

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

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 5

Legal Notice

Copyright © 2015 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.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

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 | 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 AIX

| BFF file | Size in bytes | Patches | Version |
|------------------|---------------|--------------|-----------|
| VRTSamf.bff | 12134400 | VRTSamf | 6.0.500.0 |
| VRTScavf.bff | 972800 | VRTScavf | 6.0.500.0 |
| VRTSdbac.bff | 5068800 | VRTSdbac | 6.0.500.0 |
| VRTSdbed.bff | 126771200 | VRTSdbed | 6.0.500.0 |
| VRTSfsadv.bff | 32921600 | VRTSfsadv | 6.0.500.0 |
| VRTSgab.bff | 3788800 | VRTSgab | 6.0.500.0 |
| VRTSglm.bff | 460800 | VRTSglm | 6.0.500.0 |
| VRTSilt.bff | 2508800 | VRTSilt | 6.0.500.0 |
| VRTSodm.bff | 1484800 | VRTSodm | 6.0.500.0 |
| VRTSperl.bff | 69990400 | VRTSperl | 5.14.2.20 |
| VRTSsfpci601.bff | 5120000 | VRTSsfpci601 | 6.0.500.0 |
| VRTSvcsc.bff | 361216000 | VRTSvcsc | 6.0.500.0 |
| VRTSvcscsag.bff | 19712000 | VRTSvcscsag | 6.0.500.0 |
| VRTSvcsea.bff | 1740800 | VRTSvcsea | 6.0.500.0 |
| VRTSvxfen.bff | 3174400 | VRTSvxfen | 6.0.500.0 |
| VRTSvxfsc.bff | 42956800 | VRTSvxfsc | 6.0.500.0 |
| VRTSvxvm.bff | 284620800 | VRTSvxvm | 6.0.500.0 |
| VRTSaslapm.bff | 1024000 | VRTSaslapm | 6.0.500.0 |

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

There are no changes related to Storage Foundation and High Availability in 6.0.5.

Note: For Storage Foundation and High Availability 6.0.5 to support HP IVM 6.3.5 new feature, ensure that both the VirtualBase on VSP and guest is 6.3.5.

Storage Foundation and High Availability 6.0.5 supports the new 16GB Fibre Channel **HP SN1000Q** HBA card in physical and virtual environments (HP IVM 6.3.5 through NPIV).

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` fileset upgrade and rolling upgrade and back

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

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

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 DB2 version 10.5.

Note: When updating DB2 database parameters, ensure that any deferred values are updated before using SFDB tools. The DB2 database needs to be reactivated for the parameter values to take effect.

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 AIX operating systems

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

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

Table 1-2 Supported operating systems

| Operating systems | Levels | Chipsets |
|-------------------|-----------------------------|------------------------------|
| AIX 7.1 | TL0, TL1, TL2, and TL3 | Power 5, Power 6 and Power 7 |
| AIX 6.1 | TL5, TL6, TL7, TL8, and TL9 | Power 5, Power 6 and Power 7 |

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

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

AIX 7.1 support for virtual processors

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

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 for Database features supported in database environments

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

Table 1-3 SFDB features database support for 6.0.5

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

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

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

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

Veritas Storage Foundation memory requirements

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

Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

Fixed issues

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

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 1-4 Installation and upgrades fixed issues in 6.0.5

| Fixed issues | Description |
|--------------|--|
| 3068058 | The VxFEN module and VCS fails to start after a rollback operation in the absence of VCS 6.0.1 version. |
| 3074644 | [VCS 6.0.3 - LBN - Fencing/HAD will not start after rollback from 6.0.3 to 6.0.1 if CPS based fencing is configured] |
| 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". |
| 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 1-5 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 1-6 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 |
| 2627076 | Incorrect server names sometimes display if there is a clock synchronization issue. |
| 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. |
| 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 1-7 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. |
| 3152313 | With Partitioned Directories feature enabled, removing a file may panic the system. |
| 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. |

Table 1-7 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 3022673 | Veritas File System (VxFS) is unresponsive when it changes the memory using DLPAR. |
| 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. |
| 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. |
| 3003679 | When running the fsppadm(1M) command and removing a file with the named stream attributes (nattr) at the same time, the file system does not respond. |
| 3072036 | Read operations from secondary node in CFS can sometimes fail with the ENXIO error code. |
| 3364312 | The fsadm(1M) command is unresponsive while processing the VX_FSADM_REORGLK_MSG message. |
| 3359200 | Internal test on Veritas File System (VxFS) fsdedup(1M) feature in cluster file system environment results in a hang. |
| 2735912 | The performance of tier relocation using the fsppadm(1M)enforce command degrades while migrating a large number of files. |
| 3153919 | The fsadm (1M) command may hang when the structural file set re-organization is in progress. |
| 3332902 | While shutting down, the system running the fsclustadm(1M)command panics. |
| 3046983 | Invalid CFS node number in ".__fsppadm_fclextract", causes the DST policy enforcement failure. |

Table 1-7 Veritas Storage Foundation Cluster File System High Availability 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 3364287 | Debug assert may be hit in the vx_real_unshare() function in the cluster environment. |
| 3364301 | Assert failure because of improper handling of inode lock while truncating a reorg inode. |
| 3023855 | Enable the noreuserd option of the mount(1M) command on CFS. |
| 3444775 | Internal noise testing on cluster file system results in a kernel panic in function vx_fsadm_query() with an error message. |
| 3192985 | Checkpoints quota usage on Cluster File System (CFS) can be negative. |
| 3364309 | Internal stress test on cluster file system hit debug assert in Group Lock Manager (GLM). |
| 3410837 | The error message has an issue when the user uses the cfsumount(1M) 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 fcntl lock. |

Veritas Storage Foundation Cluster File System High Availability: issues fixed in 6.0.3

[Table 1-8](#) describes the Veritas Storage Foundation Cluster File System fixed issues in 6.0.3.

Table 1-8 Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues

| Incident | Description |
|----------|---|
| 2977697 | vx_idetach generated kernel core dump while filestore replication is running. |
| 2942776 | Mount fails when volumes in vset are not ready. |
| 2923867 | Internal test hits an assert "f:xted_set_msg_pri:1". |
| 2923105 | The upgrade VRTSvxfs5.0MP4HFaf hangs at vxfs preinstall scripts. |
| 2916691 | Customer experiencing hangs when doing dedups. |

Table 1-8 Veritas Storage Foundation Cluster File System High Availability 6.0.3 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 2906018 | The vx_iread errors are displayed after successful log replay and mount of the file system. |
| 2857731 | Internal testing hits an assert "f:vx_mapdeinit:1" . |
| 2843635 | Internal testing is having some failures. |
| 2841059 | full fsck fails to clear the corruption in attribute in ode 15. |
| 2750860 | Performance issue due to CFS fragmentation in CFS cluster. |
| 2715175 | It takes 30 minutes to shut down a 4-node cluster. |

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 1-9 Veritas Storage Foundation Cluster File System High Availability fixed issues

| Incident | Description |
|----------|---|
| 2867282 | An ENOSPC error may return to the cluster file system application. |
| 2703747 | CFS failover takes up to 20 minutes due to slow log replay. |
| 2684573 | The performance of the cfsmount(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 1-10 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. |
| 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 CLIfails 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. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 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. |
| 2874810 | When you install DMP only solutions using the installdmp command, the root support is not enabled. |
| 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 departed disk group. |
| 2884122 | During the server boot, multiple wrong messages are logged into server console in Virtual I/O Server (VIOS) environment. |
| 2893530 | With no VVR configuration, when system is rebooted, it panicked. |
| 2898324 | UMR errors reported by Purify tool in "vradmind migrate" command. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 2903801 | Unnecessary messages are displayed in the syslog. |
| 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. |
| 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. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 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 reminor 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. |
| 2968845 | When the Fibre Channel devices are dynamically tracked, change in the SCSI-ID of the device results in loss of access to the device. |
| 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. |
| 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. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 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 ownername 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 vxdg 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. |
| 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. |
| 3020015 | With operating system naming scheme, the procedure of putting root disk under DMP control doesn't work properly. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 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. |
| 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 vradm 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 vradm repstatus(1M) shows "dcm contains 0 kbytes" during the Smart Autosync. |
| 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 "vxdkmpadm 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. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues *(continued)*

