2869260)

In a cluster file system for Single Instance Oracle, if an offline snapshot or Storage Checkpoint, and clone is created on the node where the database is inactive, then the cloning would fail with an error similar to SFDB vxsfadm ERROR V-81-0564 Oracle returned error.

```
... Reason: ORA-01194: file 1 needs more recovery to be consistent
ORA-01110: data file 1: /var/tmp/ikWxDkQ1Fe/data/sfaedb/system01.dbf'
(DBD ERROR: OCISmtExecute) ...
```

**Workaround:** There is no workaround for this. In case of a Single Instance database installed on a cluster file system, create the Storage Checkpoint or snapshot on the active node.

## Storage Checkpoint clone fails if the `archive log` destination is same as the datafiles destination (2869266)

Storage Checkpoint cloning fails if the `archive log` destination is the same as the datafiles destination. The error is similar to:

```
Use of uninitialized value $path in hash element
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 121.
Use of uninitialized value $path in concatenation (.) or string
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 124.
Use of uninitialized value $path in pattern match (m//)
at /opt/VRTSdbed/lib/perl/DBED/CkptOracle.pm line 126.
```

```
SFDB vxsfadm ERROR V-81-0564 Oracle returned error.
```

```
Reason: ORA-02236: invalid file name (DBD ERROR: error possibly near
<*> indicator at char 172 in 'CREATE CONTROLFILE REUSE SET DATABASE
'TClone03' RESETLOGS NOARCHIVELOG
```

**Workaround:** For the 6.0.5 release, create distinct archive and datafile mounts for the Storage Checkpoint service.

## FileSnap detail listing does not display the details of a particular snap (2846382)

FileSnap does not support displaying a detailed listing of a snapshot or clone. FileSnap only supports displaying a summary of all the snapshots or clones. For example, for the CLI `vxsfadm -s filesnap -a oracle --name=snap1 -o list`, a summary listing all the snapshots is displayed, instead of a detailed listing of a particular snapshot.

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

## Flashsnap clone fails under some unusual archivelog configuration on RAC (2846399)

In a RAC environment, when using FlashSnap, the archive log destination to snapshot must be a shared path, and must be the same across all the nodes. Additionally, all nodes must use the same archive log configuration parameter to specify the archive log destination. Configurations similar to the following are not supported:

```
tpcc1.log_archive_dest_1='location=/tpcc_arch'
tpcc2.log_archive_dest_2='location=/tpcc_arch'
tpcc3.log_archive_dest_3='location=/tpcc_arch'
```

Where tpcc1, tpcc2, and tpcc3 are the names of the RAC instances and /tpcc\_arch is the shared archive log destination.

**Workaround:** To use FlashSnap, modify the above configuration to \*.log\_archive\_dest\_1='location=/tpcc\_arch'. For example,

```
tpcc1.log_archive_dest_1='location=/tpcc_arch'
tpcc2.log_archive_dest_1='location=/tpcc_arch'
tpcc3.log_archive_dest_1='location=/tpcc_arch'
```

## Swverify error related to VRTSdbed observed after a Phase 2 rolling upgrade of SFRAC 6.0.1 on HP-UX 11.31 (2869263)

Upgrade of the SF or SFRAC stack from 5.x to 6.0.1 could display an swverify warning, as follows:

```
WARNING: Directory "/var/vx/vxdba/locks" should have mode "755" but the
actual mode is "1755".
WARNING: Directory "/var/vx/vxdba/logs" should have mode "755" but the
actual mode is "1755".
WARNING: Fileset "VRTSdbed.DBED,1=/,r=6.0.100.000" had file warnings.
```

**Workaround:** Ignore the warning, or change the directory permissions to 755 for both /var/vx/vxdba/locks and /var/vx/vxdba/logs.

## Storage Checkpoint clone fails in CFS environment if cloned using same Storage Checkpoint and same clone name on both nodes (2869268)

The Storage Checkpoint clone of an oracle database fails in a CFS environment, if you create a clone with a clone name and Storage Checkpoint name same as another clone up on a different CFS node.

