Symantec™ Storage Foundation and High Availability Solutions 6.1.1 Release Notes - Solaris

6.1.1 Maintenance Release

August 2015

Symantec

Symantec™ Storage Foundation and High Availability Solutions Release Notes

Product version: 6.1.1

Document version: 6.1.1 Rev 6

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

For information about Symantec's support offerings, you can visit our 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 Symantec Storage Foundation and High Availability Solutions

This document includes the following topics:

Introduction

This document provides information about the products in Symantec Storage Foundation and High Availability Solutions 6.1.1 Maintenance Release (6.1.1 MR).

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

http://www.symantec.com/docs/TECH211540

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/TECH211575

Before installing or upgrading Symantec 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 6.1.1 Installation Guide* at available on the Symantec website:

https://sort.symantec.com/documents

The information in the Release Notes supersedes the information provided in the product documents for SFHA Solutions.

This is "Document version: 6.1.1 Rev 6" of the *Symantec Storage Foundation and High Availability Solutions Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

https://sort.symantec.com/documents

List of products

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

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

Note: SFSYBASECE is supported on RHEL 6.

List of patches

The section lists the patches and packages for Solaris.

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

Table 1-1 Patches and packages for Solaris 10

Patch ID	Package Name	Products Affected	Patch Size
150717-05	VRTSvxvm	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VM	208 MB
150726-01	VRTSamf	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	5.0 MB
150727-01	VRTSvcsvmw	ApplicationHA	12 MB
150728-01	VRTSvcsag	ApplicationHA, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	7.0 MB
150729-01	VRTSvcs	ApplicationHA, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	230 MB
150730-01	VRTSllt	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	1.6MB
150731-01	VRTSsfcpi61	ApplicationHA, DMP,FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS, VM	13 MB
150732-01	VRTSvxfen	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	2.7MB
150733-01	VRTSvcsea	ApplicationHA, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	1.4 MB

Table 1-1 Patches and packages for Solaris 10 *(continued)*

Patch ID	Package Name	Products Affected	Patch Size
150734-01	VRTSdbac	SF Oracle RAC	4.5 MB
150735-01	VRTSdbed	SF, SF Oracle RAC, SFCFSHA, SFHA	113 MB
150736-01	VRTSvxfs	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE	38 MB
150746-01	VRTScps	SF Oracle RAC, SFCFSHA, SFHA, VCS	39 MB
	VRTSaslapm	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA	1.2 MB

Table 1-2 Package for Solaris 11

Package Name	Included Patches	Products Affected	Package Size
VRTSpkgs.p5p	VRTSamf, VRTSaslapm, VRTScps, VRTSdbac, VRTSdbed, VRTSgab, VRTSglm, VRTSgms, VRTSllt, VRTSsfcpi61, VRTSvcs, VRTSvcsag, VRTSvcsea, VRTSvxfen, VRTSvxfs, VRTSvxvm	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, VCS	206M

Important release information

- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:
  http://www.symantec.com/docs/TECH211540
- For the latest patches available for this release, go to:
  https://sort.symantec.com/
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware, visit the following URL:
  http://www.symantec.com/docs/TECH211575
- The software compatibility list summarizes each Storage Foundation and High Availability (SFHA) Solutions product stack and the product features, operating system versions, and third-party products it supports. For the latest information on supported software, visit the following URL:
  http://www.symantec.com/docs/TECH213121

Note: Before you install or upgrade SFHA Solutions products, review the current compatibility lists to confirm the compatibility of your hardware and software.

Changes introduced in 6.1.1

This section lists the changes in Symantec Storage Foundation and High Availability Solutions 6.1.1.

Changes related to Symantec Cluster Server (VCS)

Symantec Cluster Server (VCS) includes the following changes in 6.1.1.

Changes in Zpool agent

The Zpool agent doesn't monitor the ZFS (Zettabyte File System) file systems with the mountpoint property set to none or canmount set to off. The agent doesn't monitor those file systems even though the ChkZFSMounts attribute is 1. Because the agent considers users choose these file systems intentionally, it doesn't throw warning in such scenarios

Changes to hazonesetup and hazoneverify utilities

The `hazonesetup` command is able to create multiple zone resources in a service group. For more details, refer to the `hazonesetup` man page.

The `hazoneverify` command is able to verify multiple zone resources in a service group.

Changes related to Veritas File System (VxFS)

Veritas File System (VxFS) includes the following changes in 6.1.1.

VxFS improves File System check (`fsck`) performance

File System check (`fsck`) validates attribute inodes in parallel. Compared with single-thread validation, the multi-thread validation enables full `fsck` to validate attribute inodes much faster, especially on large file systems with many attribute inodes.

VxFS performance improves with Vnode Page Mapping (VPM) interface on Solaris SPARC

Vnode Page Mapping (VPM) interfaces replace legacy segmap. Because VPM uses KPM (kernel page mapping), VxFS performance improves with VPM interface on Solaris SPARC.

Patch Meta-data Standardizations

All stack products use the same PSTAMP format across Symantec. For example:

patch-version-data-timestamp

Products without changes in this release

The following Symantec products have no changes in this release:

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

- Symantec Storage Foundation for Sybase ASE CE (SFSYBASECE)

System requirements

This section describes the system requirements for 6.1.1.

Supported Solaris operating systems

For current updates, visit the Symantec Operations Readiness Tools Installation and Upgrade page: https://sort.symantec.com/land/install_and_upgrade.

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

Table 1-3 Supported operating systems

Operating systems	Levels	Chipsets
Solaris 10	Update 9, 10, and 11	SPARC
Solaris 11	Solaris 11.1 and up to Support Repository Updates (SRUs) 11.1.21.4.1 Solaris 11.2 and up to Support Repository Updates (SRUs) 11.2.6.5	SPARC

This release (version 6.1.1) is not supported on the x86-64 architecture.

This release (version 6.1.1) supports solaris and solaris10 brand zones on the Solaris 11 operating system and native brand zones on the Solaris 10 operating system.

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

Supported database software

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

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

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

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

https://support.oracle.com

Hardware compatibility list

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

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

Before installing or upgrading Veritas Cluster Server, review the current compatibility list to confirm the compatibility of your hardware and software.

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

Supported Oracle VM Server for SPARC

Supported Oracle VM Server for SPARC (OVM) versions are OVM 2.0, OVM 2.1, OVM 2.2, OVM 3.0, OVM 3.1, and OVM 3.1.1.

For supported OS version for Oracle VM Server for SPARC, refer to *Oracle VM server for SPARC Release Notes*.

The version of the Oracle Solaris operating system (OS) that runs on a guest domain is independent of the Oracle Solaris OS version that runs on the primary domain. Therefore, if you run the Oracle Solaris 10 OS in the primary domain, you can still run the Oracle Solaris 11 OS in a guest domain. Likewise if you run the Oracle Solaris 11 OS in the primary domain, you can still run the Oracle Solaris 10 OS in a guest domain.

The only difference between running the Oracle Solaris 10 OS or the Oracle Solaris 11 OS on the primary domain is the feature difference in each OS.

If necessary, upgrade Solaris before you install the SFHA Solutions products.

See http://www.symantec.com/docs/TECH202397 before you upgrade to Oracle Solaris 11.1.

Storage Foundation for Databases features supported in database environments

Storage Foundation for Databases (SFDB) product features are supported for the following database environments:

Table 1-4 SFDB features supported in database environments

Symantec Storage Foundation feature	DB2	Oracle	Oracle RAC	Sybase	Sybase ASE CE
Oracle Disk Manager	No	Yes	Yes	No	No
Cached Oracle Disk Manager	No	Yes	No	No	No
Quick I/O	Yes	Yes	Yes	Yes	Yes
Cached Quick I/O	Yes	Yes	Yes	Yes	Yes
Concurrent I/O	Yes	Yes	Yes	Yes	Yes
Storage Checkpoints	Yes	Yes	Yes	Yes	Yes
Flashsnap	Yes	Yes	Yes	Yes	Yes
SmartTier	Yes	Yes	Yes	Yes	Yes
Database Storage Checkpoints **Note:** Requires Enterprise license	No	Yes	Yes	No	No
Database Flashsnap **Note:** Requires Enterprise license	No	Yes	Yes	No	No
SmartTier for Oracle **Note:** Requires Enterprise license	No	Yes	Yes	No	No

Notes:

- SmartTier is an expanded and renamed version of Dynamic Storage Tiering (DST).
- Storage Foundation for Databases (SFDB) tools Database Storage Checkpoint, Database Flashsnap, and SmartTier for Oracle are supported with an Enterprise product license.

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

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

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

Symantec Storage Foundation memory requirements

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

Number of nodes supported

SFHA Solutions support cluster configurations with up to 64 nodes.

Fixed Issues

This section covers the fixed incidents.

See the `README_SYMC.xxxxx-xx` files in the /patches directory for the symptom, description, and resolution of the fixed issue.

- Installation and upgrade fixed issues
- Symantec Storage Foundation and High Availability fixed issues
- Symantec Storage Foundation for Databases (SFDB) tools fixed issues
- Symantec Cluster Server fixed issues
- Symantec Storage Foundation for Oracle RAC fixed issues
- Symantec Storage Foundation Cluster File System High Availability fixed issues
- Veritas File System fixed issues
- Veritas Volume Manager fixed issues
- Symantec Storage Foundation for Sybase ASE CE fixed issues
- Symantec ApplicationHA fixed issues
- LLT, GAB, and I/O fencing fixed issues

Installation and upgrade fixed issues

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

Installation and upgrade fixed issues in 6.1.1

Table 1-5 covers the incidents that are fixed related to installation and upgrade in 6.1.1.

Table 1-5 Installation and upgrade 6.1.1 fixed issues

Incident	Description
3448674 2618482	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.
3471289	CPI tries to stop vxglm even though Cluster Volume Manager (CVM) is online.
3519322	CPI exit with error when selecting rolling upgrade task from install menu.
3491158	Oracle agent configuration may impact upgrade to ApplicationHA version 6.1.
3491172 3343592	The installer exits without printing any message after 'Verifying systems' when it reruns './installer -license' on an already licensed system.
3542842	During upgrade, the Sun Microsystems Label (SMI) label changes to Extensible Firmware Interface (EFI) and makes the disk group un-importable.

Installation and upgrade fixed issues in 6.1

Table 1-6 covers the incidents that are fixed related to installation and upgrade in 6.1.

Table 1-6 Installation and upgrade 6.1 fixed issues

Incident	Description
2016346	Installed 5.0 MP3 without configuration, then upgrade to 6.0.1, installer cannot continue.
3098297	Uninstalling Symantec Storage Foundation and High Availability Solutions does not remove the `vxdcli` service from the svcs database.
3182366	Adding a node to a cluster fails if you did not set up passwordless ssh or rsh.
1215671	You must use the VCS installer program to install or upgrade VCS when the zone root is on Veritas File System (VxFS).
2737124	If you upgrade the `VRTSvlic` package manually, the product levels that were set using `vxkeyless` may be lost. The output of the `vxkeyless display` command does not display correctly.

Table 1-6 Installation and upgrade 6.1 fixed issues *(continued)*

Incident	Description
2141446	After upgrading from VCS 5.1 to higher versions of VCS, some keyless licenses may be left in the system. As a result periodic reminders get logged if Veritas Operations Manager Server is not configured.

Symantec Storage Foundation and High Availability fixed issues

This section includes issues fixed for Veritas File System and Veritas Volume Manager.

See "Veritas File System fixed issues" on page 29.

See "Veritas Volume Manager fixed issues" on page 34.

Symantec Storage Foundation for Databases (SFDB) tools fixed issues

This section describes fixed issues of Symantec Storage Foundation for Database (SFDB) tools.

Symantec Storage Foundation for Databases (SFDB) tools fixed issues in 6.1.1

Table 1-7 describes the Symantec Storage Foundation for Databases (SFDB) tools issues fixed in 6.1.1.

Table 1-7 SFDB tools 6.1.1 fixed issues

Incident	Description
3313775	Mount options are not restored after Reverse Resync Commit operation is performed.

Symantec Storage Foundation for Databases (SFDB) tools fixed issues in 6.1

Table 1-8 describes the Symantec Storage Foundation for Databases (SFDB) tools issues fixed in 6.1.

Table 1-8 SFDB tools 6.1 fixed issues

Incident	Description
2591463	Database Storage Checkpoint unmount may fail with device busy.
2534422	FlashSnap validate reports snapshot unsplittable.
2580318	dbed_vmclonedb ignores new clone SID value after cloning once.
2579929	User authentication fails.
2479901	FlashSnap resync fails if there is an existing space-optimized snapshot.
2869268	Checkpoint clone fails in a CFS environment if cloned using same checkpoint and same clone name on both nodes.
2849540	Very long off-host cloning times for large number of data files.
2715323	SFDB commands do not work with the ZHS16GBK character set.

Symantec Cluster Server fixed issues

This section describes Symantec Cluster Server fixed issues.

Symantec Cluster Server fixed issues in 6.1.1

Table 1-9 covers the fixed issues of Symantec Cluster Server in 6.1.1.

Table 1-9 Symantc Cluster Server 6.1.1 fixed issues

Incident	Description
3326591	The IPMultiNICB agent delays bringing IPv4 VIP online on Solaris version 10 by 5 seconds.
3338946	The Process resource fails to register for offline monitoring with the AMF kernel driver.
3341320	The "Cannot delete event (rid %d) in reaper" error message is repeatedly logged in the Syslog file.
3381042	Unexpected deletion of temporary files during VCS start causes the VCS agents to report an incorrect state.
3385820	Sometimes the high availability daemon (HAD) crashes if it runs for a long duration.
3389647	The hazonesetup and hazoneverify utilities do not support multiple zone resources in the same service group.