| Fixed issues | Description |
|--------------|--|
| 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. |
| 3102114 | A system crash during the 'vxsnap restore' operation can cause the vxconfigd(1M) daemon to dump core after the system reboots. |
| 3106103 | When a volume group contains a logical volume with the samename as that of any Veritas Volume Manager (VxVM) object (volumes, plexes, subdisks, etc.) in other existing disk groups, the vxconvert(1M) command on that volume group fails. |
| 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. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 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. |
| 3140211 | AIX command /etc/vx/bin/vxisforeign does not recognize LVM disks. |
| 3142315 | Disk is misidentified as clone disk with udid_mismatch flag. |
| 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 lmissing disks) go into the "ONLINE INVALID LFAILED" or"ONLINE INVALID LMISSING" state after the disk loses global disk connectivity. |
| 3147974 | Memory leak observed from VxDMP driver. |
| 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. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 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. |
| 3178267 | Stale vx* directories are not removed from /tmp. |
| 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. |
| 3211351 | DMP does not claim the not-ready devices that are controlled by EMC PowerPath(PP). |
| 3222707 | Dynamic Reconfiguration (DR) tool does not permit the removal of disks associated with a deported diskgroup(dg). |
| 3225660 | The Dynamic Reconfiguration (DR) tool does not list thin provisioned LUNs during a LUN removal operation. |
| 3231863 | The vxconvert(1M) command fails, as the vxmake(1M) command fails to create VxVM objects during conversion. |
| 3238397 | Dynamic Reconfiguration (DR) Tool's Remove LUNs option does not restart the vxattachd daemon. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 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. |
| 3244217 | Cannot reset the clone_disk flag during vxdg import. |
| 3247040 | vxdisk scandisks enables the PowerPath (PP) enclosure which was disabled previously. |
| 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. |
| 3263095 | The following message is displayed many times on AIX:NOTE VxVM vxdmp V-5-3-0 dmp_indirect_ioctl: loctl Failed for [major]/[minor]with error 22 |
| 3271595 | Veritas Volume Manager (VxVM) should prevent the disk reclaim flag from getting turned off, when there are pending reclaims on the disk. |
| 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. |
| 3279588 | System panics in uninstalling VRTSvxvm with VVR configured. |
| 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. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 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. |
| 3298401 | Running the vxdisk scandisks command causes the vxconfigd memory leak. |
| 3323548 | In the Cluster Volume Replicator (CVR) environment, a cluster-wide vxconfigd hang occurs on primary when you start the cache object. |
| 3367997 | System panics during LUN removal operations. |
| 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. |
| 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). |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 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. |
| 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. |
| 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). |
| 3205490 | OS hangs at bootup time when boot device is shared across multiple nodes. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 3199398 | Output of the command "vxddmpadm 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. |
| 3185199 | DMP fails to restore mksysb when target disk has multiple paths. |
| 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. |
| 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. |
| 3117209 | After enabling the native support or DMP root support, the CLI vxddmpadm native list does not show rootvg configuration. The lspv output for rootvg shows multiple devices with the same Port VLAN ID (PVID). |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 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. |
| 2921147 | udid_mismatch flag is absent on a clone disk when source disk is unavailable. |
| 2859470 | The Symmetrix Remote Data Facility R2 (SRDF-R2) with the Extensible Firmware Interface (EFI) label is not recognized by Veritas Volume Manager (VxVM) and goes in an error state. |
| 2857044 | System crash on voldco_getalloffset when trying to resize filesystem. |

Table 1-10 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. |
| 2567618 | The VRTSexplorer dumps core in vxcheckbaapi/print_target_map_entry. |
| 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. |
| 3423316 | The vxconfigd(1M) daemon observes a core dump while executing the vxdisk(1M) scandisks command. |
| 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. |

Table 1-10 Veritas Volume Manager 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 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. |
| 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. |
| 3376725 | EMC VNX and VPLEX devices managed by PowerPath(PP) are not shown properly in the vxdisk list output. |
| 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 vxdmp (Veritas Dynamic Multi Pathing) error messages flood syslog.</p> <p>B. DMP metanode/path under DMP metanode gets disabled unexpectedly.</p> |
| 3331765 | With the Dynamic Multi-Pathing (DMP) support enabled for the Logical Volume Manager (LVM) rootvg, the mkysyb restore using Network Installation Manager (NIM) fails when the target disk chosen for restore has multiple paths. |

Veritas Volume Manager: issues fixed in 6.0.3

[Table 1-11](#) describes the incidents that are fixed in Veritas Volume Manager in 6.0.3.

Table 1-11 Veritas Volume Manager 6.0.3 fixed issues

| Incident | Description |
|----------|--|
| 3002770 | Accessing NULL pointer in <code>dmp_aa_recv_inquiry()</code> caused system panic. |
| 2970368 | Enhancing handling of SRDF-R2 WD devices in DMP. |
| 2965910 | vxassist dump core with the <code>-o ordered</code> option. |
| 2962262 | Uninstallation of DMP fails in presence of other multi-pathing solutions. |
| 2948172 | Executing the <code>vxdisk -o thin,fssize list</code> command can result in panic. |
| 2942609 | Message displayed when user quits from Dynamic Reconfiguration Operations is shown as error message. |
| 2940446 | Full fsck hangs on I/O in VxVM when cache object size is very large |
| 2935771 | In the VVR environment, RLINK disconnects after the master is switched. |
| 2934729 | VM is claiming disks as 'online' in VIOS. |
| 2933138 | panic in <code>voldco_update_itemq_chunk()</code> due to accessing invalid buffer |
| 2930569 | The LUNs in 'error' state in output of 'vxdisk list' cannot be removed through DR(Dynamic Reconfiguration) Tool. |
| 2919720 | vxconfigd core in <code>rec_lock1_5()</code> |
| 2919714 | exit code from <code>vxevac</code> is zero when migrating on thin luns but FS is not mounted |
| 2919627 | Dynamic Reconfiguration tool should be enhanced to remove LUNs feasibly in bulk. |
| 2919318 | The I/O fencing key value of data disk are different and abnormal in a VCS cluster with I/O fencing. |
| 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. |
| 2915063 | Rebooting VIS array having mirror volumes, master node panicked and other nodes CVM FAULTED |
| 2911040 | Restore from a cascaded snapshot when its source is DETACHED leaves the volume in unusable state |
| 2910043 | Avoid order 8 allocation by <code>vxconfigd</code> in node reconfig. |

Table 1-11 Veritas Volume Manager 6.0.3 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 2899173 | vxconfigd hang after executing the <code>vradmind stopprep</code> command. |
| 2898547 | vradmind on VVR Secondary Site dumps core, when Logowner Service Group on VVR (Veritas Volume Replicator) Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes. |
| 2892983 | vxvol dumps core if new links are added while the operation is in progress. |
| 2886402 | vxconfigd hang while executing <code>tc ./scripts/ddl/dmpapm.tc#11</code> |
| 2886333 | The <code>vxdbg(1M) join</code> command should not allow mixing clone and non-clone disks in a DiskGroup |
| 2884225 | vxconvert command fails to convert 1.5TB AIX LVM diskgroup to vxvm diskgroup |
| 2882908 | Machine failed to bootup with error "PReP-BOOT : Unable to load full PReP image" |
| 2879248 | vxdisk scandisks gets hung on VIO client with <code>dmp_native_support</code> enabled |
| 2878876 | vxconfigd dumps core in <code>vol_cbr_dolog()</code> due to race between two threads processing requests from the same client. |
| 2869594 | Master node panics due to corruption if space optimized snapshots are refreshed and " <code>vxclustadm setmaster</code> " is used to select master. |
| 2866059 | Improving error messages hit during the <code>vxdisk resize</code> operation. |
| 2859470 | SRDF R2 with EFI label is not recognized by VxVM and showing in error state |
| 2858853 | vxconfigd coredumps in <code>dbf_fmt_tbl</code> on the slave node after a Master Switch if you try to remove a disk from the DG |
| 2851403 | The <code>vxportal</code> and <code>vxfs</code> processes are failed to stop during first phase of rolling upgrade. |
| 2851085 | DMP doesn't detect implicit LUN ownership changes for some of the <code>dmpnodes</code> . |
| 2837717 | The <code>vxdisk(1M) resize</code> command fails if <code>da name</code> is specified. |
| 2836798 | Prevent DLE on simple/sliced disk with EFI label |
| 2834046 | NFS migration failed due to device reminoring. |

Table 1-11 Veritas Volume Manager 6.0.3 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 2833498 | vxconfigd hangs while reclaim operation is in progress on volumes having instant snapshots |
| 2826125 | VxVM script daemon is terminated abnormally when it is invoking with exact the same process id of the last invocation. |
| 2815517 | vxdg adddisk should not allow mixing clone and non-clone disks in a DiskGroup |
| 2801962 | Grow of a volume takes significantly large time when the volume has version 20 DCO (Data Change Object) attached to it |
| 2798673 | System panics in voldco_alloc_layout() while creating volumes with instant DCO. |
| 2779580 | Secondary node gives configuration error (no Primary RVG) after reboot of master node on Primary site. |
| 2753954 | At cable disconnect on port1 of dual-port FC HBA, paths via port2 are also marked SUSPECT |
| 2744004 | vxconfigd is hung on the VVR secondary node during VVR configuration. |
| 2715129 | vxconfigd hangs during Master takeover in a CVM (Clustered Volume Manager) environment. |
| 2692012 | vxevac move error message needs to be enhanced to be less generic and give clear message for failure. |
| 2619600 | Live migration of virtual machine having SFHA/SFCFSHA stack with data disks fencing enabled, causes service groups configured on virtual machine to fault. |
| 2567618 | VRTSexplorer coredumps in vxcheckhbaapi/print_target_map_entry |
| 2510928 | Extended attributes for SRDF luns reported as Mirror with EMC (VMAX array) |
| 2482308 | Devices go into error state after unmanaging them from Powerpath |
| 2398416 | vxassist dumps core while creating volume after adding attribute "wantmirror=ctlr" in default vxassist rulefile |
| 2273190 | Incorrect setting of the UNDISCOVERED flag can lead to database inconsistency. |
| 2149922 | Record the diskgroup import and deport events in syslog |

Table 1-11 Veritas Volume Manager 6.0.3 fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 2000585 | The <code>vxrecover -s</code> command does not start any volumes if a volume is removed whilst it is running. |
| 1982965 | <code>vxvg import DG</code> fails if <code>da-name</code> is based on naming scheme which is different from the prevailing naming scheme on the host. |
| 1973983 | <code>vxunreloc</code> fails when <code>dco plex</code> is in DISABLED state. |
| 1903700 | <code>vxassist remove mirror</code> does not work if <code>nmirror</code> and <code>alloc</code> is specified on VxVM 3.5 |
| 1901838 | Incorrect setting of <code>Nolicense</code> flag can lead to <code>dmp</code> database inconsistency. |
| 1859018 | The <code>link detached from volume</code> warnings are displayed when a linked-breakoff snapshot is created. |
| 1765916 | VxVM socket files don't have proper write protection |
| 1725593 | The <code>vxndmpadm listctlr</code> command has to be enhanced to print the count of device paths seen through the controller. |
| 1289985 | <code>vxconfigd</code> core dumps upon running the <code>vxctl enable</code> command. |

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 and Cluster Volume Manager fixed issues.