### **Workaround:**

There is no workaround. Create a clone with a different clone name.

## Very long off-host cloning times for large number of datafiles (2849540)

When cloning off-host in certain Oracle database configurations, particularly with several hundred datafiles, the cloning can take a very long time, upto an hour or

more. This problem does not cause the cloning to fail. The problem applies to all services such as FlashSnap, Space-optimized snapshots, FileSnap, and Storage Checkpoint.

**Workaround:**

There is no workaround at this point of time.

**Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0.5 (2184482)**

When upgrading from SFHA version 5.0 or 5.0.1 to SFHA 6.0.5 the S\*vxdms3 startup script is renamed to NO\_S\*vxdms3. The S\*vxdms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S\*vxdms3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdms3 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\*vxdms3 to S\*vxdms3.

**sfua\_rept\_migrate fails after phased SFRAC upgrade from 5.0MP3RP5 to 6.0.1 (2874322)**

Command `sfua_rept_migrate` sometimes gives an error when upgrading to 6.0.1, and fails to unmount the repository volume. The error message is similar to:

```
./sfua_rept_migrate
Mounting SFUA Sybase ASA repository.
Unmounting SFUA Sybase ASA repository.
UX:vxfs umount: ERROR: V-3-26388: file system /rep has been mount
locked
SFORA sfua_rept_migrate ERROR V-81-5550 umount /dev/vx/dsk/repdg/repvol
failed.
SFORA sfua_rept_migrate ERROR V-81-9162 Failed to umount repository.
```

**Workaround:**

The error does not hamper the upgrade. The repository migration works fine, but the old repository volume does not get unmounted. Unmount the mount using the manual option.

For example, use `/opt/VRTS/bin/umount -o mntunlock=VCS /rep`.

For more information, see [TECH64812](#).

## Some dbed operations may fail in system configurations where the hostname “localhost” cannot be resolved [3436609]

With hostname “localhost” that fails to get resolved, many dbed operations may fail. For example, the “vxsfadm -o valid” operation fails with the following error messages:

```
bash-4.1$ /opt/VRTSdbed/bin/vxsfadm -s sos -a oracle -o valid -c \
/tmp/sn7130
Use of uninitialized value in concatenation (.) or string
at /opt/VRTSdbed/lib/perl/DBED/SfaeFsm.pm line 2119.
Use of uninitialized value in string at \
/opt/VRTSdbed/lib/perl/DBED/SfaeFsm.pm line 2120.
```

```
SFDB vxsfadm ERROR V-81-0728 The directory \
/etc/vx/vxdba/oracle/local/.sfae could not be created.
```

```
Reason: Operating system returned error: No such file or directory
```

**Workaround:** Ensure that the name "localhost" resolves to the local loopback interface address (e.g. 127.0.0.1). You can verify whether "localhost" name can be resolved on your host by using the `ping` command.

Example output on a system where "localhost" cannot be resolved:

```
bash-4.1# ping localhost
ping: unknown host localhost
```

# Software limitations

This section covers the software limitations of this release.

## Limitations related to installation

This is the limitations related to installation in the 6.0.5 release.

### Limitations related to web-based installer for SFRAC

- Web-based installer on local disk is not supported.
- If SFRAC is not configured before upgrade, the web-based installer does not support to upgrade SFRAC to 6.0.5.

## Limitations related to Install Bundles

- Web-based installer doesn't support the Install Bundles feature.
- The feature doesn't support native OS install or upgrade methods, such as ignite.
- The Install Bundles feature for 6.0.5 does not support hot fix installation.

## 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).

### Support for SmartSync with database mounted on raw volumes [3416016]

The SmartSync feature with the database configured on raw volumes depends on support from the database vendor. If supported by the database vendor, the SmartSync feature uses an extended interface between VxVM volumes and the database software to avoid unnecessary work during mirror resynchronization.

Verify with the database vendor that the database software supports the SmartSync feature.

The SmartSync feature is supported on all platforms when the database uses VxFS file systems mounted on Veritas Volume Manager volumes, through the Veritas Extension for Oracle Disk Manager (VRTSodm) interface.