Table 1-9 Symantc Cluster Server 6.1.1 fixed issues *(continued)*

Incident	Description
3422904	The Zone agent does not handle a zone in unavailable state.
3436617	VCS fails to execute triggers.
3471819	The service group fails to go online if the CurrentCount attribute value is incorrect.
3485941	A service group gets stuck in the STARTING state.
3488211	The engine log is flooded with the "Cluster Synchronization Service process is not running." error message when the ocssd process is not running on the node where ASMInst resource is in offline state.
3498072	The hazonesetup(1M) utility reports a Perl warning message when the locale is set to non-English.
3533176	Zpool resource goes into unknown state for a ZFS file system if the mountpoint property is set to none or the canmount property is set to off.
3536195	The Online operation of LDom resource fails when Disaster Recovery options are configured.

Symantec Cluster Server fixed issues in 6.1

This section describes Symantec Cluster Server fixed issues in 6.1.

VCS engine fixed issues

Table 1-10 lists the fixed issues for VCS engine.

Table 1-10 VCS engine fixed issues

Incident	Description
2858188	If you attempt to reconfigure an already configured Global Cluster Option (GCO) using `gcoconfig`, the command does not change the existing GCO IP while reconfiguring the global cluster options.
2941155	Symantec Cluster Server (VCS) does not mark a group as offline on a failed cluster when a cluster failure is declared in a GCO environment .

Table 1-10 VCS engine fixed issues *(continued)*

Incident	Description
2954319	On a heavily loaded system, the logger thread frequently picks the SIGABRT from GAB. The logger thread runs at a low priority and may not get scheduled. Hence, the SIGABRT is not processed and GAB panics the machine.
2736627	Remote cluster state remains in INIT state and Icmp heartbeat status remains UNKNOWN if IPv6 is disabled on the systems.
2848005	If you terminate the CmdServer process or if it stops due to any reason on a running VCS cluster and if you stop VCS with the SMF command (for example svcadm disable <service>), the VCS SMF service goes into maintenance state as the CmdServer fails to stop
3028644	Symantec Cluster Server notifier process dumps core if there are any issues in SNMP configuration.
3042450	Parent service group which if frozen and configured with online local hard dependency is brought offline when its child service group faults.
3079893	Symantec Storage Foundation and High Availability Solutions does not retry to online a service group when a resource in the service group faults while the group is being brought online and when OnlineRetryLimit and OnlineRetryInterval for the service group is set to non -zero values.
3090710	High Availability Daemon (HAD) starts and stops before VxFEN driver configuration completes.
3207663	When user fires 'hauser -addpriv' command to set user privileges for a group and provides any string without dash (-) instead of the '-group' option syntax error is not seen and incorrect privileges are set.
3112608	Resource is unable to come online after switch fails for a service group.
3318764	While High Availability Daemon (HAD) is running, if you empty the content of the utmp file (file name differs on different operating system (OS)) and then run `hastart -version` command, the checkboot utility fails with a segmentation fault and some agents might fail.

Bundled agents fixed issues

Table 1-11 lists the fixed issues for bundled agents.

Table 1-11 Bundled agents fixed issues

Incident	Description
2989861	Incorrect command usage is displayed for `havmconfigsync` command.
2967536	The monitor entry point invokes a test command on the MonitorProgram attribute to check if it is executable. When an application is configured with a non-default user, command is executed with `su - <user> <cmd>`. This does not work in `csh` as it requires `-c` flag to invoke the command. For example: `su - <user> -c <cmd>`.
2962270	Apache agent requires online monitoring IMF support.
2979745	MultiNICA is unable to detect loss in network connectivity.
3033290	Unnecessary zoneadm messages are seen in `engine_A.log` file.
3005729	Online function of LDom agent must not stop and unbind the already online resources in all the cases. It must check whether the requirement to online a resource is met.
3153987	In Oracle Solaris , the clean entry point ofApplication agent is reported successful even when clean program returns a non-zero value.
2964772	NFSRestart Agent may unexpectedly stop NFS processes in a local container, if an NFSRestart resource is taken offline.
2847999	Mount agent does not support BlockDevice attribute with / file system of NFS server for NFS file system.
2848020	When IP is unplumbed or in case of cable pull scenario, agent fails to offline SambasShare resource.
3039221	Converted the LDom agent entry points written in shell to Perl.
3028760	NFSRestart resource does not start NFS processes, such as `statd` and `lockd`, during online or offline operations.

AMF fixed issues

Table 1-12 lists the fixed issues of AMF in 6.1.

Table 1-12 AMF fixed issues

Incident	Description
2937673	A race condition arises in the context of amfstat, group unregistration, and event notification, which causes the AMF driver to panic.

Table 1-12 AMF fixed issues *(continued)*

Incident	Description
2848009	If an agent exits while AMF is notifying it about some events, sometimes AMF causes the node to panic.
2703641	VRTSamf patch gets installed or uninstalled when some events monitored by amf remains registered even after the patch is installed or uninstalled.
3030087	The `amfconfig -Uo` command must stop IMFD and other functions internally started or setup by AMF.
2954309	Unconfigure AMF forcefully from the AMF stop script to remove any dependencies that agents might have on the AMF.
3090229	The libusnp_vxnotify.so library used for disk group notifications, goes into an infinite loop when vxconfigd daemon is unresponsive. This causes AMF to enter an inconsistent state as a result of which AMF driver panics the node.
3145047	Due to the way AMF interacts with VXFS, AMF has access into VXFS driver even if no mounts are online, without actually holding a module reference on it. Therefore VXFS can get unloaded despite AMF having access into it.
3133181	Due to an operational error in AMF driver, in some cases an ioctl made by IMFD into AMF gets stuck inside AMF. The IMFD process cannot exit until this ioctl returns back to userspace.
3018778	Perl errors seen while using `haimfconfig` command.
2619778	In a certain error condition, all mount offline events registered with AMF are notified simultaneously. This causes the error message to get printed in the engine log for each registered mount offline event.
3259682	If `vxconfigd` hangs, then registration thread of imfd trying to get disk group status from `vxconfigd` also hangs. Therefore, the `amfregister` command waiting for IMFD gets stuck.
3279336	If AMF is unconfigured while a disk group resource registration with AMF is going on, then both the contexts may enter hung state.
3177476	If a process online registration with AMF is unregistered after it has already been triggered, the machine panics.
3274145	AMF must not load if the File System itself is not yet loaded.

Table 1-12 AMF fixed issues *(continued)*

Incident	Description
3322153	A race case between registration and unregistration of any event in AMF causes soft lockup causing machine panic.

Enterprise agents fixed issues

Table 1-13 lists the fixed issues for enterprise agents.

Table 1-13 Enterprise agents fixed issues

Incident	Description
1938138	The health check monitoring in Oracle agent for VCS does not work due to incompatibility of the health check APIs provided by Oracle.
3088915	VCS reports the status of Oracle resources configured inside the container as OFFLINE even when Oracle processes are running inside the container.
2847994	The ASMDG agent delays the exit of offline entry point when it finds the device (any one of the volume) busy as indicated by the user command. For each of the disk group mentioned in ASMDG agent's DiskGroups attribute, agent runs an SQL command and gets the list of volumes used by it.
3240209	During the Oracle online operation, the Oracle agent unnecessarily tries to back up the database due to an incorrect pattern match.
1805719	Due to issues with health check monitoring, Intentional Offline does not work for VCS agent for Oracle.

Fixed operational issues

Table 1-14 lists the fixed issues for enterprise agents.

Table 1-14 Fixed operational issues

Incident	Description
3210553	If the system tags are modified without selecting the fencing option in a Replicated Data Cluster (RDC) setup, the Stretch site wizard fails to modify tags.

Symantec Storage Foundation for Oracle RAC fixed issues

This section describes the fixed issues of Symantec Storage Foundation for Oracle RAC.

Symantec Storage Foundation for Oracle RAC fixed issues in 6.1.1

Table 1-15 covers the fixed issues of SF Oracle RAC in 6.1.1.

Table 1-15 SF Oracle RAC 6.1.1 fixed issue

Incident	Description
3561361	The CRS resource changes to online or unknown state on Oracle 12c systems after the listener is manually stopped using the srvctl command.

Symantec Storage Foundation for Oracle RAC fixed issues in 6.1

Table 1-16 lists the issues fixed in 6.1.

Table 1-16 Symantec Storage Foundation for Oracle RAC 6.1 fixed issues

Incident	Description
3090447	The CRSResource agent does not support the C shell (csh) environment.
2873102	When you install, configure, or uninstall SF Oracle RAC, the installer prompts you to optionally upload installation logs to the Symantec Web site. If the installer encounters connectivity problems, you may see an error similar to the following: `Status read failed: Connection reset by peer at <media_path>/../perl/lib/5.14.2/Net/HTTP/Methods.pm line 269.`
3186840	Running the command `zoneadm -u` to upgrade zones removes some of the package binaries.
2851403	Veritas File System modules may fail to unload if SmartMove is enabled and a break-off snapshot volume has been reattached.

Symantec Storage Foundation Cluster File System High Availability fixed issues

This section describes the fixed issues of Symantec Storage Foundation Cluster File System of High Availability.

Symantec Storage Foundation Cluster File System High Availability fixed issues in 6.1.1

Table 1-17 lists the incidents that are fixed in Symantec Storage Foundation Cluster File System High Availability in 6.1.1.

See "Veritas File System fixed issues in 6.1.1" on page 29.

See "Veritas Volume Manager fixed issues in 6.1.1" on page 34.

Table 1-17 Symantec Storage Foundation Cluster File System High Availability 6.1.1 fixed issues

Incident	Description
3463464	Internal kernel functionality conformance test hits a kernel panic due to null pointer dereference.
3348520	In a Cluster File System (CFS) cluster having multi volume file system of a smaller size, execution of the fsadm command causes system hang if the free space in the file system is low.
3413926	Internal testing hangs due to high memory consumption resulting in fork failure.
3092114	The information output displayed by the "df -i" command may be inaccurate for cluster mounted file systems.
3529860	The package verification using the Apkg verifyA command fails for VRTSglm, VRTSgms, and VRTSvxfs packages on Solaris 11.
3332902	While shutting down, the system running the fsclustadm(1M) command panics.
1949445	System is unresponsive when files were created on large directory.
3534779	Internal stress testing on Cluster File System (CFS) hits a debug assert.
3449152	Failed to set 'thin_friendly_alloc' tunable in case of cluster file system (CFS).

Symantec Storage Foundation Cluster File System High Availability fixed issues in 6.1

Table 1-18lists the incidents that are fixed in Symantec Storage Foundation Cluster File System High Availability in 6.1.

See "Veritas File System fixed issues in 6.1" on page 30.

See "Veritas Volume Manager fixed issues in 6.1" on page 39.

Table 1-18 Symantec Storage Foundation Cluster File System High Availability 6.1 fixed issues

Incident	Description
3331093	The mount agent hangs during repeated switchovers due to incorrect callback mechanism between VxFS and Asynchronous Montioring Framework (AMF).
3331050	Panic in vx_recvpremnt due to null pointer dereference during cluster reconfiguration.
3331017	Service group creation using the `cfsshare addvip` command fails with -a *nodename* option.
3330991	The `vxprint` command fails during online of CVMVolDg resource, resulting in the CVMVoldg resource getting faulted, as the volume in the CVMVolume list does not exist for clones.
3331029	The CFSMount resource fails when the default resource name is already set to an existing mount resource.
3312897	The deadlock between thaw and disable-recovery during a file system freeze leads to the file system hang in CFS.
3274592	File system hangs during the fsadm reorg due to a deadlock resulting from the out-of-order processing of broadcast messages.
3263336	Clusterwide file system hangs due to a deadlock between the file system freeze and the worklist thread.
3259634	A Cluster File System having more than 4G blocks gets corrupted because the blocks containing some file system metadata get eliminated.
3235274	During the cfsshare share operation, the VCS resources are reconfigured leading to AMF related errors.
3192985	Checkpoints quota usage on Cluster File System (CFS) can be negative.

Table 1-18 Symantec Storage Foundation Cluster File System High Availability 6.1 fixed issues *(continued)*

Incident	Description
3079215	Oracle RAC Database creation fails with the `Ora-00600[ksfd_odmio1]` error when Veritas Oracle Disk Manager (ODM) is linked.
3047134	The system panics during the internal testing due to a Group Atomic Broadcast (GAB) callback routine in an interrupt context with the following message: `Kernel panic - not syncing: GLM assert GLM_SPIN_LOCK:13018485`.
3046983	The invalid CFS node number in .__fsppadm_fclextract, causes the SmartTier policy enforcement failure.
2972183	The `fsppadm`(1M) enforce command takes a long time on the secondary nodes compared with the primary nodes.
2956195	The `mmap` command in Cluster File System (CFS) environment takes a long time to complete.
2942776	CFS mount fails with the error ENXIO or EIO on volume vset device.
2923867	Cluster hangs due to deadlock during RCQ processing.
2912089	The system becomes unresponsive while growing a file through vx_growfile in a fragmented cluster file system.
2895743	Accessing named attributes for some files stored in CFS seems to be slow.
2857629	File system corruption occurs requiring a full `fsck`(1M) after cluster reconfiguration.
2834192	Unable to mount the cluster file system after the full `fsck`(1M) utility is run.
2756779	The read and write performance needs to be improved on Cluster File System (CFS) for applications that rely on the POSIX file-record using the fcntl lock.
2715175	The `cfsumount` command runs slowly on large file systems, there are some file system reconfigure threads in the kernel.
2689326	The `mount` command may hang when there are a large number of inodes with extops and a small vxfs_ninode, or a full fsck cannot fix the link count table corruptions

Table 1-18 Symantec Storage Foundation Cluster File System High Availability 6.1 fixed issues *(continued)*

Incident	Description
2647519	VxFS requires large memory for cluster mount of large file systems.
2590918	Delay in freeing unshared extents upon primary switchover.
2107152	The system panics when you unmount a mntlock protected Veritas File System, if that device is duplicately mounted under different directories.

Veritas File System fixed issues

This section describes the fixed issues of Veritas File System.