Table 1-12 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>vxndmpadm list dmpnode dmpnodename=dmpnode name</code> . |
| 2818840 | Enhance the <code>vxndmpraw</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. |
| 2792242 | I/O hang after performing zone remove/add operations. |

Table 1-12 Veritas Volume Manager fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 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_drl_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. |
| 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 sourcedg. |
| 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. |
| 2727590 | The <code>vxconfigd</code> daemon dumped core after renaming iSCSI device from the OS side. |
| 2713166 | Suppress <code>vxio errpt</code> messages relating to write-disabled (WD) & NR devices (EMC SRDF-R2) & EMC BCV. |
| 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. |
| 2709767 | Thin provisioning reclaim should not work with MPIO. |
| 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>vradmind</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). |

Table 1-12 Veritas Volume Manager fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 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. |
| 2675538 | The <code>vxdisk resize</code> command may cause data corruption. |
| 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>vxdmppadm list dmpnode</code> printing incorrect path type. |
| 2620555 | I/O hang due to SRL overflow & CVM reconfig. |
| 2612960 | The <code>vxconfigd</code> daemon core dumps after upgrading, due to GPT/AIX label disk. |
| 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> . |

Table 1-12 Veritas Volume Manager fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 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. |
| 2533248 | SAN Boot: A node can reboot and bring volumes online with a failed path, but failed to boot up after restoration of the failed path again in next reboot. |
| 2526606 | Memory leaks in DMP. |
| 2441283 | The <code>vxsnap admir</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>vxvg move</code> transaction not completing and backups fail. |
| 2227678 | The second rlink gets detached and does not connect back when overflowed in a multiple-secondaries environment. |
| 1675482 | The <code>vxvg list dgname</code> command gives error 'state=new failed'. |
| 1190117 | <code>vxdisk -f init</code> can overwrite some of the public region contents. |
| 2698035 | Tunable values do not change as per the values applied through <code>vxtune</code> . |

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 1-13 Veritas Dynamic Multi-Pathing fixed issues

| Incident | Description |
|----------|--|
| 2826958 | <code>pwnn</code> no is not displayed in the output of command <code>"vxmpadm list dmpnode dmpnodename="</code> . |
| 2818840 | Enhance the <code>vxmprow</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. |

Table 1-13 Veritas Dynamic Multi-Pathing fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 2729501 | exclude path not working properly and can cause system hang while coming up after enabling native support. |
| 2727590 | STRATUS: vxconfigd dumped core after renaming iscsi device from OS side. |
| 2700086 | EMC BCV (NR) established devices are resulting in multiple dmp events messages (paths being disabled/enabled). |
| 2652485 | Inactive snapshot luns cause trespassing. |
| 2626199 | vxdmadm list dmpnode printing incorrect path-type. |
| 2564092 | [VxVM][Usability]Automate the lun provisioning (addition) / removal steps using vxdiskadm /or new VxVM CLI command. |
| 2556467 | DMP-ASM: disable all paths and reboot host cause /etc/vx/.vxdmprawdev 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 1-14 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. |
| 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. |
| 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. |

Table 1-14 Veritas File System 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 2963763 | When the thin_friendly_alloc() and deliache_enable() functionality is enabled, VxFS may enter a deadlock. |
| 2966277 | Systems with high file system activity like read/write/open/lookup may panic the system. |
| 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. |
| 2983248 | The vxrepquota(1M) command dumps core. |
| 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 'metadatok' flag on a volume. |
| 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. |
| 3079215 | Oracle RAC Database creation failed with the Ora-00600[ksfd_odmio1] error when Veritas ODM links. |
| 3089211 | When adding or removing CPUs, Veritas File System (VxFS) may crash with Data Storage Interrupt (DSI) stack trace. |
| 3089314 | The Workload Partitions (WPARs) become unresponsive while running the pwdck command. |
| 3340286 | After a file system is resized, the tunable setting of dalloc_enable gets reset to a default value. |
| 3096834 | Intermittent vx_disable messages are displayed in the system log. |
| 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. |
| 3131360 | The vxfsconvert(1M) command fails with error messages, indicating that unallocated inodes cannot be found. |

Table 1-14 Veritas File System 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 3149174 | Veritas Oracle Disk Manager (ODM) clone shutdown fails with the ORA-03113: end-of-file on communication channel error. |
| 3164418 | Internal stress test on locally mounted VxFS filesystem 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. |
| 3197901 | Prevent duplicate symbol in VxFS libvxfspriv.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). |
| 3237204 | The vxfsstat(1M) statistical reporting tool displays inaccurate memory usage information. |
| 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. |
| 3265538 | System panics because Veritas File System (VxFS) calls the lock_done kernel service at intpri=A instead of intpri=B. |
| 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. |

Table 1-14 Veritas File System 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 3323866 | Some ODM operations may fail with "ODM ERROR V-41-4-1-328-22 Invalid argument" |
| 3331045 | Kernel Oops in unlock code of map while referring freed mlink 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. |
| 3340286 | After a file system is resized, the tunable setting of dalloc_enable gets reset to a default value. |
| 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. |
| 3364293 | On a cluster, one of the node panics in the odm_io_deinit()function, because of a race between ODM mount thread and unmount thread.[0] |
| 3364306 | Stack overflow seen in extent allocation code path. |
| 3394803 | A panic is observed in VxFS routine vx_upgrade7() functionwhile running the vxupgrade command(1M). |
| 3395692 | Double deletion of a pointer in the vx_do_getacl() function causes abend_trap(). |
| 3417076 | The vxtunefs(1M) command fails to set tunables. |
| 3417321 | The vxtunefs(1M) tunable man page gives an incorrect |
| 3418997 | The vxtunefs(1M) command accepts the garbage value for the 'write_throttle' tunable. |
| 3436699 | An assert failure occurs because of a race condition between clone mount thread and directory removal thread while pushing data on clone. |

Table 1-14 Veritas File System 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|--|
| 3448627 | The vxtunefs(1M) command accepts the garbage value for the discovered_direct_iosz tunable. |
| 3449150 | The vxtunefs(1M) command accepts garbage values for certain tunables. |
| 3449152 | Failed to set 'thin_friendly_alloc' tunable in case of cluster file system (CFS). |
| 3463717 | Information regarding Cluster File System (CFS) that does not support the 'thin_friendly_alloc' tunable is not updated in the vxtunefs(1M) command man page. |

Veritas File System: issues fixed in 6.0.3

[Table 1-15](#) describes the incidents that are fixed in Veritas File System in 6.0.3.

Table 1-15 Veritas File System 6.0.3 fixed issues

| Incident | Description |
|----------|--|
| 2978326 | Changing value of the dalloc_enable/dalloc_limit tunables fails on a cluster mounted filesystem. |
| 2895743 | Accessing named attributes for some files seems to be slow. |
| 2887423 | Severe lock contention in vx_sched. HF1e needs to be addressed. |
| 2885592 | vxdump to the vxcompress file system is aborted. |
| 2881211 | File ACLs not preserved in checkpoints properly if file has hardlink. |
| 2878164 | VxFS consuming too much pinned heap. |
| 2874054 | The vxconvert command fails to convert LVM diskgroup to vxvm diskgroup due to malloc issue |
| 2858683 | Reserve extent attributes changed after vxrestore, only for files greater than 8192bytes. |
| 2857751 | The internal testing hits the asert "f:vx_cbdnlc_enter:1a". |
| 2857629 | File system corruption can occur requiring a full fsck of the system. |
| 2857568 | Performance issues seen during back-up operations reading larges files sequentially. |

Table 1-15 Veritas File System 6.0.3 fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 2848948 | VxFS buff cache consumption increased significantly after running over 248 days. |
| 2806466 | fsadm -R resulting in panic at LVM layer due to vx_ts.ts_length set to 2GB. |
| 2624262 | fsdedup.bin hit oops at vx_bc_do_brelse. |
| 2590918 | Delay in freeing unshared extents upon primary switch over. |

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 1-16 Veritas File System fixed issues

| Incident | Description |
|----------|---|
| 2764861 | Uncompress by vxcompress ignores quota limitation. |
| 2753944 | The file creation threads can hang. |
| 2735912 | The performance of tier relocation using fspadm enforce is poor when moving a large amount of files. |
| 2712392 | Threads hung in VxFS. |
| 2709869 | System panic with redzone violation when vx_free() tried to free fiostat. |
| 2682055 | On AIX, If drefund_supported is set to 1 (i.e. drefund_supported=1) and drefund is not enabled (i.e. drefund_enable=0) then tuning of 'vmmbufs_resv_disable' to 1 fails. |
| 2674639 | The cp(1) command with the -p option may fail on a file system whose File Change Log (FCL) feature is enabled. The following error messages are displayed: cp: setting permissions for 'file_name': Input/output error cp: preserving permissions for 'file_name': No data available. |
| 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 nlinks does not update the vx_maxlink and maxlink_enable tunables. |
| 2651922 | ls -l command on local VxFS file system is running slow and high CPU usage is seen. |

Table 1-16 Veritas File System fixed issues (*continued*)

| Incident | Description |
|----------|---|
| 2650354 | Allow 8MB and 4MB values for chunk_flush_size tunable on AIX. |
| 2650330 | Accessing a file with O_NSHARE mode by multiple process concurrently on AIX could cause file system hang. |
| 2626390 | Freeing a large number of pages at once can induce a small I/O latency. |
| 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. |
| 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. |
| 2086902 | Spinlock held too long on vxfs spinlock, and there is high contention for it. |
| 1529708 | Formatting issue with the output of vxrepquota. |

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 1-17 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 1-17 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. |
| 3088915 | The Oracle resource fails to come online in the container after VCS is upgraded from 6.0.1 to 6.0.3 version. Additionally, the monitor fails to detect the processes. |
| 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. |
| 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. |

Table 1-17 Veritas Cluster Server 6.0.5 fixed issues (*continued*)

| Fixed issues | Description |
|--------------|---|
| 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. |
| 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. |
| 3315273 | LVMVG resource fails to come online if volume group configured in the resource has more than 128 disks. |
| 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. |
| 3406176 | The WPAR aware agents configured in the WPAR fail to log messages in the secure cluster. |
| 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: issues fixed in 6.0.3

[Table 1-18](#) describes the incidents that are fixed in Veritas Cluster Server in 6.0.3.