Veritas File System fixed issues in 6.1.1

Table 1-19lists the incidents that are fixed in Veritas File System (VxFS) in 6.1.1.

Table 1-19 Veritas File System 6.1.1 fixed issues

Incident	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.
3297840	A metadata corruption is found during the file removal process.
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.
3383147	The ACA operator precedence error may occur while turning off delayed allocation.
3394803	A panic is observed in VxFS routine vx_upgrade7() function while running the vxupgrade command(1M).
3417321	The vxtunefs(1M) tunable man page gives an incorrect
3434811	The vxfsconvert(1M) in VxFS 6.1 hangs.
3448492	In Solaris SPARC, introduced Vnode Page Mapping (VPM) interface.

Table 1-19 Veritas File System 6.1.1 fixed issues *(continued)*

Incident	Description
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.
3472551	The attribute validation (pass 1d) of full fsck takes too much time to complete.
3478017	Internal test hits assert in voprwunlock.
3499886	The patch ID from the VxFS in Solaris 10 patch is displayed in PSTAMP.
3541083	The vxupgrade(1M) command for layout version 10 creates 64-bit quota files with inappropriate permission configurations.

Veritas File System fixed issues in 6.1

Table 1-20lists the incidents that are fixed in Veritas File System (VxFS) in 6.1.

Table 1-20 Veritas File System 6.1 fixed issues

Incident	Description
3331134	File system hangs due to a race condition when inodes are re-used from the delicache list.
3331125	Enhancement to handle partial compressed extents during dedupe operation.
3331109	Additional checks in `fsck` to prevent file system metadata corruption with filesnap.
3331105	The `fsck` command does not validate if multiple reorg inodes point to the same source inode.
3331095	The `fsppadm` utility dumps core when an incorrect policy is specified during enforcement.
3331071	The `fsppadm` query and enforcement should honor the `-P` option to exclude private files.
3331045	The system panics in `vx_unlockmap` due to null pointer dereference.
3331032	If a sparse zone is configured on the system and is in running state, the installation of `VRTSvxfs` package fails and the vxfsdlic SMF service inside the local zones goes into maintenance state.

Table 1-20 Veritas File System 6.1 fixed issues *(continued)*

Incident	Description
3331010	File system full `fsck` fails as it erroneously accesses freed memory during RCT processing.
3330982	The VxFS-AMF integration needs to be enhanced to support mounts inside the non-global zone.
3310755	`fsck` fix to handle ZFOD extents while processsing the `VX_RCQ_OP_DEC_ALL` operation.
3308673	A fragmented FS may get disabled when delayed allocations are enabled.
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.
3291635	The file system hangs when RCQ is full.
3261462	Mapbad corruption due to buffer overrun of VX_TYPED_4 to VX_TYPED_DEV8 conversion.
3253210	The file system hangs when it has reached the space limit.
3252983	During the test after having ported from 2486597, you see a dead loop situation where CPU is taken 100%, and the system barely responds.
3249958	When /usr is mounted as a separate file system, the VxFS fails to load.
3233284	The `fsck` (1M) command hangs while checking Reference Count Table (RCT).
3228955	Some `fsck` enhancements to check that invalid extops are not present in older file system layouts.
3224101	After the optimization is enabled for updating the `i_size` across the cluster nodes lazily, the system panics.
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.
3194635	File system metadata corruption involving ZFOD extents and filesnap or compression.
3189562	Oracle daemons get hang with the `vx_growfile()` kernel function.
3164418	Data corruption happens due to the ZFOD split during ENOSPC conditions.

Table 1-20 Veritas File System 6.1 fixed issues *(continued)*

Incident	Description
3153919	The `fsadm` shrink may hang, waiting for the hlock ownership while structural file set reorg is in progress.
3152313	With the Partitioned Directories feature enabled, removing a file may panic the system.
3150368	The `vx_writesuper()` function causes the system to panic in `evfsevol_strategy()`.
3142045	With Oracle 12c version, Veritas ODM library gives a version mismatch error.
3140990	Requirement for the ability to turn off VxFS's invalidation of pages for some Network File System (NFS) workloads.
3137886	Thin Provisioning Logging does not work for reclaim operations triggered via `fsadm`.
3101418	The current time returned by the operating system (Oracle error code ORA-01513) during Oracle startup is invalid.
3096834	Intermittent `vx_disable` messages display in the system log.
3089211	When you add or remove CPUs, Veritas File System (VxFS) may crash with the Data Storage Interrupt (DSI) stack trace.
3069695	Default Access Control Lists (ACLs) are handled on named attributes.
3068902	In case of stale NFS mounts, the `statfs()` function calls on non-VxFS file systems may cause df commands to hang.
3066116	The system panics due to the NULL pointer dereference at the `vx_worklist_process()` function.
3042485	The fix to address file system metadata corruption involves named attribute directories.
3040944	The file system hangs due to a deadlock between the `dalloc` flusher thread and `dalloc` freeze under ENOSPC conditions.
3029093	The `fsck` command fails to repair a file system with inconsistencies in RCT/RCQ records.
3011959	The system may panic because of the file system locking or unlocking using the `fsadm`(1M) or the `vxumount`(1M) command.

Table 1-20 Veritas File System 6.1 fixed issues *(continued)*

Incident	Description
3003679	When you run the `fsppadm`(1M) command and remove a file with the named stream attributes `(nattr)` at the same time, the file system does not respond.
2999493	The file system check validation fails after a successful full `fsck` during the internal testing with the following message: `run_fsck : First full fsck pass failed, exiting`.
2983248	The `vxrepquota`(1M) command dumps core.
2977697	A core dump is generated while the clone is being removed.
2966277	The high file system activities like read, write, open and lookup may panic the system.
2926684	In rare cases, the system may panic while performing a logged write.
2924447	Full `fsck` perfermance needs to be improved to reduce the amount of disk I/O.
2923105	Removal of the VxFS module from the kernel takes a longer time.
2916691	The `fsdedup` command hangs with an infinite loop in `vx_dedup_extents`.
2908391	It takes a long time to remove checkpoints from the VxFS file system, when there are a large number of files present.
2906018	The `vx_iread` errors are displayed after successful log replay and mount of the file system.
2905820	If the file is being read via the NFSv4 client, then removing the same file on the NFSv4 server may hang if the file system is VxFS.
2885592	The `vxdump` operation is aborted on file systems which are compressed using the `vxcompress` command.
2881211	File ACLs are not preserved in checkpoints properly if the file has a hardlink.
2878164	VxFS consumes too much pinned heap.
2864471	The file system hangs during clone removal with Partition directory turned on.
2858683	For files greater than 8192 bytes, the reserve-extent attribute is changed after you run the command `vxrestore`.

Table 1-20 Veritas File System 6.1 fixed issues *(continued)*

Incident	Description
2841059	The file system is marked for a full `fsck` operation and the attribute inode is marked as `bad ondisk`.
2839871	On a system with DELICACHE enabled, several file system operations may hang.
2833450	The `fstyp` command returns a negative value for the ninode on file systems larger than 2 TB.
2827751	High kernel memory allocation occurs when Oracle Disk Manager (ODM) is used with non-VxVM devices.
2825125	VxFS does not support for sub-directories larger than 64K.
2781552	Mount detects the file system not being clean and hence sets the `fullfsck` flag. `fsck` is not able to clean the system.
2750860	Performance of the write operation with small request size may degrade on a large file system.
2720034, 2689195	The `vxfsckd` daemon does not restart after being manually killed.
2667658	The `fscdsconv endian` conversion operation fails because of a macro overflow.
2624262	System panics while executing dedup operation.
2444146	The Oracle Disk Manager read returns EINTR while running unspecified Oracle jobs.
2417858	VxFS quotas do not support 64 bit limits.

Veritas Volume Manager fixed issues

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

Veritas Volume Manager fixed issues in 6.1.1

Table 1-21lists the incidents that are fixed in Veritas Volume Manager (VxVM) in 6.1.1.

Table 1-21 Veritas Volume Manager 6.1.1 fixed issues

Incident	Description
2866299	The vxrecover command does not automatically recover layered volumes in an RVG.
3107699	VxDMP (Veritas Dynamic MultiPathing) causes system panic after a shutdown/reboot.
3339195	Execution of the vxdiskadm command displays an error.
3344796	Execution of the vxdiskadm command displays an error.
3421322	LDOM guest experiences intermittent I/O failures.
3421326	Symantec Dynamic Multi-Pathing (DMP) keeps on logging 'copyin failure messages' in the system log repeatedly.
3421330	vxfentsthdw utility tests fail in LDOM guest for DMP backed virtual devices.
3422504	Paths get unexpectedly disabled/enabled in the Oracle VM Server for SPARC (LDOM) guest.
3424704	When the user uses the localized messages, the vxbootsetup(1M) command fails.
2422535	Changes on the Veritas Volume Manager (VxVM) recovery operations are not retained after the patch or package upgrade.
2573229	On RHEL6, the server panics when Dynamic Multi-Pathing (DMP) executes PERSISTENT RESERVE IN command with REPORT CAPABILITIES service action on powerpath controlled device.
2812161	In a Veritas Volume Replicator (VVR) environment, after the Rlink is detached, the vxconfigd(1M) daemon on the secondary host may hang.
2847520	The resize operation on a shared linked-volume can cause data corruption on the target volume.
2999871	The vxinstall(1M) command gets into a hung state when it is invoked through Secure Shell (SSH) remote execution.
3077582	A Veritas Volume Manager (VxVM) volume may become inaccessible causing the read/write operations to fail.
3087893	EMC PowerPath pseudo device mappings change with each reboot with VxVM (Veritas Volume Manager).

Table 1-21 Veritas Volume Manager 6.1.1 fixed issues *(continued)*

Incident	Description
3236772	Heavy I/O loads on primary sites result in transaction/session timeouts between the primary and secondary sites.
3259732	In a CVR environment, rebooting the primary slave followed by connect-disconnect in loop causes rlink to detach.
3271315	The vxdiskunsetup command with the shred option fails to shred sliced or simple disks on Solaris X86 platform.
3281004	For DMP minimum queue I/O policy with large number of CPUs a couple of issues are observed.
3300418	VxVM volume operations on shared volumes cause unnecessary read I/Os.
3314647	The vxcdsconvert(1M)command fails with error: Plex column offset is not strictly increasing for column/plex.
3326964	VxVM hangs in Clustered Volume Manager (CVM) environments in the presence of FMR operations.
3340923	For Asymmetric Logical Unit Access (ALUA) array type Logical Unit Numbers (LUN), Dynamic Multi-Pathing (DMP) disables and enables the unavailable asymmetric access state paths on I/O load.
3368361	When site consistency is configured within a private disk group and CVM is up, the reattach operation of a detached site fails.
3372724	When the user installs VxVM, the system panics with a warning.
3373142	Updates to vxassist and vxedit man pages for behavioral changes after 6.0.
3373208	DMP wrongly sends the SCSI PR OUT command with APTPL bit value as A0A to arrays.
3374200	A system panic or exceptional IO delays are observed while executing snapshot operations, such as, refresh.
3377383	The vxconfigd crashes when a disk under Dynamic Multi-pathing (DMP) reports device failure.
3385753	Replication to the Disaster Recovery (DR) site hangs even though Replication links (Rlinks) are in the connected state.
3399323	The reconfiguration of Dynamic Multipathing (DMP) database fails.

Table 1-21 Veritas Volume Manager 6.1.1 fixed issues *(continued)*

Incident	Description
3400504	Upon disabling the host side Host Bus Adapter (HBA) port, extended attributes of some devices are not seen anymore.
3403390	After a crash, the linked-to volume goes into NEEDSYNC state.
3408320	Thin reclamation fails for EMC 5875 arrays.
3415188	I/O hangs during replication in Veritas Volume Replicator (VVR).
3417044	System becomes unresponsive while creating a VVR TCP connection.
3435225	In a given CVR setup, rebooting the master node causes one of the slaves to panic.
3435475	The vxcdsconvert(1M) conversion process gets aborted for a thin LUN formatted as a simple disk with Extensible Firmware Interface (EFI) format.
3437852	The system panics when Symantec Replicator Option goes to PASSTHRU mode.
3440790	The vxassist(1M) command with parameter mirror and the vxplex command(1M) with parameter att hang.
3441356	Pre-check of the upgrade_start.sh script fails on Solaris.
3444765	In Cluster Volume Manager (CVM), shared volume recovery may take long time for large configurations.
3446415	A pool may get added to the file system when the file system shrink operation is performed on FileStore.
3450758	The slave node was not able to join CVM cluster and resulted in panic.
3455460	The vxfmrshowmap and verify_dco_header utilities fail with an error.
3462171	When SCSI-3 persistent reservation command 'ioctls' are issued on non-SCSI devices, dmpnode gets disabled.
3475521	During a system reboot, the following error message is displayed on the console: es_rcm.pl:scripting protocol error
3482026	The vxattachd(1M) daemon reattaches plexes of manually detached site.
3485907	Panic occurs in the I/O code path.

Table 1-21 Veritas Volume Manager 6.1.1 fixed issues *(continued)*

Incident	Description
3492062	Dynamic Multi-Pathing (DMP) fails to get page 0x83 LUN identifier for EMC symmetrix LUNS and continuously logs error messages.
3520991	The vxconfigd(1M) daemon dumps core due to memory corruption.
3524376	Removed Patch ID from the VxVM Solaris 10 patch PSTAMP.
3526500	Disk IO failures occur with DMP IO timeout error messages when DMP (Dynamic Multi-pathing) IO statistics demon is not running.
3547931	List of modifications made to the VRTSaslapm package
3445120	Change tunable VOL_MIN_LOWMEM_SZ value to trigger early readback.
3428025	When heavy parallel I/O load is issued, the system that runs Symantec Replication Option (VVR) and is configured as VVR primary crashes.
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.
3418830	A node boot-up hangs while starting the vxconfigd(1M) daemon.
3417185	Rebooting the host, after the exclusion of a dmpnode while I/O is in progress on it, leads to the vxconfigd(1M) to dump core.
3409612	The value of reclaim_on_delete_start_time cannot be set to values outside the range: 22:00-03:59.
3399131	For PowerPath (PP) enclosure, both DA_TPD and DA_COEXIST_TPD flags are set.
3358904	The system with Asymmetric Logical Unit Access (ALUA) enclosures sometimes panics during path fault scenarios.
3353211	A. After EMC Symmetrix BCV (Business Continuance Volume) device switches to read-write mode, continuous vxdmp (Veritas Dynamic Multi Pathing) error messages flood syslog. B. DMP metanode/path under DMP metanode gets disabled unexpectedly.
3338208	The writes from fenced out LDOM guest node on Active-Passive (AP/F) shared storage device fails with an unexpected error.
3317430	The vxdiskunsetup utility throws error after upgradation from 5.1SP1RP4.

Table 1-21 Veritas Volume Manager 6.1.1 fixed issues *(continued)*

Incident	Description
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.
3197987	When vxddladm assign names file=<filename> is executed and the file has one or more invalid values for enclosure vendor ID or product ID, vxconfigd(1M) dumps core.

Veritas Volume Manager fixed issues in 6.1

Table 1-22lists the incidents that are fixed in Veritas Volume Manager (VxVM) in 6.1.

Table 1-22 Veritas Volume Manager fixed issues

Incident	Description
2787713	CVM fails to start if the first node joining the cluster has no connectivity to the storage.
2866299	The vxrecover command does not automatically recover layered volumes in an RVG.
3325371	When using snapshots, there is a panic in `vol_multistepsio_read_source`.
3312162	VVR:DV: Verification of the remote volumes found differences with `vradmin verifydata`.
3301470	All CVR nodes panic repeatedly due to a null pointer dereference in vxio.
3283525	Data Change Object (DCO) corruption after volume resize leads to `vxconfigd` hang.
3271595	VxVM should not allow turning off disk reclaim flag when there are pending reclaims on the disk.
3261601	`dmp_destroy_dmpnode` trying to free an already freed address.
3258276	DMP paths keep huge layer open number which causes SSD driver's total open number overflows (0x80000000).
3254311	The system panics when reattaching the site to a site-consistent disk group having a volume larger than 1 TB.

Table 1-22 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
3249264	Disks get into the `ERROR` state after being destroyed with the command `vxdg destroy dg-name`.
3240858	The `/etc/vx/vxesd/.udev_lock` file may have different permissions at different instances.
3237503	System hang may happen after creating space-optimized snapshot with large size cache volume.
3236773	Multiple error messages of format `vxdmp V-5-3-0 dmp_indirect_ioctl: Ioctl Failed` can be seen during set/get failover-mode for EMC ALUA disk array.
3235350	System panic by `vxiod` process.
3230148	Panic in `volmv_cvm_serialize` due to mismatch in active and serial sio counts.
3218013	Dynamic Reconfiguration (DR) Tool does not delete stale OS (operating system) device handles.
3199398	Output of the command `vxdmpadm pgrrereg` depends on the order of DMP node list where the terminal output depends on the last LUN (DMP node).
3199056	Veritas Volume Replicator (VVR) primary system panics in the `vol_cmn_err` function due to the VVR corrupted queue.
3194358	Continuous I/O error messages on OS device and DMP node can be seen in the syslog associated with the EMC Symmetrix not-ready (NR) LUNs.
3188154	`vxconfigd` is down after enabling native support on and reboot.
3185471	iSCSI Luns visible to the host triggers VM to import every disk group visible to the host, regardless of the `autoimport` flag. VCS imported disk groups on another host and triggered a split-brain situation.
3182350	If there are more than 8192 paths in the system, the `vxassist`(1M) command hangs when you create a new VxVM volume or increase the existing volume's size.
3182175	The `vxdisk -o thin,fssize list` command can report incorrect file system usage data.

Table 1-22 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
3178029	The value of "different blocks" is more than 100% while syncing the rvg.
3162418	The `vxconfigd`(1M) command dumps core due to an incorrect check in the `ddl_find_cdevno()` function.
3160973	The `vxlist`(1M) command hangs while executing on Extensible Firmware Interface (EFI) formatted disk is attached to host.
3152769	DMP Path failover takes time in Oracle VM Server for SPARC environments when one I/O domain is down.
3146715	Rlinks do not connect with Network Address Translation (NAT) configurations on Little Endian Architecture.
3138849	Sometimes `es_rcm.pl` is triggered before `vxconfigd` is up.
3131071	The Veritas Volume Manager (VxVM) patch installation in Solaris Alternate Boot Environment (ABE) results in data corruption.
3130876	`vxconfigd` hangs on master after you remove and add data and Data Change Object (DCO) disk from all the node and wait for all site to active state(cc setup).
3130379	The `vxplex` command core dumped under random memory allocation failures.
3126204	[VVR] : machine panics when SRL is full.
3125631	With latest train snapshot fails for dbdst setup with error `vxsnap ERROR V-5-1-6433 Component volume has changed`.
3121380	I/O of replicated volume group (RVG) hangs after one data volume is disabled.
3116990	The syslog is filled with extra write protected messages.
3114999	An operating system boot fails after an encapsulated upgrade of SF 6.0.1.
3114134	Smart(sync) Autosync fails to work and instead replicates the entire volume size for larger sized volumes.
3111062	Make the `vxrsync` socket connection mechanism more robust.
3107741	`vxrvg snapdestroy` fails with a `Transaction aborted waiting for io drain` error and `vxconfigd` hangs for about 45 minutes.

Table 1-22 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
3103168	Primary master oops after repeatedly disconnecting and connecting replication link and reboot the primary slave.
3101419	In a CVR environment, when the SRL overflows, the replicated volume group (RVG) I/O hangs for a long time.
3090667	System panics/hangs while executing `vxdisk -o thin,fssize list` as part of VOM SF discovery.
3086627	The `VxVM vxdisk ERROR V-5-1-16282 Cannot retrieve stats: Bad address` error message displays while using vxdisk -o thin,fssize list for hitachi_usp-vm0 enclosure on the array configured for truecopy P-VOLs.
3076093	The patch upgrade script `installrp` can panic the system while doing a patch upgrade.
3063378	Some VxVM commands run slowly when EMC PowerPath presents and manages "read only" devices such as EMC SRDF-WD or BCV-NR.
3046560	Ioctl DKIOCGVTOC to raw character volume fails.
3041014	Beautify error messages which are seen on the execution of relayout command.
3038382	The `vxlufinish`(1M) command runs `fuser -k` on non-root file systems, which is unexpected.
3026977	Dynamic Reconfiguration (DR) operation of `vxdiskadm` removes LUNs even those which are not in Failing/Unusable state.
3019684	IO hang is observed when SRL is about to overflow after logowner switch from slave to master.
3015181	IO hangs on both nodes of cluster when you disable the diskarray.
3012929	`vxconfigbackup` keeps old disk names in its files and gives errors, when disk names are changed.
3011405	The `vxtune -o export` command failed with `V-5-1-8826 (EXDEV).`
3010191	Previously excluded paths are not excluded after upgrade to VxVM 5.1SP1RP3.
3002770	While issuing a SCSI inquiry command, NULL pointer dereference in DMP causes system panic.

Table 1-22 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
3002498	When a disk is initialized with the `vxdisk -f init daname` command, `vxconfigd`(1M) dumps core.
2994976	BAD TRAP panic in `vxio:vol_mv_pldet_callback`.
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.
2979824	The `vxdiskadm`(1M) utility bug results in the exclusion of the unintended paths.
2970368	Enhance handling of SRDF-R2 Write-Disabled devices in DMP.
2969844	The device discovery failure should not cause the DMP database to be destroyed completely.
2966990	In a Veritas Volume Replicator (VVR) environment, the I/O hangs at the primary side after multiple cluster reconfigurations are triggered in parallel.
2964547	Cannot load module '`misc/ted`' during system reboot.
2959733	Handling the device path reconfiguration in case the device paths are moved across LUNs or enclosures to prevent the `vxconfigd`(1M) daemon coredump.
2959325	The `vxconfigd`(1M) daemon dumps core while performing the disk group move operation.
2957556	The '`vxdisksetup`' command fails when tpdmode is native and enclosure based naming scheme is on.
2948172	Executing the `vxdisk -o thin,fssize list` command can result in panic.
2946440	Add back the support for "INF" for LSI and ENGENIO VIDs to the LSI ASL.
2940446	A full file system check (fsck) hangs on I/O in Veritas Volume Manager (VxVM) when the cache object size is very large.
2925893	Make changes to Huawei APM to skip re-registering the keys on Secondary during failover.

Table 1-22 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2921816	System panics while starting replication after disabling the DCM volumes.
2919714	On a thin Logical Unit Number (LUN), the `vxevac`(1M) command returns 0 without migrating the unmounted-VxFS volumes.
2919318	The I/O fencing key value of data disk are different and abnormal in a VCS cluster with I/O fencing.
2915836	`vxnotify` does not report volume enabled message.
2915751	Solaris machine panics during dynamic lun expansion of a CDS disk.
2915063	During the detachment of a plex of a volume in the Cluster Volume Manager (CVM) environment, the system panics.
2911040	The restore operation from a cascaded snapshot leaves the volume in unusable state if any cascaded snapshot is in the detached state.
2910367	When SRL on the secondary site is disabled, the secondary node panics.
2899173	The `vxconfigd`(1M) daemon hangs after executing the `vradmin stoprep` command.
2898547	The `vradmind` process dumps core on the Veritas Volume Replicator (VVR) secondary site in a Clustered Volume Replicator (CVR) environment, when Logowner Service Group on VVR Primary Site is shuffled across its CVM (Clustered Volume Manager) nodes
2898324	UMR errors reported by Purify tool in `vradmind migrate` command.
2882566	On Solaris, you can successfully add a disk which is removed from a disk group using the `vxdg rmdisk -k` command to another disk group without any error messages
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.
2880981	Thin Reclamation of EMC Symmetrix array with Microcode 5876 could fail with error EIO.
2878876	The `vxconfigd` daemon dumps core in `vol_cbr_dolog()` due to race between two threads processing requests from the same client.

Table 1-22 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2876706	VxVM commands hang when a LUN is changed to `not_ready` state from the array.
2875962	During the upgrade of VRTSaslapm package, a conflict is encountered with VRTSvxvm package because an APM binary is included in VRTSvxvm package which is already installed.
2869514	Issue with a configuration with large number of disks when the joining node is missing disks.
2866997	VxVM Disk initialization fails as an un-initialized variable gets an unexpected value after OS patch installation.
2866299	The `vxrecover` command does not automatically recover layered volumes in an RVG.
2860230	Shared disk remains as opaque after vxdiskunsetup it on master node.
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
2858853	After master switch, `vxconfigd` dumps core on old master.
2857044	System crashes while resizing a volume with Data Change Object (DCO) version 30.
2851403	System panics while unloading `vxio` module when SmartMove feature is used and the `vxportal` module is reloaded (for example, during VxFS package upgrade).
2845383	The site gets detached if the plex detach operation is performed with the site- consistency set to off.
2836528	Unable to grow LUN dynamically on Solaris x86 using `vxdisk resize` command.
2815517	The `vxdg adddisk` command allows mixing of clone and non-clone disks in a disk group.
2807158	On Solaris platform, sometimes system can hang during VM upgrade or patch installation.
2779580	The `vradmin repstatus` operation may display a configuration error after cluster reconfiguration in a CVR environment.

Table 1-22 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2762147	I/O hangs on the primary node when running the `vxrvg` snaprestore operation.
2753954	At cable disconnect on port1 of dual-port FC HBA, paths via port2 marked SUSPECT.
2751423	`vxconfigd` core dumps in `ddl_migration_devlist_removed` during execution of internal testing.
2742706	Panic due to mutex not being released in vxlo_open.
2737686	The `vxddladm list [devices|hbas|ports|targets]` command shows invalid output in some platforms and in some platforms the output fields are empty.
2715129	`vxconfigd` hangs during Master takeover in a CVM (Clustered Volume Manager) environment.
2643506	`vxconfigd` core dumps when different LUNs of same enclosure are configured with different array modes.
2567618	The VRTSexplorer dumps core in `vxcheckhbaapi/print_target_map_entry`.
2510928	The extended attributes reported by `vxdisk -e list` for the EMC SRDF luns are reported as `tdev mirror`, instead of `tdev srdf-r1`.
2422535	Changes on the Veritas Volume Manager (VxVM) recovery operations are not retained after the patch or package upgrade.
2398954	The system panics while performing I/O on a VxFS mounted instant snapshot with the Oracle Disk Manager (ODM) SmartSync enabled.
2366066	The VxVM (Veritas Volume Manager) `vxstat` command displays absurd statistics for READ & WRITE operations on VxVM objects.
2165920	The `vxrelocd`(1M) daemon creates a defunct (zombie) process.
2152830	In a multilevel clone disks environment, a regular disk group import should be handled properly. In the case of a disk group import failure, it should report the correct error message.
2149922	Importing a disk group using clone disks fails with a "wrong usage" or "invalid attribute" error.
2123677	Expanding a LUN to a size greater than 1 TB fails to show the correct expanded size.

Table 1-22 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
2106530	`vxresize`(1M) fails for a data volume in rootdg if filesystem is mounted using block device reference as `bootdg`.
2101093	A system panic is observed in the dmp_signal_event() function.
2000585	`vxrecover` does not start remaining volumes if one of the volumes is removed during `vxrecover` command run.
1973983	The `vxunreloc`(1M) command fails when the Data Change Object (DCO) plex is in DISABLED state.
1953257	Panic in voldiodone, because a disk with hung IO is moved out of the disk group.
1952197	Running vxtrace against a volume shows response times as negative.
1942051	I/O hangs on master node after disabling secondary paths from slave node and rebooting slave node.
1903700	Removing mirror using `vxassist` does not work.
1902483	Unique PGR key per group is not needed.
1765916	VxVM socket files do not have proper write protection.
3139983	Fix the design issues w.r.t to fixed, timebound and path_busy error retry.
3137603	VxDMP module not getting loaded after bundle installation and reboot.
3122546	System panicked while booting up from mirror when ZFS DMP root support is on.
3115206	An LDOM panicked while booting up after enabling ZFS root support.
3055891	The guest domain cannot recognize the DMP disk
3038684	Restore daemon enables the paths of BCV NR devices.
2874810	Installing DMP with a keyless license or DMP-only license does not enable DMP native support for LVM root volumes
2421823	Turning off the DMP native support does not reset the preferred_names field in lvm.conf to the original values
1289985	`vxconfigd` core dumps upon running the `vxdctl enable` command.