Table 1-18 Veritas Cluster Server 6.0.3 fixed issues

| Incident | Description |
|----------|--|
| 3013962 | Monitor fails to detect DB2 resource online for DB2 version 10.1. |
| 3013940 | If DB2 is installed in NON-MPP mode and UseDB2Start attribute is set to 0, we still use db2start command to start the DB2 process instead of using db2gcf command. |
| 2964772 | If you take an NFSRestart resource offline, the NFSRestart agent may unexpectedly stop NFS processes in a local container WPARs. |
| 2941155 | Group is not marked offline on faulted cluster in a GCO environment after a cluster failure is declared. |
| 2937673 | AMF driver panics the machine when amfstat is executed. |
| 2861253 | In vxfen driver log statement, jeopardy membership is printed as garbage. |
| 2848009 | AMF panicks the machine when an Agent is exiting |
| 2737653 | Incorrect descriptions about the RVGPrimary online script . |
| 2736627 | REMOTE CLUSTER STATE remains in INIT state and lcmp heartbeat status is UNKNOWN. |
| 2619600 | Use of Live Partition Mobility on an SFHA or SFCFSHA node with SCSI-3 fencing enabled for data disks causes service groups on that node to fault. |

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 1-19](#) lists the fixed issues for LLT, GAB, and I/O fencing.

Table 1-19 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 1-20](#) lists the fixed issues for bundled agents.

Table 1-20 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. |
| 2773376 | On a VCS node running the AIX operating system, the <code>Netlsnr</code> agent may fail to start a database if your system uses LDAP for authentication and if the default shell for <code>Netlsnr</code> users is CSH. |
| 2728802 | If the <code>'httpd'</code> binary or the <code>'ab'</code> binary is not present at the location that you specified in the <code>'httpdDir'</code> 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. |
| 2822920 | <code>DNSAgent</code> goes to UNKNOWN state if the Top Level Domain (TLD) is more than 4 characters in length. |
| 2850900 | <code>VSS_HOME</code> gets set to wrong directory in <code>WPAR</code> entry point, because of which the <code>WPAR</code> agent displays errors. |
| 2850902 | The <code>lpar_sysoffline</code> scripts are not updated as part of the package update. |

Table 1-20 Bundled agents fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 2779780 | When the monitor of a MultiNICB resource and online of an IPMultiNICB resource are scheduled at the same time, one of the entry point fails. This leads to fault of one of the resources and may also trigger an incorrect failover. |
| 2850858 | Error observed when ContainerInfo attribute for service group gets updated while running the <code>hawparsetup.pl</code> script. |
| 2846389 | In releases prior to VCS 6.0.1, the upper bound value of FaultTolerance attribute of the CoordPoint 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 FaultTolerance attribute of CoordPoint agent had to be set to less than the majority of the coordination points. Subsequent to VCS 6.0.1, the FaultTolerance attribute of CoordPoint agent is less than the majority of coordination points. |
| 2850903 | Following error message appears in the engine log when Application agent is configured inside WPAR and the agent tries to online the application: <pre>VCS ERROR V-16-1-10600 Cannot connect to VCS engine</pre> |

VCS engine fixed issues in 6.0.1

[Table 1-21](#) lists the fixed issues for VCS engine.

Table 1-21 VCS engine fixed issues

| Incident | Description |
|----------|--|
| 2832754 | When a Global Cluster Option (GCO) is configured across clusters having duplicate system names, command-line utility <code>hagrpd</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. |

Table 1-21 VCS engine fixed issues (*continued*)

| Incident | Description |
|----------|--|
| 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 haclus -status <cluster> command is run. |
| 2746802 | When failover group is probed, VCS engine clears the MigrateQ and TargetCount. |
| 2746816 | The syslog call used in gab_heartbeat_alarm_handler and gabsim_heartbeat_alarm_handler functions is not async signal safe. |

Enterprise agents fixed issues in 6.0.1

[Table 1-22](#) lists the fixed issues for enterprise agents.

Table 1-22 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. |
| 2773376 | Oracle agent does not function when the user authentication is performed through LDAP and the default shell is CSH. |
| 2831044 | Sybase agent script entry points must handle large process command line. |
| 2699800 | Db2udb resource is reported OFFLINE unexpectedly by the monitor entry point of the Db2udb agent. |

Agent framework fixed issues in 6.0.1

[Table 1-23](#) lists the fixed issues for agent framework.

Table 1-23 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 1-24 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 1-25](#) lists the issues fixed in 6.0.1.

Table 1-25 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> |

Table 1-25 Issues fixed in 6.0.1 (*continued*)

| Incident | Description |
|----------|--|
| 2329580 | 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. |
| 2392741 | Policy-managed Oracle RAC databases fail to come online on some of the nodes in the server pool. 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. |
| 2749412 | Setting the <code>UseVirtualIP</code> attribute to 1 overwrites the IP address of the virtual interface on some nodes in the cluster. |
| 2757032 | The PrivNIC/MultiPrivNIC 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. |
| 2580393 | Removal of SAN cable from any node in a global cluster setup takes application service groups offline on all nodes. 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. |
| 2734745 | The PrivNIC resource faults after the <code>UseVirtualIP</code> attribute is set to 1. |
| 2740150 | The SF Oracle RAC installer fails to set the value of the CSSD resource attribute <code>OfflineWaitLimit</code> 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 <code>add_drv/rem_drv</code> race between our modules which are independent of each other. |

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 1-26 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. |

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

[Table 1-27](#) describes the incidents that are fixed in Veritas Storage Foundation for Databases (SFDB) tools in 6.0.3.

Table 1-27 Veritas Storage Foundation for Databases (SFDB) tools 6.0.3 fixed issues

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

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

[Table 1-28](#) describes the Veritas Storage Foundation for Databases (SFDB) tools issues fixed in this release.

Table 1-28 SFDB tools fixed issues

| Incident | Description |
|----------------------|---|
| 2585643 | <p>If you provide an incorrect host name with the <code>-r</code> option of <code>vxsfadm</code>, 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> |
| 2582694 (2580318) | <p>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.</p> |
| 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 <code>umask</code> value, which results in the required files and directories being inaccessible to the non-root users.</p> |

Known issues

This section covers the known issues in this release.

- [Issues related to installation and upgrade](#)
- [Veritas Volume Manager known issues](#)
- [Veritas Cluster Server known issues](#)
- [Veritas Dynamic Multi-pathing known issues](#)
- [Veritas Storage Foundation known issues](#)

- [Veritas Storage Foundation and High Availability known issues](#)
- [Veritas Storage Foundation Cluster File System High Availability known issues](#)
- [Veritas Storage Foundation for Oracle RAC known issues](#)

Issues related to installation and upgrade

This section describes the known issues during installation and 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.

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.

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 filesets `VRTSspbx`, `VRTSat`, and `VRTSicsco`. 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 `VRTSspbx`, `VRTSat`, and `VRTSicisco` filesets after the upgrade process completes.

Performing an upgrade or rolling upgrade to SFHA 6.0.5 using NIM ADM may fail if the OS version is incorrect (2869221)

You may see the following error during an upgrade or rolling upgrade using NIM ADM:

```
CPI ERROR V-9-40-4782 Cannot install SFCFSHA on system
sfibmblch4-9-v07 since its oslevel is 6.1 TL 00. Upgrade the system
to 6.1 TL5 or later to install SFCFSHA
```

Workaround:

If you see the above error, upgrade the operating system to the correct technology level (TL5). To check the technology level prior to upgrading, run the `oslevel -s` command.

Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups [2574731]

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

Workaround:

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

To unfreeze the service groups manually

- 1 List all the frozen service groups

```
# hagr -list Frozen=1
```

- 2 Unfreeze all the frozen service groups:

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

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

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

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

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

After a locale change restart the vxconfig daemon (2417547)

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

Workaround: See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more information on `vxconfig` daemon recovery.