Symantec Storage Foundation for Sybase ASE CE fixed issues

This section describes the incidents that are fixed in Symantec Storage Foundation for Sybase ASE CE.

Symantec Storage Foundation for Sybase ASE CE fixed issues in 6.1.1

There is no fixed issue for Symantec Storage Foundation for Sybase ASE CE in 6.1.1.

Symantec Storage Foundation for Sybase ASE CE fixed issues in 6.1

This section describes the fixed issues of Symantec Storage Foundation for Sybase ASE CE in 6.1

Table 1-23 Symantec Storage Foundation for Sybase ASE CE 6.1 fixed issues

Incident	Description
2615341	AutoFailOver = 0 attribute absent in the sample files at `/etc/VRTSagents/ha/conf/Sybase`.

Symantec ApplicationHA fixed issues

This section describes the fixed issues of Symantec ApplicationHA.

Symantec ApplicationHA 6.1.1 fixed issues

lists the fixed issue for Symantec ApplicationHA in 6.1.1.

Table 1-24 Symantec ApplicationHA 6.1.1 fixed issues

Incident	Description
3560240	If a node is restarted after upgrade, VCS is started even when an application is not configured.

Symantec ApplicationHA 6.1 fixed issues

Table 1-25 Symantec ApplicationHA 6.1 fixed issues

Incident number	Description
2141382	When you install ApplicationHA on a virtual machine, and then try to install Symantec Storage Foundation (SF), you may notice errors in the SF installation.

LLT, GAB, and I/O fencing fixed issues

This section describes the fixed issues of LLT, GAB and I/O fencing.

LLT, GAB, and I/O fencing fixed issues in 6.1.1

Table 1-26lists the fixed issues for LLT, GAB, and I/O fencing in 6.1.1.

Table 1-26 LLT, GAB, and I/O fencing 6.1.1 fixed issues

Incident	Description
3031216	The dash (-) in a disk group name causes vxfentsthdw(1M) and Vxfenswap(1M) utilities to fail.
3302091	VCS and SFRAC fail to start after a reboot on Solaris version 11.
3379052	SMF services for LLT, GAB, and VXFEN packages go into maintenance state when a system is shut down.
3393407	On Solaris 11, LLT, GAB and VXFEN drivers may fail to start after upgrading VCS on an alternate boot environment.
3460406	The /opt/VRTScps/bin/vxcpserv process generates a core dump.
3471571	Cluster nodes may panic if you stop the HAD process by force on a node and reboot that node.
3532859	The Coordpoint agent monitor fails if the cluster has a large number of coordination points.

LLT, GAB, and I/O fencing fixed issues in 6.1

Table 1-27lists the fixed issues for LLT, GAB, and I/O fencing in 6.1.

Table 1-27 LLT, GAB, and I/O fencing 6.1 fixed issues

Incident	Description
2869763	When you run the `addnode -responsefile` command, if the cluster is using LLT over UDP, then the `/etc/lltab` file generated on new nodes is not correct. So, the procedure fails and you cannot add nodes to a cluster using CPI response files.
2991093	The preferred fencing node weight does not get reset to the default value when HAD is terminated. In spite of lack of high availability on that node, fencing may give preference to that node in a network partition scenario.
2995937	The default value of preferred fencing node weight that vxfen uses is 1 (one). However, when HAD starts without any service group or if HAD is stopped or terminated, the node weight is reset to 0 (zero). Since vxfen resets the preferred fencing weight to its default value when HAD gets terminated, stopping HAD and killing HAD shows different preferred fencing weight.
3025931	There is a corner case where while shutting down the system, if the GAB service stop script is not run successfully, then, on the next reboot the GAB service fails to load. The GAB driver remains added into the system but module does not get loaded. In such cases, devlink entries are not created for the GAB driver and configuration of gab fails.
2110148	Installer is unable to split a cluster that is registered with one or more CP servers.
2802682	Server-based fencing may fail to start if you use the existing configuration files after reinstalling the stack.
2858190	If VRTSvxfen package 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 package installed on the system you cannot run the utility from the install media.
2724565	In SFRAC environments, sometimes GAB might fail to start because of the race between GAB and LMX in calling add_drv.
3140359	Port a does not come up due to race between `gabconfig -cx` and `gabconfig -x`.
3101262	GAB queue is overloaded causing memory pressure during I/O shipping.
3218714	GAB does not log messages about changing tunable values.
2858076	Changing the module parameter `gab_conn_wait` had no effect.

Known issues

This section covers the known issues in this release.

Issues related to installation and upgrade

This section describes the known issues during installation and upgrade.

The installer prompts warning during rollback of SFSYBASECE on Solaris 10 [3566593]

During rollback on Solaris 10, the installer prompts the following warning message despite SFSYBASECE 6.1.1 is completely installed:

```
CPI WARNING V-9-30-2634 6.1.1 does not appear to be completely\
installed on system_name
```

Workaround: No workaround. You can safely ignore the error message.

Some modules fail to unload during installation or upgrade of SFCFSHA and SF Oracle RAC packages [3451707, 3560458]

Some modules fail to unload during installation or upgrade of SFCFSHA and SF Oracle RAC packages. The issue is seen with the recent versions (SRUs, updates) of Solaris 11 operating system. During installation or upgrade of SFCFSHA and SF Oracle RAC packages, Global Atomic Broadcast (GAB), Group Lock Manager

(GLM) and Group Messaging Services (GMS) cannot be unloaded. Consequently the installer fails with the following error messages:

```
Stopping vxgms
.....................................................................................
Failed
Stopping vxglm
.....................................................................................
Failed"
..
vxgms failed to stop on node1
vxglm failed to stop on node1
..
"Symantec Storage Foundation Cluster File System HA
Shutdown did not complete successfully
```

Workaround: Restart the system.

Node panics after upgrade from Solaris 11 to Solaris 11.1 on systems running version 6.0.1 or earlier

Nodes running version 6.0.1 or earlier panic after you upgrade the operating system from Solaris 11 to Solaris 11.1. This is due to changes introduced in the Solaris operating system.

The issue is observed for SFHA, VCS, SFCFSHA, and SF Oracle RAC.

Workaround: Perform the following steps during the operating system upgrade from Solaris 11 to Solaris 11.1 before you boot to the Solaris 11.1 boot environment. This will prevent the product from starting on the Solaris 11.1 boot environment.

Open the file `/etc/default/llt` on the new boot environment and set `LLT_START` to 0.

Open the file `/etc/default/gab` on the new boot environment and set `GAB_START` to 0

Open the file `/etc/default/amf` on the new boot environment and set `AMF_START` to 0

Open the file `/etc/default/vxfen` on the new boot environment and set `VXFEN_START` to 0

After the operating system is upgraded to Solaris 11.1, upgrade the product to a version that support Solaris 11.1.

LLT, GAB, and fencing configuration fails after reinstalling 6.1 on Solaris 11 systems [3559175]

After upgrading to version 6.1.1 on Solaris 11 systems, if you uninstall VCS, SFHA, SFCFSHA, SF Oracle RAC, or SF Sybase CE and reinstall version 6.1, the LLT, GAB, and fencing configuration fails.

Workaround: Perform the following steps to resolve the issue:

1. Add the module drivers to the system:

```
# /lib/svc/method/llt-postinstall
# /lib/svc/method/gab-postinstall
# /lib/svc/method/vxfen-postinstall
```

2. Enable SMF services for each module:

```
# svcadm disable system/llt; svcadm enable system/llt
# svcadm disable system/gab; svcadm enable system/gab
# svcadm disable system/vxfen; svcadm enable system/vxfen
```

   For SF Oracle RAC, perform the following additional step to enable SMF services for VCSMM:

```
# svcadm disable system/vcsmm; svcadm enable system/vcsmm
```

Oracle Disk Manager (ODM) service may fail to start [3532633]

The Oracle Disk Manager (ODM) service may fail to start during jumpstart installation or normal installation. Jumpstart installation uses the `pkgadd` command with the ROOT path as an argument. If the ROOT path is set to “/a”, the ODM service fails to start.

Jumpstart uses the following command to install packages:

```
pkgadd -v -a ${PKGDIR}/admin  -d ${PKGDIR}/${PKG}.pkg -R ${ROOT} ${PKG}
```

Workaround: If the ODM service fails, configure ODM manually. Then, manually start the ODM service by using the following commands:

```
# /lib/svc/method/odm start
# svcadm -v enable -r vxodm
```

Incorrect `vol_min_lowmem_sz` tunable setting after upgrade to version 6.1.1 [3540898]

The `vol_min_lowmem_sz` tunable may be set to a value less than its default value of 32 MB after you upgrade to version 6.1.1. Additionally, the `vxtune` command may allow the tunable value to be thus modified without displaying an error.

Workaround: The problem has no critical functionality impact. However, for performance considerations, it is recommended that you verify that the value of the `vol_min_lowmem_sz` tunable is set to at least its default value. Use the `vxtune` command to modify the tunable value.

LLT and GAB fail to startup after rebooting from the alternate boot envrionment [3548629]

After you install LLT and GAB on the alternate boot environment and reboot the system, LLT and GAB fails to start as the corresponding drivers are not loaded. This is because the LLT and GAB startup script runs before the corresponding postinstall scripts.

Workaround: Reboot the system.

Web-based installer doesn't support or block rolling upgrade from releases before 6.1 to 6.1.1. [3555013]

The web-based installer doesn't support or block rolling upgrade from releases before 6.1 to 6.1.1. Issues may occur if you use the web-based installer for rolling upgrade from releases before 6.1 to 6.1.1.

Workaround: Use the script-based installer with the -base_path option for the rolling upgrade from releases before 6.1 to 6.1.1.

Drivers may not be loaded after reboot [3536921]

When the installer stops the processes such as vxportal, vxio, vxglm, and vxgms, it uses the `rem_drv` command to ensure that the OS doesn't load the drivers back. Consequently, after reboot the unloaded drivers make the product unavailable.

Workaround: If drivers such as vxdmp, vxio are not loaded after reboot, enter the following command to start your product:

`/opt/VRTS/install`*`prod`* `-start`

Relink Oracle Database Binary menu is missing from upgrade and rollback tasks [3518852]

If you use the response file that is generated by the `-makeresponsefile` command, the installer omits to relink Oracle database.

Workaround: Add the relink info to `responsefile` manually. For example, add:

```
$CFG{crs_home}="/u01/app/grid/product/11.2.0/gridhome";
$CFG{db_home}="/u01/app/oracle/product/11.2.0/dbhome_1";
$CFG{oracle_group}="oinstall";
$CFG{oracle_user}="oracle";
$CFG{relink_oracle_database}=1;
```

Refresh keys fail on a non-secure cluster with fencing configured [3544942]

In a non-secure cluster, after the fencing is configured, if you use the `-fencing` option and select **refresh keys/registrations on the existing coordination pints**, the installer reports `could not find the domain name!.`

Workaround: Use the `-security` option to enable security cluster, and then use the `-security` option again to change the cluster to non-secure mode.

6.1 hot fix scripts deleted after rollback [3539990]

If you upgrade from 6.1 to 6.1.1 with 6.1 hot fixes installed, and then roll back to 6.1 using `/opt/VRTS/install/uninstallmr`, the scripts of the hot fixes also get deleted.

The issue happens to Solaris 10 SPARC.

Workaround: If you want to remove the 6.1 hot fixes, uninstall them manually. Refer to the README files of the hot fixes about the procedure of uninstallation.

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:

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

2 Unfreeze all the frozen service groups:

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

On Solaris 10 xprtld will not be started if user use jumpstart to install product (3325954)

If you install the operating system plus the Symantec product using the JumpStart method and after installation, reboot the machine then configure and start the product, all the processes will be started except for `xprtld` process.

Workaround:

After reboot, manually execute the following command to start `xprtld`:

```
# /opt/VRTSsfmh/adm/xprtldctrl start
```

Flash Archive installation not supported if the target system's root disk is encapsulated

Symantec does not support SFHA Solutions installation using Flash Archive if the target system's root disk is encapsulated.

Make sure that the target system's root disk is unencapsulated before starting the installation.

Upgrade or uninstallation of Symantec Storage Foundation HA may encounter module unload failures (2159652)

When you upgrade or uninstall Symantec Storage Foundation HA, some modules may fail to unload with error messages similar to the following messages:

```
fdd failed to stop on node_name
vxfs failed to stop on node_name
```

The issue may be observed on any one or all the nodes in the sub-cluster.

Workaround: After the upgrade or uninstallation completes, follow the instructions provided by the installer to resolve the issue.

SF Oracle RAC installer does not support use of `makeresponsefile` option (2577669)

The SF Oracle RAC installer does not support the use of `makeresponsefile` option for configuring Oracle RAC settings. The following message is displayed when you attempt to configure Oracle RAC using the option:

```
Currently SFRAC installer does not support -makeresponsefile option.
```

Workaround: Configure Oracle RAC by editing the response file manually.

During Live Upgrade, installer displays incorrect message about VRTSaa package removal (1710504)

If you use Live Upgrade to upgrade SFHA Solutions 5.0MP1 to SFHA Solutions 6.1.1, the installer may display a message that the VRTSaa package failed to uninstall.

Workaround: Verify whether the `VRTSaa` package was removed correctly from the alternate boot disk.

```
# pkginfo -R alternate_root_path -l VRTSaa
```

For example, run the following command

```
# pkginfo -R /altroot.5.10 -l VRTSaa
```

If the VRTSaa package was removed, you can ignore this error.

If the VRTSaa package was not removed, remove the package manually:

```
# pkgrm -R alternate_root_path -l VRTSaa
```

For example, run the following command

```
# pkgrm -R /altroot.5.10 -l VRTSaa
```

After Live Upgrade to Solaris 10 Update 10, boot from alternate boot environment may fail (2370250)

If your setup involves volumes in a shared disk group that are mounted as CFS in a cluster, then during Live Upgrade using the `vxlustart` command from any supported Solaris version to Solaris 10 Update 10, boot from an alternate boot environment may fail.

Workaround: Run the `vxlufinish` command. Before rebooting the system, manually delete the entries of all the volumes of shared disks that are mounted as CFS in the `/altroot.5.10/etc/vfstab` directory.

Live Upgrade to Solaris 10 Update 10 fails in the presence of zones (2521348)

SFCFSHA Live Upgrade from Solaris 10 Update 7 5.1SP1 to Solaris 10 Update 10 using the `vxlustart` commands fails in the presence of zones with the following error message:

```
ERROR: Installation of the packages from this media of the media failed;
pfinstall returned these diagnostics:
Processing default locales
        - Specifying default locale (en_US.ISO8859-1)
Processing profile
ERROR: This slice can't be upgraded because of missing usr packages for
the following zones:
ERROR:    zone1
ERROR:    zone1
ERROR: This slice cannot be upgraded because of missing usr packages for
one or more zones.
The Solaris upgrade of the boot environment <dest.27152> failed.
```

This is a known issue with the Solaris `luupgrade` command.

Workaround: Check with Oracle for possible workarounds for this issue.

On Sparc, Live Upgrade from Solaris 9 to Solaris 10 Update 10 may fail (2424410)

On Sparc, Live Upgrade from Solaris 9 to Solaris 10 Update 10 may fail with the following error:

```
Generating file list.
Copying data from PBE <source.24429> to ABE <dest.24429>.
99% of filenames transferredERROR: Data duplication process terminated
unexpectedly.
ERROR: The output is </tmp/lucreate.13165.29314/lucopy.errors.29314>.

29794 Killed
Fixing zonepaths in ABE.
Unmounting ABE <dest.24429>.
100% of filenames transferredReverting state of zones in PBE
<source.24429>.
ERROR: Unable to copy file systems from boot environment <source.24429>
to BE <dest.24429>.
ERROR: Unable to populate file systems on boot environment <dest.24429>.
```

```
Removing incomplete BE <dest.24429>.
ERROR: Cannot make file systems for boot environment <dest.24429>.
```

This is a known issue with the Solaris `lucreate` command.

Workaround: Install Oracle patch 113280-10,121430-72 or higher before running `vxlustart`.

After a locale change, you need to restart the vxconfig daemon (2417547)

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

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

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

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

If the CVM protocol version does not upgrade successfully, upgrade the CVM protocol on the CVM master node.

To upgrade the CVM protocol on the CVM master node

1 Find out which node is the CVM master:

```
# vxdctl -c mode
```

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

```
# vxdctl upgrade
```

During upgrade from 5.1SP1 to 61 with an encapsulated root disk, splitting the root mirror fails if the target disk group name is used by a deported disk group (2280560)

During an upgrade from SFHA Solutions 5.1 SP1 to SFHA Solutions 6.1 with an encapsulated root disk, splitting the root mirror fails if the target disk group name for the split operation is used by an existing deported disk group.

Workaround:

Specify a different disk group name as a target for the split operation.

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

If you install or configure SFHA Solutions 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.

On Solaris 10, a flash archive installed through JumpStart may cause a new system to go into maintenance mode on reboot (2379123)

If a Flash archive is created on a golden host with encapsulated root disks, when this Flash archive is installed onto another host through JumpStart, the new system may go to maintenance mode when you initially reboot it.

This problem is caused by the predefined root disk mirror in the Flash archive. When the archive is applied to a clone system, which may have different hard drives, the newly cloned system may get stuck at root disk mirroring during reboot.

Workaround: Create the Flash archive on a golden host with no encapsulated root disks. Run `vxunroot` to clean up the mirrored root disks before you create the Flash archive.

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

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

Workaround: Do one of the following:

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

Erroneous resstatechange trigger warning [2277819]

You may encounter the following warning when you restart resources:

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

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

Upgrade with zones installed on CFS is not supported if CFS is under VCS control (3322276)

If CFS is under VCS control, then upgrade with zones installed on CFS is not supported if you perform phased upgrade.

Workaround: Unmount the CFS before performing the phased upgrade. After the upgrade is complete, re-mount the CFS and reinstall the zone(s).

Rolling upgrade may encounter a problem if open volumes from different disk groups have the same name (3326196)

When you perform a rolling upgrade, the installer may block the rolling upgrade even if all the open volumes are under VCS control. This may occur if there are volumes with the same name under different disk groups although they are not mounted.

Workaround: Avoid creating volumes from different disk groups with the same name. If they already exist, umount all the VxFS mount points. After the upgrade is finished, remount the volumes.

Upgrade fails because there is zone installed on the VxFS file system which is offline. The packages in the zone are not updated. (3319753)

If the zone installed on VxFS file system is under VCS control, and the VxFS file system is in offline state, the upgrade fails because it's not able to update the packages in the zones.

Workaround:

Check the status of the mounted file system which has the zones on it. If the file system is offline, you need to first bring it online, then do the upgrade, so that the packages in the local zone can be updated.

A Storage Foundation ABE upgrade fails while upgrading from 6.1PR1 to 6.1 (3325054)

If you perform an upgrade from SFHA or SFRAC 6.1 PR1 to 6.1 in a Solaris Alternate Boot Environment (ABE), you may see the following error when you uninstall 6.1 PR1:

```
pkg: An unexpected error happened during uninstall:
[('/mnt/etc/vx/vxesd/vxesd.socket',
'/mnt/var/pkg/lost+found/etc/vx/vxesd-20130927T101550Z/vxesd.socket',
"[Errno122] Operation not supported on transport endpoint:
'/mnt/etc/vx/vxesd/vxesd.socket'")]
```

Workaround: Remove */mnt*`/etc/vx/vxesd/vxesd.socket`, where */mnt* is the ABE mountpoint. After you remove the socket, you can successfully uninstall 6.1 PR1 and install 6.1.

If you choose to upgrade nodes without zones first, the rolling upgrade or phased upgrade is not blocked in the beginning, but fails later (3319961)

If you choose to upgrade nodes without zones first, the rolling upgrade or phased upgrade is not blocked in the beginning, but fails later when you start to upgrade the nodes that have zones installed.

This issue occurs in the following scenarios:

- A zone is installed on a Cluster File System (CFS) on one of the nodes.
- A node is installed on a Veritas File System (VxFS) on one of the nodes, and node is under Symantec Cluster Server (VCS) control.

Workaround:

1. Before you upgrade, uninstall the zones on the nodes which have zones installed. Enter:.

   ```
   zoneadm -z zonename uninstall
   ```

2. Run the installer to run the upgrade.
3. After the upgrade completes, reinstall the zones.

The vxdisksetup command fails to initialize disks in cdsdisk format for disks in logical domains greater than 1 TB (2557072)

The `vxdisksetup` command fails to initialize disks in `cdsdisk` format for disks in logical domains greater than 1 TB. This issue is due to an Oracle VM Server

command which fails when the number of partitions in the GUID partition table (GPT) label is greater than 9. The `cdsdisk` format requires at least 128 partitions to be compatible with Linux systems.

Workaround: There is no workaround for this issue.

Rolling upgrade of VCS from pre-6.0 versions fails with CP server in secure mode [3262900]

If the CP server is configured in secure mode, rolling upgrade of VCS from versions lower than 6.0 to 6.1 is not supported. Since the `vxcpserv` process is not compatible with shared authentication, CP server service group fails to come online after performing phase 1 of the rolling upgrade.

Workaround: Use full upgrade or phased upgrade instead of rolling upgrade.

Upgrade or uninstallation of Symantec Storage Foundation HA may encounter module unload failures

When you upgrade or uninstall Symantec Storage Foundation HA, some modules may fail to unload with error messages similar to the following messages:

```
llt failed to stop on node_name
gab failed to stop on node_name
```

The issue may be observed on any one or all the nodes in the sub-cluster.

Workaround: After the upgrade or uninstallation completes, follow the instructions provided by the installer to resolve the issue.

Installing VRTSvlic package on Solaris system with local zones displays error messages [2555312]

If you try to install VRTSvlic package on a Solaris system with local zones in installed state, the system displays the following error messages:

```
cp: cannot create /a/sbin/vxlicinst: Read-only file system
cp: cannot create /a/sbin/vxlicrep: Read-only file system
cp: cannot create /a/sbin/vxlictest: Read-only file system
```

Workaround: On the Solaris system, make sure that all non-global zones are started and in the running state before you install the VRTSvlic package.

VRTSvcsea package cannot be uninstalled from alternate disk in manual live upgrade [2481391]

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

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

VCS Zone users must be added after upgrade to VCS 6.0 or later

If you upgrade your configuration containing Zone resources to VCS 6.0 or later from:

- VCS 5.1SP1RP1 or later VCS releases with DeleteVCSZoneUser attribute of Zone agent set to 1
- VCS 5.1SP1 or ealier VCS releases

You may see the following issue.

Zone agent offline/clean entry points delete VCS Zone users from configuration. After upgrade to VCS 6.0, VCS Zone users need to be added to the configuration. VCS Zone users can be added by running `hazonesetup` utility with new syntax after upgrade. See the *Symantec Storage Foundation and High Availability Solutions Virtualization Guide* for Solaris for more information on `hazonesetup` utility and see the *Symantec Storage Foundation and High Availability Solutions Virtualization Guide* for Solaris.

VCS installation with CPI fails when a non-global zone is in installed state and zone root is not mounted on the node (2731178)

On Solaris 10, CPI tries to boot a zone in installed state during installation/ or uninstallation. The boot fails if the underlying storage for zone root is not imported and mounted onto the node, causing the installation or uninstallation to fail.

Workaround: Make sure that the non-global zones are in running or configured state when CPI is invoked for installation or uninstallation.

Log messages are displayed when VRTSvcs is uninstalled on Solaris 11 [2919986]

The following message is displayed when you uninstall VRTSvcs package on Solaris 11 OS.

```
The following unexpected or editable files and directories were
salvaged while executing the requested package operation; they
have been moved to the displayed location in the image:

  var/VRTSvcs/log -> /var/pkg/lost+found/var/VRTSvcs/log-20111216T122049Z
  var/VRTSvcs/lock -> /var/pkg/lost+found/var/VRTSvcs/lock-20111216T122049Z
  var/VRTSvcs -> /var/pkg/lost+found/var/VRTSvcs-20111216T122049Z
  etc/VRTSvcs/conf/config
->/var/pkg/lost+found/etc/VRTSvcs/conf/config-20111216T122049Z
```

You can safely ignore this message as this is an expected behavior of IPS packaging. The files mentioned in the above message are not part of the package. As a result, uninstallation moves them to `/var/pkg/lost+found` directory.

Cluster goes into `STALE_ADMIN_WAIT` state during upgrade from VCS 5.1 to 6.1 [2850921]

While performing a manual upgrade from VCS 5.1 to VCS 6.1, cluster goes in `STALE_ADMIN_WAIT` state if there is an entry of DB2udbTypes.cf in main.cf.

Installation of VRTSvcsea package in VCS 5.1 creates a symbolic link for `Db2udbTypes.cf` file inside `/etc/VRTSvcs/conf/config` directory which points to `/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf`. During manual upgrade, the VRTSvcsea package for VCS 5.1 gets removed, which in turn removes the symbolic link for file `Db2udbTypes.cf` inside `/etc/VRTSvcs/conf/config` directory. After the complete installation of VRTSvcsea for VCS 6.1, because of absence of file `Db2udbTypes.cf` inside `/etc/VRTSvcs/conf/config`, cluster goes into STALE ADMIN WAIT state.

Workaround: Manually copy `DB2udbTypes.cf` from `/etc/VRTSagents/ha/conf/Db2udb` directory to the `/etc/VRTSvcs/conf/config` directory after the manual upgrade before starting HAD.

Verification of Oracle binaries incorrectly reports as failed during Oracle Grid Infrastructure installation

The verification of Oracle binaries may incorrectly report as failed during the Oracle Grid Infrastructure installation using the SFRAC installer. The message is erroneously reported due to a break in passwordless SSH communication. The SSH communication fails because execution of the `root.sh` script changes the owner of the operating system root directory to the grid user directory.

Symantec Cluster Server known issues

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

Operational issues for VCS

This section describes the Operational known issues for VCS.

Some VCS components do not work on the systems where a firewall is configured to block TCP traffic [3545338]

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 *Symantec 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.

Stale legacy_run services seen when VCS is upgraded to support SMF [2431741]

If you have VCS 5.0MPx installed on a Solaris 10 system, VCS uses RC scripts to manage starting services. If you upgrade VCS to any version that supports SMF for VCS, you see stale legacy_run services for these RC scripts in addition to the SMF services.

Workaround: There are two ways to remove these legacy services:

- Open svccfg console using `svccfg -s smf/legacy_run` and delete the legacy services.
  For example:

```
svccfg -s smf/legacy_run
svc:/smf/legacy_run> listpg *
rc2_d_S70llt    framework     NONPERSISTENT
rc2_d_S92gab    framework     NONPERSISTENT
svc:/smf/legacy_run> delpg rc2_d_S70llt
svc:/smf/legacy_run> delpg rc2_d_S92gab
svc:/smf/legacy_run> exit
```

- Reboot the system.