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

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

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

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

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
```

Issues with keyless licensing reminders after upgrading VRTSvlic [2141446]

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

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

To resolve this issue, perform the following steps:

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

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

```
# vxkeyless set NONE
```

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

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

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

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

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

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

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

VRTSvcsea fileset cannot be uninstalled from alternate disk in manual live upgrade

Description: In manual live upgrade procedure from 5.1x to 5.1SP1, all filesets are copied to an alternate root disk. However, VRTSvcsea fileset cannot be uninstalled from alternate disk to upgrade it to 5.1SP1.

Workaround: Instead of removing the VRTSvcsea fileset, you must apply a patch to upgrade this fileset to 5.1SP1 version.

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

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

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

Workaround: Complete the rolling upgrade to VCS 6.0.

SF failed to upgrade when Application HA is installed (3088810)

If you have installed both ApplicationHA 6.0 and SF 6.0.1, the installer can't upgrade SF 6.0.1 to 6.0.5. The following error message is displayed:

CPI ERROR V-9-30-1303 SFHA 6.0.100 does not appear to be installed on <your system>

Workaround:

Use the following command to specify the exact product for the upgrade:

```
# ./installmr -prod SF60
```

Ignore the warning message that SFHA is already installed after the pre-check and continue the upgrade.

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

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

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

1. Log into the node as the root user.
2. Kill all `vxfscd` processes:

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

3. Remove the `vxfscd-pid` file:

```
# rm /var/adm/cfs/vxfscd-pid
```

4. Bring the `vxfscd` resource online:

```
# hares -online vxfscd_resname -sys node_name
```

Veritas Volume Manager known issues

This section describes the known issues in this release of Veritas Volume Manager (VxVM).

VxVM commands may respond slowly when you disable the primary paths and run the `vxdisk scandisks` command (3450060)

EMC mirror-view LUNs can be in the Advanced Policy Firewall (APF) mode and managed by DMP. In this case, if you disable the primary paths and run the `vxdisk scandisks` command, I/O is delayed on the secondary path until it's timeout. As a result, VxVM commands respond slowly.

Workaround:

To solve this issue, use any of the workarounds:

Configure LUNs in the Asymmetric Logical Unit Access (ALUA) mode.

OR

Set the error recovery policy to `fixed retry` with the following command:

```
# vxddmpadm setattr enclosure ENCLOSURE_NAME  
recoveryoption=fixedretry retrycount=5
```

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
```

Disk write failure may occur when you convert Logical Volume Manager (LVM) configurations to VxVM configurations with vxconvert (3396489)

LVM does SCSI2 reservation on disk devices to avoid concurrency violation. But SCSI2 reservation doesn't support multi-pathing operations. If DMP is enabled, the I/O operations to the given LVM disk are allowed only on the path containing volume groups. During disk initialization, if you try to do I/O operations to LVM disks over the paths which don't contain volume groups, those I/O operations fail. Hence, the disk write failure occurs.

Workaround:

Before conversion, manually disable the paths which are to LVM disks but don't contain volume groups:

- 1 List the Logical Volume Group (LVG) on the machine:

```
# lsvg
```

- 2 Find the path that the LVG is created on:

```
# lspv | grep volume_group_name
```

- 3 Find vxdisks that are using this LVG:

```
# vxddmpadm getsubpaths all | grep path_name
```

Where *path_name* specifies the name of the path that the LVG is created on.

- 4 Find all the paths of the vxdisk:

```
# vxdisk path | grep vxdisk_name
```

- 5 Disable the paths which are to LVM disks but don't contain volume groups:

```
# vxddmpadm disable path=path_name
```

- 6 To confirm all paths are disabled except the one which has volume groups, check vxdisk path:

```
# vxdisk path | grep vxdisk_name
```

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.

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
```

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.

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

Workaround:**To recover from this situation**

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

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

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

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

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

Co-existence check might fail for CDS disks

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

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

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

Workaround:

There is no workaround for this issue.

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

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

Workaround:

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

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.

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.

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.

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

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

Workaround:

Use the following procedure to clear the clone disk flag.

To clear the clone disk flag

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

```
# vxdisk rm device_name
```

- 2 Deport the disk group.

```
# vxdg deport dg_name
```

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

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

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

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

- 5 Clear the clone_disk tag from these devices.

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

Upgrading from Veritas Storage Foundation and High Availability Solutions 5.x to 6.0.1 may fail for IBM XIV Series arrays (2715119)

Starting in the Veritas Storage Foundation and High Availability Solutions 5.1 SP1 release, the Array Support Library (ASL) for the IBM XIV enclosures converts the LUN Serial Number from Hexadecimal to Decimal. Because of this change, the enclosure names differ from releases prior to the 5.1 SP1 releases. When you upgrade Veritas Storage Foundation and High Availability Solutions from a release prior to that release to the current 6.0.1 release, XIV LUNs may go into an error state. Note that the latest RPs on 5.1/5.1SP1 are already modified to use the same logic for enclosure naming.

Workaround:

After the upgrade, run `vxddladm assign names`.

After devices that are managed by EMC PowerPath lose access to storage, Veritas Volume Manager commands are delayed (2757198)

In an environment which includes devices that are managed by EMC PowerPath, a storage loss causes Veritas Volume Manager commands to be delayed. In the event of storage loss, VxVM sends SCSI inquiry from each LUN path to check the health of path, which are delayed by the presence of EMC PowerPath.

Performance impact when a large number of disks are reconnected (2802698)

If the storage connectivity is lost to part of the storage, the disk group configuration copy is rebalanced to the disks that have connectivity. For example, if the storage for an entire enclosure is removed from a disk group with multiple enclosures. The rebalancing process takes time, during which time the `vxconfigd` daemon is busy and does not respond to commands.

Plex synchronization is not completed after resuming synchronization on a new master when the original master lost connectivity (2788077)

When you run `vxrecover -o force`, it recovers only one subvolume and it cannot detect that the rest of the volume needs recovery.

When you run the `vxassist mirror` command, you run the `vxplex att` command serially on each subvolume. If the failure happens before you start the `attach` operation (need to mark the concerned plex as the attach operation is in progress), `vxrecover` will not redo the attach operation because it cannot find any record of the attach operation in progress.

Workaround:

Run the following command on each subvolume to manually recover the complete volume:

```
# /usr/lib/vxvm/type/fsgen/vxplex -U fsgen -g diskgroup \  
-o force useopt att volume plex
```

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

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

Workaround:

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

CVMVoIDg agent may fail to deport CVM disk group

The CVM disk group is deported based on the order in which the CVMVoIDg resources are taken offline. If the CVMVoIDg 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 CVMVoIDg resource taken offline. If the attribute value is 0 for the last CVMVoIDg resource taken offline, the disk group is not deported.

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

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

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

Set the following attributes for LUNs

1 Set the following attributes:

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

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

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

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

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

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

2 Reboot the system for the changes to take effect.

If SF is installed in global system and WPAR is created with the -t option, /usr/sbin/sync gets removed from WPAR (3007163)

With AIX 6.1 TL8 and AIX 7.1 TL2, when WPARs are created with the new option `-t`, all the file systems are copied from global to WPAR. Also the files and filesets that are not visible in WPAR are deleted, so `sync` file gets deleted as part of this operation.

Any utilities that use `sync` internally might get the following error:

```
# oslevel -s
rpm_share: 0645-025 Unable to locate command module sync.
rpm_share: 0645-007 ATTENTION: init_baselib() returned an
```

```
unexpected result.  
7100-02-01-1245
```

Workaround :

Use one of the two ways to avoid this issue:

- Append -T preserve_private=yes option while creating WPAR with the new option -t.
- Copy `/usr/sbin/sync` from global system to WPAR.

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.

Veritas Cluster Server known issues

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

- [Operational issues for VCS](#)
- [Issues related to the VCS engine](#)
- [Issues related to the bundled agents](#)
- [Issues related to the VCS database agents](#)
- [Issues related to the agent framework](#)
- [Issues related to global clusters](#)
- [LLT known issues](#)
- [I/O fencing known issues](#)
- [Veritas Cluster Server agents for Veritas Volume Replicator known issues](#)
- [Issues related to Intelligent Monitoring Framework \(IMF\)](#)
- [Issues related to the Cluster Manager \(Java Console\)](#)

Operational issues for VCS

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

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

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

Connecting to the database outside VCS control using sqlplus takes too long to respond [704069]

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

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.

If a nested VxFS mount point is registered with Asynchronous Monitoring Framework (AMF), un-mount for its parent directory fails (3456400)

When agent registers for mount online monitoring of a mount point with AMF, due to a bug in AMF, a reference on the mount point is left unreleased. As a result, unmounting of parent directory fails.

Workaround:

- Do not register nested mount points with AMF

- In case you have already hit this issue, then restart the computer to move back to the stable state

Issues related to the VCS engine

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

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

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

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

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

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

Service group is not auto started on the node having incorrect value of EngineRestarted [2653688]

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

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

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

If the top level resource which does not have any 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.

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.

System sometimes displays error message with `vcscrypt` or `vcscdecrypt` [2850899]

If random number generator is not configured on your system and you run `vcscrypt` or `vcscdecrypt`, the system sometimes displays the following error message:

```
VCS ERROR V-16-1-10351 Could not set FIPS mode
```

Workaround: Ensure that the random number generator is defined on your system for encryption to work correctly. Typically, the files required for random number generator are `/dev/random` and `/dev/urandom`.

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.

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).

Every `ha` command takes longer time to execute on secure FIPS mode clusters [2847997]

In secure FIPS mode cluster, `ha` commands take 2-3 seconds more time than in secure cluster without FIPS mode for non-root users. This additional time is required to perform the FIPS self-tests before the encryption module can be used in FIPS mode.

Workaround: No workaround.

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`:

```
hagrp -offline -force ClusterService -any
```

or

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

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.

Extremely high CPU utilization may cause HAD to fail to heartbeat to GAB [1744854]

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

The `hacf -cmdtoconf` command generates a broken `main.cf` file [1919951]

The `hacf -cmdtoconf` command used with the `-dest` option removes the include statements from the types files.

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

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.

Issues related to the bundled agents

Mount agent fails to bring the Mount resource online due to OS issue [3508584]

Mount resource is configured with `jfs2` file system. While bringing the Mount resource online, the resource may fault with the following log message due to the IBM issue `IZ773575`:

```
mount: <Device_Path> on <Mount_Path>: Device busy
The current volume is: <Device_Path>
Open volume exclusive read or write returned, rc = 16
fsck: 0507-289 Device unavailable or locked by another process.
      Cannot continue.
```

Workaround: The online operation for Mount resource internally calls the `fsck` command if `mount` command fails. IBM recommends re-running the `fsck` command to resolve the issue. Hence, increase the `OnlineRetryLimit` value for Mount resource to a higher value than its default value.

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

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

Workaround: Set the `LinkTestRatio` attribute to 0.

Application agent cannot handle a case with user as root, envfile set and shell as csh [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 `csch` shell and `EnvFile` is written accordingly.

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

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

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

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

Bringing the LPAR resource offline may fail [2418615]

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

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

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

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

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

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

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

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

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

Workaround:

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

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.

MemCPUAllocator agent fails to come online if DLPAR name and hostname do not match [2954312]

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

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

HA commands inside WPAR agent get stuck due to the login/password prompt [2431884]

After upgrade of secure clusters from VCS versions lower than VCS 6.0, the HA commands that run from within the WPAR display login/password prompts. Hence, agents trying to run HA commands inside WPAR get stuck because of the prompt, as the WPAR credentials are not upgraded because of change of architecture of VxAT in VCS 6.0.

Workaround: Run `hawparsetup.pl` again for each WPAR resource. This will create new credentials for the WPAR which can be used by HA commands in VCS 6.0.

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

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

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

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

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

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

LPAR agent may dump core when all configured VIOS are down [2850898]

When using Virtual Input Output Servers (VIOS), the LPARs need a restart after VIOS restart/reboot/crash. If management LPAR is not restarted after VIOS is rebooted, then LPAR agent may dump core.

Workaround: Restart the management LPAR which was depended on the rebooted VIOS.

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.

Prevention of Concurrency Violation (PCV) is not supported for applications running in a container [2536037]

For an application running in a container, VCS uses a similar functionality as if that resource is not registered to IMF. Hence, there is no IMF control to take a resource offline. When the same resource goes online on multiple nodes, agent detects and reports to engine. Engine uses the offline monitor to take the resource offline. Hence, even though there is a time lag before the detection of the same resource coming online on multiple nodes at the same time, VCS takes the resource offline.

PCV does not function for an application running inside a WPAR on AIX.

Workaround: No workaround.

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

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

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

When two IPs of different subnets are assigned to a virtual NIC, the NIC resource might go into faulted state (2919101)

When two IPs of different subnets are assigned to a virtual NIC, the NIC resource might go into faulted state.

Workaround:

Change the order of plumbing the IPs on the interface. The base IP of the interface which belongs to the same subnet as the NetworkHosts (of NIC resource) should be plumbed last.

For example, for the following configuration of `nic1` resource of NIC agent:

```
NIC nic1(  
    Device = en0  
    NetworkHosts = { "10.209.76.1" }  
    Protocol = IPv4  
)
```

The IPs "10.209.78.46" and "192.168.1.29" are plumbed on `en0`. The order of plumbing IPs on the device should be "192.168.1.29" first and "10.209.78.46" later. This is because the NetworkHosts IP (10.209.76.1) and 10.209.78.46 are of the same subnet.

Workload Partition (WPAR) and WPAR-aware agents cannot run in a non-shared WPAR (3313690)

Non-shared WPARs have writable `/usr` file systems and `/opt` file systems, which are local to the WPAR. The product installer installs VCS packages in `/opt/VRTSvcs` and libraries in `/usr/lib` in a global environment. You cannot synchronize those VCS packages with the local copy of the `/usr` and the `/opt` file systems on the non-shared WPAR. As a result, those VCS packages are not available to the non-shared WPAR. Therefore, in absence of the VCS packages, agents that are configured to monitor applications inside the non-shared WPAR cannot run.

Workaround:

No workaround available.

The Workload Partition (WPAR) cannot reassign the unlimited shares of memory or CPU to WPAR resources (3276104)

If you specify no shares of memory, CPU, or both by the WorkLoad attribute when you create WPAR resources, the WPAR agent assigns unlimited shares of memory, CPU, or both to WPAR. When you specify the memory value as x and CPU as y by the WorkLoad attribute, the WorkLoad attribute is like the following:

WorkLoad = {Memory = x , CPU = y }.

The range of x and y is from 1 to 65535.

Meanwhile, if you make the WPAR resource offline and then online, the WPAR resource comes online with memory of x and CPU of y . After this action, if you specify no shares (**WorkLoad = {}**) again, the WPAR resource doesn't get unlimited memory or CPU on the system for a WPAR. Instead, it uses previously specified x and y .

The valid range for shares of memory and CPU should be -1 or between 1 to 65535. The value of -1 indicates assigning unlimited shares. But currently, the agent does not support negative values to the attributes. As a result, you cannot reassign unlimited shares of CPU or memory to WPAR through VCS.

Workaround:

To assign unlimited shares of CPU, enter the following command outside of VCS:

```
# chwpar -R shares_CPU=-1 wpar_name
```

To assign unlimited shares of memory, enter the following command outside of VCS:

```
# chwpar -R shares_memory=-1 wpar_name
```

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
```

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 VCS database agents**Netlsnr agent monitoring can't detect tnslnsr running on Solaris if the entire process name exceeds 79 characters [3784547]**

If the Oracle listener process is configured with a long name, consequently the tnslnsr process starts with a name longer than 79 characters. As a result, the proc structure doesn't show the full name of the Oracle listener process, and fails the Netlsnr agent monitoring.

Workaround: Configure shorter path or listener name, which does not exceed 79 characters.

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

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

Workaround:

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

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

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

NOFAILOVER action specified for certain Oracle errors

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

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

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

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

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

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

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

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

Issues related to the agent framework

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.

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.

The agent framework does not detect if service threads hang inside an entry point [1442255]

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

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

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

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

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

Workaround: No workaround.

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`
- `# hagrp -online`
- `# hagrp -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

The engine log file receives too many log messages on the secure site in global cluster environments [1919933]

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

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

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

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

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

LLT known issues

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

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

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

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

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

```
# lsattr -El SEA
```

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

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

LLT port stats sometimes shows `recvcnt` larger than `recvbytes` (1907228)

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

GAB known issues

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

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.

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 deinitiated 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.

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.

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 VRTSats fileset 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
```

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.

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

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

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

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

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

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

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

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

Workaround: Start VxFEN again after some time.

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 fileset is not installed on the system (2858190)

If VRTSvxfen fileset 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 fileset installed on the system you cannot run the utility from the install media.

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

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.

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.

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

After you run the `vxfenswap` 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.

Veritas Cluster Server agents for Veritas Volume Replicator known issues

The following are new additional Veritas Cluster Server agents for Veritas Volume Replicator known issues in 6.0.5 release.

Stale entries observed in the sample main.cf file for RVGLogowner agent [2872047]

Stale entries are found in sample `main.cf` file for RVGLogowner agent. The stale entries are present in `main.cf.seattle` file on the RVGLogowner agent which includes CFSQlogckd resource. However, CFSQlogckd is not supported since VCS 5.0.

Workaround: In the `cvm` group remove the following two lines:

```
CFSQlogckd qlogckd (  
    Critical = 0  
)
```

Issues related to Intelligent Monitoring Framework (IMF)

Process offline monitoring issues with Asynchronous Monitoring Framework[3540463]

Process offline monitoring registrations with Asynchronous Monitoring Framework (AMF) freeze nodes of Symantec Cluster Server (VCS) on AIX 7.1 TL3 SP3 and AIX 6.1 TL9 SP3 versions.

Workaround: Install the IBM authorized program analysis report (APAR) IV63274 on all nodes to fix this issue. For information related to APAR, contact IBM Support.

You can also refer to the following TechNote:

<http://www.symantec.com/docs/TECH222286>

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

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

There is no functionality impact.

Workaround: There is no workaround.

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.

Engine log shows Asynchronous Monitoring Framework (AMF) error message on using the `cfsshare` command [3235274]

When you use the `cfsshare` command, the IMF plugin may report the following error in the engine log:

```
AMF libvxamf ERROR V-292-2-139 This event is already registered
```

Workaround: This message is harmless and can be ignored. There is no workaround of the issue.

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.

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.

Veritas Dynamic Multi-pathing known issues

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

Migration of I/O fencing-enabled disks of VxVM disk group from EMC PowerPath TPD to VxVM DMP fails [3528561]

If I/O Fencing is enabled on some disks from VxVM disk group, migration of those disks from EMC PowerPath TPD to VxVM DMP fails with the following error messages:

```
VXFEN vxfenconfig NOTICE Driver will use SCSI-3 compliant disks.
VXFEN vxfenconfig ERROR V-11-2-1090 Unable to register with a
Majority of the coordination points.
```

Workaround: Restart the server.

Symantec has reported the issue to EMC PowerPath Engineering.

DS4K series array limitations

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

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

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

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.

Continuous trespass loop when a CLARiiON LUN is mapped to a different host than its snapshot (2761567)

If a CLARiiON LUN is mapped to a different host than its snapshot, a trespass on one of them could cause a trespass on the other. This behavior could result in a

loop for these LUNs, as DMP tries to fail back the LUNs if the primary paths are available.

Workaround:

To avoid this issue, turn off the `dmp_monitor_ownership` tunable:

```
# vxddmpadm settune dmp_monitor_ownership=off
```

The `dmp_monitor_fabric` is not persistent across reboot and upgrades (2975623)

The `dmp_monitor_fabric` parameter is not persistent across reboot and upgrades. Even if you have changed its value, it will be changed back to the previous value after system reboot or product upgrade.

Workaround:

Change it again after system reboot or product upgrade.

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

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

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

Example of this issue:

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

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

# vxddisk rm emc0_00bc

# mkvg -y dmxvg hdisk110
dmxvg
```

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

# vxdisk scandisks

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

To recover from this situation

- 1 Varyoff the LVM volume group:

```
# varyoffvg dmxvg
```

- 2 Remove the disk from VxVM control.

```
# vxdisk rm emc0_00bc
```

- 3 Trigger DMP reconfiguration.

```
# vxdisk scandisks
```

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

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

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 1-29](#) shows the Hitachi arrays that have new array names.