The hastop -all command on VCS cluster node with AlternateIO resource and StorageSG having service groups may leave the node in LEAVING state [2523142]

On a VCS cluster node with AlternateIO resource configured and StorageSG attribute contain service groups with Zpool, VxVM or CVMVolDG resources, `hastop -local` or `hastop -all` commands may leave the node in "LEAVING" state.

This issue is caused by lack of dependency between service group containing LDom resource and service groups containing storage resources exported to logical domain in alternate I/O domain scenarios. In this scenario VCS may attempt to stop the storage service groups before stopping logical domain which is using the resources.

Workaround: Stop the LDom service group before issuing `hastop -local` or `hastop -all` commands.

Missing characters in system messages [2334245]

You may see missing characters, especially in long system messages in response to certain commands.

Workaround: No workaround.

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.

After OS upgrade from Solaris 10 update 8 or 9 to Solaris 10 update 10 or 11, Samba server, SambaShare and NetBios agents fail to come online [3321120]

On Solaris 10 update 8 and update 9, default path of Samba binaries is `/usr/sfw/sbin/smbd` and default samba configuration file location is `/etc/sfw/smb.conf`. On Solaris 10 update 10 and update 11, the default path of Samba binaries is changed to `/usr/sbin/smbd` and default Samba configuration file location is `/etc/samba/smb.conf`. Therefore, after OS upgrade from Solaris 10 update 8 or update 9 to Solaris 10 update 10 or update 11, Samba server, SambaShare and NetBios agents are unable to locate binaries and configuration file.

Workaround: After the OS upgrade from Solaris 10 update 8 or update 9 to Solaris 10 update 10 or update 11, update the SambaTopDir and ConfFile attributes of the Samba server resources appropriately to reflect the correct location.

CP server does not allow adding and removing HTTPS virtual IP or ports when it is running [3322154]

CP server does not support adding and removing HTTPS virtual IPs or ports while the CP server is running. However, You can add or remove the IPM virtual IPs or ports.

Workaround: No workaround. If you want to add a new virtual IP for HTTPS, you must follow the entire manual procedure for generating HTTPS certificate for the CP server (server.crt), as documented in the *Symantec Cluster Server Installation Guide*.

CP server does not support IPv6 communication with HTTPS protocol [3209475]

CP server does not support IPv6 communication when using the HTTPS protocol. This implies that in VCS 6.1, CP servers listening on HTTPS can only use IPv4. As a result, VCS 6.1 fencing clients can also use only IPv4.

Workaround: No workaround.

CP server service group fails to come online with the default database path after the CP server is upgraded from 6.0 to 6.1 on a multi-node cluster [3326639]

If the CP server is configured on a multi-node cluster before the upgrade with security enabled, you must reconfigure the CP server after the CP server upgrade. If you reuse the old credentials with the old database path, the CP server service group does not come online. Since the default database paths of CP server in 6.0 and 6.1 are different, reusing the old credentials and default database path prevents the CP server service group from coming online.

Workaround:

If the CP server multi-node cluster is configured with security enabled and if the old credentials such as database path are expected to be reused in reconfiguration of the CP server after the upgrade of the CP server, use the same database path before and after the upgrade.

Issues related to the VCS engine

This section describes the knonw issues about the VCS engine.

VCS enters into admin_wait state when Cluster Statistics is enabled with load and capacity defined [3199210]

VCS enters into admin_wait state when started locally if:

1. Statistics attribute value is set to Enabled, which is its default value.
2. Group Load and System Capacity values are defined in units in `main.cf`.

Workaround:

1. Stop VCS on all nodes in the cluster.
2. Perform any one of the following steps:
   - Edit the `main.cf` on one of the nodes in the cluster and set the Statistics attribute to Disabled or MeterHostOnly.
   - Remove the Group Load and System Capacity values from the `main.cf`.
3. Run `hacf -verify` on the node to verify that the configuration is valid.
4. Start VCS on the node and then on the rest of the nodes in the cluster.

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.

Missing host names in engine_A.log file (1919953)

The GUI does not read the `engine_A.log` file. It reads the `engine_A.ldf` file, gets the message id from it, and then queries for the message from the bmc file of the appropriate locale (Japanese or English). The bmc file does not have system names present and so they are read as missing.

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

The `hacf -cmdtocf` 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 `hacf -cmdtocf` command.

Character corruption observed when executing the uuidconfig.pl -clus -display -use_llthost command [2350517]

If password-less ssh/rsh is not set, the use of `uuidconfig.pl` command in non-English locale may print garbled characters instead of a non-English string representing the Password prompt.

Workaround: No workaround.

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 dependancy is disabled then the other resources do not come online and the following message is displayed:

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

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

NFS resource goes offline unexpectedly and reports errors when restarted [2490331]

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

Workaround: Terminate the agent process.

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

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

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

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

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

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

If secure and non-secure WAC are connected the engine_A.log receives logs every 5 seconds [1919933]

Two WACs in a global service group 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 a global service group.

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.

Oracle service group faults on secondary site during failover in a disaster recovery scenario [2653704]

Oracle service group fails to go online in the DR site when disaster strikes the primary site. This happens if the AutoFailover attribute on the Service Group is set to 1 and when the corresponding service group's FireDrill is online in the DR site. Firedrill Service group may remain ONLINE on the DR site.

Workaround: If the service group containing the Oracle (or any database) resource faults after attempting automatic DR failover while FireDrill is online in the DR site, manually offline the FireDrill Service Group. Subsequently, attempt the online of the Oracle Service Group in the DR site.

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

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

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

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

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

Workaround: No workaround.

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

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

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

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

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

GCO clusters remain in INIT state [2848006]

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

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

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

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

The high availability 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 the high availability command with same non-root user (this will pass).

Startup trust failure messages in system logs [2721512]

If you configure a cluster with security enabled, there might be some messages logged in system message logs related to Symantec authentication. These messages can be ignored and have no effect on functionality.

Workaround: No workaround.

Running the `-delete -keys` command for any scalar attribute causes core dump [3065357]

Running the `-delete -keys` command for any scalar attribute is not a valid operation and must not be used. However, any accidental or deliberate use of this command may cause engine to core dump.

Workaround: No workaround.

Agent reports incorrect state if VCS is not set to start automatically and `utmp` file is empty before VCS is started [3326504]

If you have not configured VCS to start automatically after a reboot and have tmptied the `utmp` file before starting VCS manually with the hastart command, some agents might report an incorrect state.

The `utmp` file (file name may differ on different operating systems) is used to maintain a record of the restarts done for a particular machine. The checkboot utility used by `hastart` command uses the functions provided by the OS which in turn use the `utmp` file to find if a system has been restarted so that the temporary files for various agents can be deleted before agent startup. If OS functions do not return correct value, High Availability Daemon (HAD) starts without deleting the stale agent files. This might result in some agents reporting incorrect state.

Workaround: If a user wishes to delete the `utmp` file this should be done only when VCS is already running or the customer should delete the temporary files in`/var/VRTSvcs/lock/volatile/` manually before starting VCS.

Site preference fencing policy value fails to set on restart of a site-aware cluster [3380584]

If you restart VCS on a site-aware cluster, the PreferredFencingPolicy fails to reset to the value 'Site' assigned to it before the restart.

Workaround: Reassign the fencing policy value manually to the cluster.

Issues related to the bundled agents

This section describes the known issues of the bundled agents.

After reboot of ldom, Apache service group fails to get online [3543906]

When you restart the ldom, the `/var/run` directory is emptied, hence the directory for Apache PID file is not present on the system. For example, `/var/run/apache2` is missing. Therefore the PID file cannot be created. Consequently a startup command of Apache agent fails with the following error message in Apache error logs:

```
[error] (2)No such file or directory: \
could not create /var/run/apache2/httpd.pid
```

The Startup command is `$HttpdDir/httpd -f` *`Config File`* `-k start`.

Since the startup command fails outside VCS control, the Apache agent fails to bring Apache service group online.

Workaround: Write a preonline trigger for Apache service Group to create the required directory, for example:

```
mkdir /var/run/apache2
```

Refer to *Symantec Cluster Server 6.1 Administrator's Guide* for more details on VCS triggers.

Entry points that run inside a zone are not cancelled cleanly [1179694]

Cancelling entry points results in the cancellation of only the `zlogin` process. The script entry points that run inside a zone are forked off using the `zlogin` command. However, the `zlogin` command forks off an `sh` command, which runs in the context of the Solaris zone. This shell process and its family do not inherit the group id of the `zlogin` process, and instead get a new group id. Thus, it is difficult for the agent framework to trace the children or grand-children of the shell process, which translates to the cancellation of only the `zlogin` process.

Workaround: Oracle must provide an API or a mechanism to kill all the children of the `zlogin` process that was started to run the entry point script in the local-zone.

Solaris mount agent fails to mount Linux NFS exported directory [2098333]

The Solaris mount agent mounts the mount directories. At this point, if it tries to mount a Linux NFS exported directory, the mount fails showing the following error:

```
nfs mount: mount: <MountPoint>: Not owner
```

This is due to system NFS default version mismatch between Solaris and Linux.

Workaround: Configure `MountOpt` attribute in mount resource and set `vers=3` for it.

Example

```
root@north $ mount -F nfs south:/test /logo/
nfs mount: mount: /logo: Not owner
root@north $
Mount nfsmount (
                MountPoint = "/logo"
                BlockDevice = "south:/test"
                FSType = nfs
                MountOpt = "vers=3"
               )
```

The zpool command runs into a loop if all storage paths from a node are disabled [2010892]

The Solaris Zpool agent runs `zpool` commands to import and export zpools. If all paths to the storage are disabled, the zpool command does not respond. Instead, the zpool export command goes into a loop and attempts to export the zpool. This

continues till the storage paths are restored and zpool is cleared. As a result, the offline and clean procedures of Zpool Agent fail and the service group cannot fail over to the other node.

Workaround: You must restore the storage paths and run the zpool clear command for all the pending commands to succeed. This will cause the service group to fail over to another node.

Zone remains stuck in down state if tried to halt with file system mounted from global zone [2326105]

If zone halts without unmounting the file system, the zone goes to down state and does not halt with the `zoneadm` commands.

Workaround: Unmount the file system manually from global zone and then halt the zone. For VxFS, use following commands to unmount the file system from global zone.

To unmount when VxFSMountLock is 1

```
umount -o mntunlock=VCS <zone root path>/<Mount Point>
```

To forcefully unmount when VxFSMountLock is 1:

```
# umount -f -o mntunlock=VCS <zone root path>/<Mount Point>
```

To unmount when VxFSMountLock is 0:

```
# umount <zone root path>/<Mount Point>
```

To forcefully unmount when VxFSMountLock is 0:

```
# umount -f <zone root path>/<Mount Point>
```

To halt the zone, use following command:

```
# zoneadm -z <zone_name>  halt
```

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

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

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

The zpool commands hang and remain in memory till reboot if storage connectivity is lost [2368017]

If the FailMode attribute of `zpool` is set to continue or wait and the underlying storage is not available, the `zpool` commands hang and remain in memory until the next reboot.

This happens when storage connectivity to the disk is lost, the `zpool` commands hang and they cannot be stopped or killed. The zpool commands run by the monitor entry point remains in the memory.

Workaround: There is no recommended workaround for this issue.

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

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

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

Offline of zone resource may fail if `zoneadm` is invoked simultaneously [2353541]

Offline of zone EP uses `zoneadm` command to offline a zone. Therefore, if `zoneadm` is invoked simultaneously for multiple zones, the command may fail. This is due to Oracle bug 6757506 that causes a race condition between multiple instances of `zoneadm` command and displays the following message:

```
zoneadm: failed to get zone name: Invalid argument
```

Workaround: No workaround.

Password changed while using `hazonesetup` script does not apply to all zones [2332349]

If you use the same user name for multiple zones, updating password for one zone does not updated the password of other zones.

Workaround: While updating password for VCS user which is used for multiple zones, update password for all the zones.

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.

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 local Zone on Solaris

Workaround: No workaround.

Share resource goes offline unexpectedly causing service group failover [1939398]

Share resource goes offline unexpectedly and causes a failover when NFSRestart resource goes offline and UseSMF attribute is set to 1 (one).

When NFSRestart resource goes offline, NFS daemons are stopped. When UseSMF attribute is set to 1, the exported file systems become unavailable, hence Share resource unexpectedly goes offline.

Workaround: Set the value of ToleranceLimit of Share resource to a value more than 1.

Mount agent does not support all scenarios of loopback mounts [2938108]

For a mount point under VCS control, you can create loop back mounts for the mount point. For example, mount point /mntpt is mounted on /a as loop back mount and /a is mounted on /b as loop back mount, then offline and online of the mount resource fails.

Workaround: Mount the mount point /mntpt on /b as loop back mount.

Invalid Netmask value may display code errors [2583313]

If you specify invalid Netmask value for the IP resource attribute, you may see the code errors similar to the following when you try to online the resource.

```
==================================
Illegal hexadecimal digit 'x' ignored at
/opt/VRTSperl/lib/site_perl/5.12.2/Net/Netmask.pm line 78.
ifconfig: <Netmask_value>: bad address
==========================================
```

Workaround: Make sure you specify a valid Netmask value.

Zone root configured on ZFS with ForceAttach attribute enabled causes zone boot failure (2695415)

On Solaris 11 system, attaching zone with `-F` option may result in zone boot failure if zone root is configured on ZFS.

Workaround: Change the ForceAttach attribute of Zone resource from 1 to 0. With this configuration, you are recommended to keep the default value of DetachZonePath as 1.

Error message is seen for Apache resource when zone is in transient state [2703707]

If the Apache resource is probed when the zone is getting started, the following error message is logged:

```
Argument "VCS ERROR V-16-1-10600 Cannot connect to VCS engine\n"
isn't numeric in numeric ge (>=) at /opt/VRTSvcs/bin/Apache/Apache.pm
line 452.
VCS ERROR V-16-1-10600 Cannot connect to VCS engine
LogInt(halog call failed):TAG:E:20314 <Apache::ArgsValid> SecondLevel
MonitorTimeOut must be less than MonitorTimeOut.
```

Workaround: You can ignore this message. When the zone is started completely, the `halog` command does not fail and Apache agent monitor runs successfully.