Table 1-29 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

Veritas Storage Foundation known issues

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

- [Veritas Storage Foundation known issues](#)
- [Veritas File System known issues](#)
- [Replication known issues](#)
- [Veritas Storage Foundation for Databases \(SFDB\) tools known issues](#)

Veritas Storage Foundation known issues

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

In an IPv6 environment, db2icrt and db2idrop commands return a segmentation fault error during instance creation and instance removal (1602444)

When using IBM DB2 `db2icrt` command to create a DB2 database instance on a pure IPv6 environment, the `db2icrt` command returns segmentation fault error message. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2icrt -a server -u db2fen1 db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 4700: 26182 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -addfcm -i ${INSTNAME?}
```

The `db2idrop` command also returns segmentation fault, but the instance is removed successfully after the `db2idrop` command is issued. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2idrop db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 3599: 7350 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -remove -s DB2_${INSTNAME?} 2> /dev/null
```

```
DBI1070I Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

This issue has been identified as an IBM issue. Once IBM has fixed this issue, then IBM will provide a hotfix for this segmentation problem.

At this time, you can communicate in a dual-stack to avoid the segmentation fault error message until IBM provides a hotfix.

To communicate in a dual-stack environment

- ◆ Add an IPv6 hostname as an IPv4 loopback address to the `/etc/hosts` file. For example:

```
127.0.0.1 swlx20-v6
```

Or

```
127.0.0.1 swlx20-v6.punipv6.com
```

127.0.0.1 is the IPv4 loopback address.

`swlx20-v6` and `swlx20-v6.punipv6.com` are the IPv6 hostnames.

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

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

<http://www-01.ibm.com/support/docview.wss?uid=isg1I252585>

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

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

- For version 6.1.3.1:
<http://www-01.ibm.com/support/docview.wss?uid=isg1I252959>
For the fix:
<ftp://ftp.software.ibm.com/aix/efixes/iz52959/>
- For version 6.1.2.4:
<http://www-01.ibm.com/support/docview.wss?uid=isg1I252720>
For the fix:
<ftp://ftp.software.ibm.com/aix/efixes/iz52720/>
- For version 6.1.2.5 :
<http://www-01.ibm.com/support/docview.wss?uid=isg1I252975>
For the fix:
<ftp://ftp.software.ibm.com/aix/efixes/iz52975/>

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

Oracle 11gR1 may not work on pure IPv6 environment [1819585]

There is problem running Oracle 11gR1 on a pure IPv6 environment.

Running AIX 6.1, you may receive the following error message when using sqlplus:

```
$ sqlplus "/ as sysdba"  
SQL> startup nomount  
SQL> ORA 0-0-0-0
```

Workaround: There is no workaround for this, as Oracle 11gR1 does not fully support pure IPv6 environment.

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

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

Workaround:

To resolve this known issue

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

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

DB2 databases are not visible from the VOM Web console (1850100)

If you upgraded to SF 5.1, DB2 databases will be not visible from the VOM web console.

This will be fixed in the SF 5.1 Patch 1 release.

Workaround: Reinstall is required for VOM DB2-Hotfix (`HF020008500-06.sfa`), if the host is upgraded to SF 5.1. Use the deployment framework and reinstall the hotfix for DB2 (`HF020008500-06.sfa`) on the managed host.

To resolve this issue

- 1 In the Web GUI, go to **Settings > Deployment**.
- 2 Select **HF020008500-06 hotfix**.
- 3 Click **Install**.
- 4 Check the **force** option while reinstalling the hotfix.

A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a SmartTier placement policy (1880081)

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

Workaround:

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

Veritas File System known issues

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

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

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

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

This issue is verified with the following commands:

- `cp -r`
- `du`

However, this issue might occur with other commands.

Workaround: Run the command from a different directory.

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:

```
mesg 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.

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

At times, the performance on a VxFS file system can be slower than on a JFS file system.

Workaround:

There is no workaround for this issue.

Deduplication can fail with error 110 (2591473)

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

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

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

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

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

Workaround:

Make more space available on the file system.

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

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

Workaround:

Unexport the file system prior to unmounting.

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

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

```
UX:vxfs fsadm: ERROR: V-3-20343: cannot shrink /dev/vx/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.

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.

Severe impact in read performance (sequential and random) on compressed files compared to uncompressed files (2609152)

The read throughput is highly degraded for compressed files. The difference is seen for sequential I/O and random I/O. For sequential reads, the degradation is visible even when the amount of data read compressed files is one-third of the uncompressed files (compression ratio).

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 (2909206, 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.

The vxcompress operation with the multithreaded option takes longer than the single threaded one (3031878)

The `vxcompress` operation with the multithreaded option takes longer than the single threaded one.

Workaround:

There is no workaround for this issue.

The `fsvoladm` command fails to clear the `metadataok` flag (2999560)

The `fsvoladm` command fails to clear the `metadataok` flag on some volumes of a VXFS file system, which is mounted on a volume set with 6 volumes.

Workaround:

There is no workaround for this issue.

Replication known issues

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

I/O hangs on the new master node when the primary master node is down (3431253)

In a CVR environment, if the primary master is set as the logowner for an RVG, bringing down the primary master may result in an I/O hang on the new master.

Workaround:

No workaround is available.

`vradm syncvol` command compatibility with IPv6 addresses (2075307)

The `vradm syncvol` command does not work with the compressed form of IPv6 addresses if the target disk group and volume names are not specified.

Workaround:

In IPv6 environments, if you run the `vradm syncvol` command and identify the target host using the compressed form of the IPv6 address, then you also need to specify the target disk group and volume names.

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

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

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

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

Workaround:**To resolve this issue**

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

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

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

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

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

Workaround:**To resolve this issue**

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

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

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

Workaround:

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

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

Issue 1:

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

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

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

Issue 2:

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

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

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

Workaround:

The following workarounds resolve these issues.

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

For example:

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

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

For example:

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

In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon (1672410, 1672417, 1825031)

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

Workaround:

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

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

When VVR is used for replicating shared disk groups in a Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) or Veritas Storage Foundation for Oracle RAC (SFRAC) environment consisting of three or more nodes, a logowner change event might, in rare instances, render `vradmin` commands unusable on some or all of the cluster nodes. In such instances, the following message appears in the "Config Errors:" section of the output of the `vradmin repstatus` and `vradmin printrvg` commands:

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

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

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

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

Workaround:

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

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

While vradadmin commands are running, vradmind may temporarily lose heart beats (2071568, 2275444)

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

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

Workaround:

To resolve this issue

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

```
export IPM_HEARTBEAT_TIMEOUT  
IPM_HEARTBEAT_TIMEOUT=120
```

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

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

vxassist layout removes the DCM (145413)

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

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

Cannot relayout data volumes in an RVG from concat to striped-mirror (2129601)

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

Workaround:

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

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

```
# vxrlink -g diskgroup status rlink
```
- 3 Stop the affected RVG. Enter the following:

```
# vxrvlg -g diskgroup stop rvlg
```
- 4 Disassociate the volumes from the RVG. Enter the following:

```
# vxvol -g diskgroup dis vol
```
- 5 Relayout the volumes to striped-mirror. Enter the following:

```
# vxassist -g diskgroup relayout vol layout=stripe-mirror
```
- 6 Associate the data volumes to the RVG. Enter the following:

```
# vxvol -g diskgroup assoc rvlg vol
```
- 7 Start the RVG. Enter the following:

```
# vxrvlg -g diskgroup start rvlg
```
- 8 Resume or start the applications.

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

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

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

Workaround:

There are two workarounds for this issue.

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

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

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

Workaround:

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

vradmin verifydata may report differences in a cross-endian environment (2834424)

When replicating between two nodes in a cross-platform environment, and performing an autosync or replication, the `vradmin verifydata` command may report differences. This is due to different endianness between the platforms. However, the file system on the secondary node will be consistent and up to date.

I/O hangs on the primary node when running vxrvg snaprestore operation (2762147)

In a CVR environment, if a secondary node is set as the logowner for an RVG, issuing the `vxrvg snaprestore` command on the primary node may result in an I/O hang.

The vxrecover command does not automatically recover layered volumes in an RVG (2866299)

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 layered volumes correctly. The recovery process fails to recover the subvolumes, which remain in the NEEDSYNC state.

Workaround:

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

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

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

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

Workaround:

To resize layered volumes that are associated to an RVG

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

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

- 3 Stop the affected RVG. Enter the following:

```
# vxrvrg -g diskgroup stop rvrg
```

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

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

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

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

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

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

- 7 Start the RVG. Enter the following:

```
# vxrvrg -g diskgroup start rvrg
```

- 8 Resume or start the applications.

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 vradmind functionality after a master switch operation**

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

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

- 2 Re-enter the command that failed.

Veritas Storage Foundation for Databases (SFDB) tools known issues

The following are known issues in this release of Veritas Storage Foundation for Databases (SFDB) tools.

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
```

Flashsnap clone fails under some unusual archive log 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'
```

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 5.0x to 6.0.5 (2184482)

The `sfua_rept_migrate` command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 or 5.0MP3 to SFHA or SF for Oracle RAC 6.0.5.

When upgrading from SFHA version 5.0 or 5.0MP3 to SFHA 6.0.5 the `S*vxdbsms3` startup script is renamed to `NO_S*vxdbsms3`. The `S*vxdbsms3` startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the `S*vxdbsms3` startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbsms3 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*vxdbsms3` to `S*vxdbsms3`.

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`.

Workaround

There is no workaround for this issue.

Clone fails with error "ORA-01513: invalid current time returned by operating system" with Oracle 11.2.0.3 (2804452)

While creating a clone database using any of the point-in-time copy services such as Flashsnap, SOS, Storage Checkpoint, or Filesnap, the clone fails. This problem appears to affect Oracle versions 11.2.0.2 as well as 11.2.0.3.

You might encounter an Oracle error such as the following:

```
/opt/VRTSdbed/bin/vxsfadm -s flashsnap -o clone
-a oracle -r dblxx64-16-v1 --flashsnap_name TEST11 --clone_path
/tmp/testRecoverdb --clone_name clone1
USERNAME:  oragrid
STDOUT:
Retrieving snapshot information ...           Done
Importing snapshot diskgroups ...           Done
Mounting snapshot volumes ...               Done
```

```
ORA-01513: invalid current time returned by operating system
```

This is a known Oracle bug documented in the following Oracle bug IDs:

- Bug 14102418: DATABASE DOESNT START DUE TO ORA-1513
- Bug 14036835: SEEING ORA-01513 INTERMITTENTLY

Workaround:

Retry the cloning operation until it succeeds.

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.

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.

dbed_update command failed after upgrading a Storage Foundation product from 5.1SP1RP1 to 6.0.1 on AIX 6.1 (2846434)

`dbed_update` might fail under some Oracle configurations, even when the database is up and running, with the following error message.

```
dbed_update -S apr1 -H /opt/oracle/app/oracle/product/11.2/db_1
```

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

```
Reason: ORA-01041: internal error. hostdef extension doesn't exist  
(DBD ERROR: OCI_SessionBegin)
```

You are able to connect to the database manually using `sqlplus`. This problem is because the version of `DBD::Oracle perl` module, used by the SFDB tools, uses somewhat older Oracle instant client libraries. With these, the SFDB tools are unable to connect to the Oracle database even when the database is up and running.

Workaround: There is no workaround at this point of time.

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.

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: OCIStmtExecute) ...
```

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.

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.

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](#).

Veritas Storage Foundation and High Availability known issues

For known issues of Veritas Storage Foundation and High Availability, refer to [Veritas Storage Foundation known issues](#) and [Veritas Cluster Server known issues](#).

Veritas Storage Foundation Cluster File System High Availability known issues

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

ClearClone attribute not added automatically to `CVMTypes.cf` after upgrade [3483394]

The ClearClone attribute is not added automatically to the CVMVolDg agent configuration file after upgrade.

Workaround:

Manually add ClearClone attribute to the `/etc/VRTSvcs/conf/config/CVMTypes.cf` file as follows:

```
# /opt/VRTSvcs/bin/haconf -makerw
# /opt/VRTSvcs/bin/haattr -add CVMVolDg ClearClone -integer 0
# /opt/VRTSvcs/bin/hatype -modify CVMVolDg ArgList -add ClearClone
# /opt/VRTSvcs/bin/haconf -dump -makero
```

If you roll back the software from 6.0.5 Maintenance Release, use the following steps:

```
# /opt/VRTSvcs/bin/haconf -makerw
# /opt/VRTSvcs/bin/hatype -modify CVMVolDg ArgList -delete ClearClone
# /opt/VRTSvcs/bin/haattr -delete CVMVolDg ClearClone
# /opt/VRTSvcs/bin/haconf -dump -makero
```

When you set a zero value for tunables `max_buf_data_size` and `discovered_direct_iosz`, their values change to the default ones [3449606]

When you set a zero value for tunables `max_buf_data_size` and `discovered_direct_iosz`, their values change to the default ones.

For example, the default value for the `discovered_direct_iosz` tunable is 256K.

```
# vxtunefs mount_point | grep discovered
discovered_direct_iosz = 262144
```

If we change it to 64K:

```
# vxtunefs -o discovered_direct_iosz=64k mount_point

# vxtunefs mount_point | grep discovered
discovered_direct_iosz = 65536
```

If we again change it to zero (which is invalid), it is changed to the default value, instead of keeping the previous value as it is.

Workaround:

Run the `vxtunefs` command again to set the tunable value to the previous one.

```
# vxtunefs -o discovered_direct_iosz=64k mount_point
```

On AIX, performance issue may occur if the application performs sequential unaligned large reads (3064877)

AIX uses 128K uiomove size to move data from page cache to user buffer. When the application performs sequential reads, the readahead task is triggered. However, in some scenarios, readahead may not be able to catch up with the uiomove size. In this case, page fails during uiomove, resulting in the performance degradation.

Workaround:

No workaround for this issue.

The vxfsckd resource fails to start when vxfsckd is killed manually and the cluster node is rebooted (2720034)

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

Workaround:

Use the following commands for this situation:

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

Kill all `vxfsckd` processes:

```
fsclustadm cfsdeinit
hastart
```

NFS issues with VxFS Storage Checkpoint (2027492)

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

```
Stale NFS file handle
```

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

There is no workaround at this time.

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.

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

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

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

Workaround

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

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 `fscckptadm` command:

```
# fscckptadm 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:

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

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() (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.

Inode access and modification times are not getting updated on the primary node when a file owned by the primary node is accessed from a secondary node (2170318)

The inode access times and inode modification itimes (collectively known as itimes) are not getting updated on the primary node when a file owned by the primary node is accessed from a secondary node. The primary node has a stale value for those itimes. A cluster file system requires consistent itimes on all the nodes at the same time. The system performance has a minimal impact even if itimes are not same on all nodes.

Workaround: There is no workaround for this issue.

File system check daemon fails to restart after abnormal termination (2942464)

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 `vxfsckd` processes:

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

3. Remove the `vxfsckd-pid` file:

```
# rm /var/adm/cfs/vxfsckd-pid
```

4. Bring the `vxfsckd` resource online:

```
# hares -online vxfsckd_resname -sys node_name
```

Veritas Storage Foundation for Oracle RAC known issues

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

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

Oracle 10g R2 patchset Clusterware configuration fails on AIX 6.1 TL9 due to Oracle Bug ID 18275474 [3436854]

Oracle 10g R2 patchset Clusterware configuration fails on AIX 6.1 TL9 with the following error message in system logs:

```
Duplicate clsomon found
```

Workaround:

No workaround is available. Check Oracle metalink of Bug ID 18275474 for the latest updates.

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.

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.

11gr2 Grid Infrastructure Node does not join the cluster after eviction (3667491)

The 11gr2 Grid Infrastructure Node does not join the cluster after eviction, and shows the following error in ocspd.log:

```
sgipcnUdpSend No buffer space available (74)
```

This issue is due to an Oracle Bug 12720728. Oracle GIPC has problems identifying the remote interface and binding the endpoints.

Workaround: Refer to Oracle metalink of Doc ID 1352887.1, and apply 11.2.0.2 GI PSU5 / 11.2.0.3 GI PSU3 or upgrade to 11.2.0.4.

11gr2 Grid Infrastructure Second Node Fails to Join the Cluster as CRSD and EVMD are in INTERMEDIATE State on AIX 6.1 TL8 or 7.1 TL2 (3667943)

This is an issue that truncates multicast packets on both AIX 6.1 TL08 SP01 and AIX 7.1 TL02 SP01.

Workaround: Refer to Oracle metalink of Doc ID 1528452.1, and apply APAR IV35888 for AIX 6.1 TL08 and APAR IV35893 for AIX 7.1 TL02.

SF Oracle RAC issues

This section lists the known issues in SF Oracle RAC for this release.

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>

Node fails to join the SFHA cluster if the file system containing Oracle Clusterware is not mounted (2611055)

The sequence number of the startup script for Oracle High Availability Services daemon (ohasd) is lower than some of the SFHA components such as VXFEN and VCS. During system startup, if the file system containing Oracle Clusterware does not get mounted before the ohasd startup script is executed, the script continuously waits for the file system to become available. As a result, the other scripts (including those of SFHA components) are not executed and the node being started does not join the SFHA cluster.

Workaround: If the rebooted node does not join the SFHA cluster, the cluster can be started manually using the following command:

```
# installsfrac -start node1 node2
```

Issue with format of the last 8-bit number in private IP addresses (1164506)

The PrivNIC/MultiPrivNIC resources fault if the private IP addresses have a leading 0 in any of the octets that comprise the IP address, for example X.X.X.01 or X.X.0X.1. or X.0X.X.1 or 0X.X.X.1, where X is an octet of the IP address.

When you configure private IP addresses for Oracle Clusterware, ensure that the IP addresses have a format as displayed in the following two-node example:

- On galaxy: 192.168.12.1
- On nebula: 192.168.12.2

Confirm the correct format by viewing the PrivNIC or MultiPrivNIC resource in the `/etc/VRTSvcs/conf/config/main.cf` file.

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 Basic (SF Basic) known issues

This section describes the known issues in this release of SF Basic.

SF Basic doesn't support the Install Bundles feature [3477800]

SF Basic doesn't support the Install Bundles feature.

Workaround:

No workaround is available

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 NIM.
- The Install Bundles feature for 6.0.5 does not support hot fix installation.

Limitations related to rolling upgrade

Rolling upgrade with response files to 6.0.5 is not supported for VCS, SFHA, SFCFSHA, SFRAC, SFSYBASECE, and SVS.

Limitations related to SFRAC

If you use Automatic Storage Management (ASM) storage for Oracle database, there is no need to perform the Relink operation.

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.