Monitor falsely reports NIC resource as offline when zone is shutting down (2683680)

If a NIC resource is configured for an Exclusive IP zone, the NIC resource is monitored inside the zone when the zone is functional. If the NIC monitor program is invoked when the zone is shutting down, the monitor program may falsely report the NIC resource as offline. This may happen if some of the networking services are offline but the zone is not completely shut down. Such reports can be avoided if you override and set the ToleranceLimit value to a non-zero value.

Workaround: When a NIC resource is configured for an Exclusive IP zone, you are recommended to set the ToleranceLimit attribute to a non-zero value.

Calculate the ToleranceLimit value as follows:

Time taken by a zone to completely shut down must be less than or equal to NIC resource's MonitorInterval value + (MonitorInterval value x ToleranceLimit value).

For example, if a zone take 90 seconds to shut down and the MonitorInterval for NIC agent is set to 60 seconds (default value), set the ToleranceLimit value to 1.

Apache resource does not come online if the directory containing Apache pid file gests deleted when a node or zone restarts (2680661)

The directory in which Apache http server creates PidFile may get deleted when a node or zone restarts. Typically the PidFile is located at `/var/run/apache2/httpd.pid`. When the zone reboots, the `/var/run/apache2` directory may get removed and hence the http server startup may fail.

Workaround: Make sure that Apache http server writes the PidFile to an accessible location. You can update the PidFile location in the Apache http configuration file (For example: `/etc/apache2/httpd.conf`).

Online of LDom resource may fail due to incompatibility of LDom configuration file with host OVM version (2814991)

If you have a cluster running LDom with different OVM versions on the hosts, then the LDom configuration file generated on one host may display error messages when it is imported on the other host with a different OVM version. Thus, the online of LDom resource may also fail.

For example, if you have a cluster running LDom with OVM versions 2.2 on one and OVM 2.1 on the other node, the using XML configuration generated on the host with OVM 2.2 may display errors when the configuration is imported on the host with OVM 2.1. Thus, the online of LDom resource fails.

The following error message is displayed:

```
ldm add-domain failed with error Failed to add device
/ldom1/ldom1 as ld1_disk1@primary-vds0 because this device
```

```
is already  exported on LDom primary. Volume ld1_disk1
already exists in vds primary-vds0.
```

Workaround: If the CfgFile attribute is specified, ensure that the XML configuration generated is compatible with the OVM version installed on the nodes.

Online of IP or IPMultiNICB resource may fail if its IP address specified does not fit within the values specified in the allowed-address property (2729505)

While configuring an IP or IPMultiNICB resource to be run in a zone, if the IP address specified for the resource does not match the values specified in the **allowed-address** property of the zone configuration, then the online of IP resource may fail. This behavior is seen only on Solaris 11 platform.

Workaround: Ensure that the IP address is added to **allowed-address** property of the zone configuration.

Application resource running in a container with PidFiles attribute reports offline on upgrade to VCS 6.0 or later [2850927]

Application resource configured to run in a container configured with PidFiles attribute reports state as offline after upgrade to VCS 6.0 or later versions.

When you upgrade VCS from lower versions to 6.0 or later, if application resources are configured to run in a container with monitoring method set to PidFiles, then upgrade may cause the state of the resources to be reported as offline. This is due to changes introduced in the Application agent where if the resource is configured to run in a container and has PidFiles configured for monitoring the resource then the value expected for this attribute is the pathname of the PID file relative to the zone root.

In releases prior to VCS 6.0, the value expected for the attribute was the pathname of the PID file including the zone root.

For example, a configuration extract of an application resource configured in VCS 5.0MP3 to run in a container would appear as follows:

```
Application apptest (
User = root
StartProgram = "/ApplicationTest/app_test_start"
StopProgram = "/ApplicationTest/app_test_stop"
PidFiles = {
      "/zones/testzone/root/var/tmp/apptest.pid" }
ContainerName = testzone
)
```

Whereas, the same resource if configured in VCS 6.0 and later releases would be configured as follows:

```
Application apptest (
User = root
StartProgram = "/ApplicationTest/app_test_start"
StopProgram = "/ApplicationTest/app_test_stop"
PidFiles = {
      "/var/tmp/apptest.pid" }
)
```

Note: The container information is set at the service group level.

Workaround: Modify the PidFiles pathname to be relative to the zone root as shown in the latter part of the example.

```
# hares -modify apptest PidFiles /var/tmp/apptest.pid
```

NIC resource may fault during group offline or failover on Solaris 11 [2754172]

When NIC resource is configured with exclusive IP zone, NIC resource may fault during group offline or failover. This issue is observed as zone takes long time in shutdown on Solaris 11. If NIC monitor is invoked during this window, NIC agent may treat this as fault.

Workaround: Increase ToleranceLimit for NIC resource when it is configured for exclusive IP zone.

NFS client reports error when server is brought down using `shutdown` command [2872741]

On Solaris 11, when the VCS cluster node having the NFS share service group is brought down using `shutdown` command, NFS clients may report "Stale NFS file handle" error. During shutdown, the SMF service svc:/network/shares un-shares all the shared paths before taking down the virtual IP. Thus, the NFS clients accessing this path get stale file handle error.

Workaround: Before you shutdown the VCS cluster node, disable the svc:/network/shares SMF service, so that only VCS controls the un-sharing of the shared paths during the shutdown operation.

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 resource) are still online on the failing node. The disk group on the failing node may also get disabled but IP resource on the same node continues to be online.

Workaround: Configure the preonline trigger for the service groups containing DiskGroup resource 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/` as T0preonline_ipc:

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

2 Enable the preonline trigger for the service group.

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

The CoordPoint agent faults after you detach or reattach one or more coordination disks from a storage array (3317123)

After you detach or reattach a coordination disk from a storage array, the CoordPoint agent may fault because it reads an older value stored in the I/O fencing kernel module.

Workaround: Run the `vxfenswap` utility to refresh the registration keys on the coordination points for both server-based I/O fencing and disk-based I/O fencing. But, even if the registrations keys are not lost, you must run the `vxfenswap` utility to refresh the coordination point information stored in the I/O fencing kernel module.

For more information on refreshing registration keys on the coordination points for server-based and disk-based I/O fencing, refer to the *Symantec Cluster Server Administrator's Guide*.

Mount resource does not support spaces in the MountPoint and BlockDevice attribute values [3335304]

Mount resource does not handle intermediate spaces in the configured MountPoint or BlockDevice attribute values.

Workaround: No workaround.

Issues related to the VCS database agents

This section describes the known issues about VCS database agents.

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

The ASMInstAgent does not support having pfile or 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 [2125453]

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

NOFAILOVER action specified for certain Oracle errors

The Symantec 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 *Symantec 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.

ASMInstance resource monitoring offline resource configured with OHASD as application resource logs error messages in VCS logs [2846945]

When the Oracle High Availability Services Daemon (OHASD) is configured as an application resource to be monitored under VCS and if this resource is offline on the failover node then the ASMInstance resource in the offline monitor logs the following error messages in the VCS logs:

```
ASMInst:asminst:monitor:Cluster Synchronization Service
process is not running.
```

Workaround: Configure the application in a separate parallel service group and ensure that the resource is online.

Issues related to the agent framework

This section describes the known issues about the agent framework.

Agent may fail to heartbeat under heavy load [2073018]

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

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

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

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

Agent framework cannot handle leading and trailing spaces for the dependent attribute (2027896)

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

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

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

In rare cases, the agent framework does not detect if all service threads hang inside a C-based 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 [3208239]

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.

CFSMount agent may fail to heartbeat with VCS engine and logs an error message in the engine log on systems with high memory load [3060779]

On a system with high memory load, CFSMount agent may fail to heartbeat with VCS engine resulting into V-16-1-53030 error message in the engine log.

VCS engine must receive periodic heartbeat from CFSMount agent to ensure that it is running properly on the system. The heartbeat is decided by AgentReplyTimeout attribute. Due to high CPU usage or memory workload (for example, swap usage greater than 85%), agent may not get enough CPU cycles to schedule. This causes heartbeat loss with VCS engine and as a result VCS engine terminates the agent and starts the new agent. This can be identified with the following error message in the engine log:

```
V-16-1-53030 Termination request sent to CFSMount
agent process with pid %d
```

Workaround: Increase the AgentReplyTimeout value and see if CFSMount agent becomes stable. If this does not resolve the issue then try the following workaround. Set value of attribute NumThreads to 1 for CFSMount agent by running following command:

```
# hatype -modify CFSMount NumThreads 1
```

Even after the above command if CFSMount agent keeps on terminating, report this to Symantec support team.

Issues related to VCS in Japanese locales

This section covers the known issues about VCS in a Japanese locale.

The hares -action command displays output in English [1786742]

The `hares -action` command incorrectly displays output in English.

Character corruption issue

Character corruption occurs if installer is run with HIASCII option on French locale. [1539754, 1539747]

Workaround: No workaround.

Messages inside the zone are not localized [2439698]

Locale is not set correctly for Solaris zone. Therefore, you may not see localized messages inside the zone.

Workaround: No workaround.

System messages having localized characters viewed using `hamsg` may not be displayed correctly

If you use `hamsg` to view system messages, the messages containing a mix of English and localized characters may not be displayed correctly. [2405416]

Workaround: No workaround. However, you can view English messages in the VCS log file.

Standalone utilities display output in English [2848012]

The following utilities display output in English:

- `-haping`
- `-hamultinicb`
- `-haipswitch`

Workaround: No workaround.

English error messages displayed by the gcoconfig wizard [3018221]

Whenever the gcoconfig wizard calls a command internally, the messages from that command are displayed in English.

Workaround: No workaround.

Issues related to global clusters

This section describes the known issues about 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.

Issues related to Intelligent Monitoring Framework (IMF)

This section describes the known issues of Intelligent Monitoring Framework (IMF).

Package verification check for VRTSamf may fail [3548433]

If you perform a verification check on the VRTSamf package by running the `pkg verify VRTSamf` commands, it may fail with the following error message:

```
PACKAGE
STATUS
pkg://Symantec/VRTSamf
ERROR
        file: etc/default/amf
                Size: 235 bytes should be 234
                Hash: effd9981aaba8d8f52f9417186e27fd18e42544e should be
a2de7947bd99603711ac7d81eef0b1763dea77d7
```

Workaround: No workaround. You can ignore the error as it does not affect AMF functionalities.

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.

IMF does not fault zones if zones are in ready or down state [2290883]

IMF does not fault zones if zones are in ready or down state.

IMF does not detect if zones are in ready or down state. In Ready state, there are no services running inside the running zones.

Workaround: Offline the zones and then restart.

IMF does not detect the zone state when the zone goes into a maintenance state [2535733]

IMF does not detect the change in state. However, the change in state is detected by Zone monitor in the next cycle.

Workaround: No workaround.

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 one registered with the AMF.

Direct execution of the `linkamf` command 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.

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.

VCS engine displays error for cancellation of reaper when Apache agent is disabled [3043533]

When `haimfconfig` script is used to disable IMF for one or more agents, the VCS engine logs the following message in the engine log:

```
AMF imf_getnotification ERROR V-292-2-193
Notification(s) canceled for this reaper.
```

This is an expected behavior and not an issue.

Workaround: No workaround.

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

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

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

For example:

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

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

Agent cannot become IMF-aware with agent directory and agent file configured [2858160]

Agent cannot become IMF-aware if Agent Directory and Agent File are configured for that agent.

Workaround: No workaround.

AMF may panic the system if it receives a request to unregister an already unregistered resource [3333913]

If AMF encounters any internal error, it unregisters all the resources which it cannot support. During such an event, if any agent calls unregister for one of such resources, AMF may panic the machine.

Workaround: No Workaround.

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.

Unable to log on to secure VCS clusters on Solaris 11 using Java GUI (2718943)

Connecting to secure clusters deployed on Solaris 11 systems using VCS Java GUI is not supported in VCS 6.0PR1. The system displays the following error when you attempt to use the Java GUI:

```
Incorrect username/password
```

Workaround: No workaround.

Issues related to live migration

This section describes the known issues about live migration.

Operating system in guest domain with multiple IO services hangs when guest migrates back [3127470]

Operating system inside the guest domain hangs when the guest domain is provided with IO services from multiple IO domains but not from primary domain and guest domain is migrated to another node and back to the source node.

Workaround: Make sure that firmware of the physical system is upgraded to latest version.

Issues related to virtualization

This section describes the known issues about virtualization.

Locale message displayed on Solaris 11 system for solaris10 brand zones [2695394]

When you run the `zlogin` command on a Solaris 11 system, the system logs the following error message:

```
Could not set locale correctly.
```

The default locale for Solaris 11 is en_US.UTF-8 and that of Solaris 10 is C. With solaris10 brand zone, en_US.UTF-8 is not installed inside the zone by default. Therefore, the error message is logged.

Workaround: This message can be safely ignored as there is no functionality issue. To avoid this message, install en_US.UTF-8 locale on solaris10 brand zone.

Symantec Dynamic Multi-Pathing known issues

This section describes the known issues in this release of Symantec Dynamic Multi-pathing (DMP).

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.

Creating a zpool fails with a incorrect disk size error (2277875)

When the tunable parameter dmp_native_support is turned on, creating a zpool on DMP devices may fail with the following error:

```
one or more devices is less than the minimum size (64 M)
```

This error may occur even if the device size is greater than the required minimum size.

Workaround:

To resolve this issue, use one of the following commands:

- `# vxdisk scandisks`
- `# format -e dmp_device`

DMP aggregates EFI labelled LUNS to a 0_0 disk (2558408)

While performing `vxdiskunsetup` of some luns, if you format and label the disks as EFI, all the EFI labelled luns are aggregated to a 0_0 disk.

Workaround:

When changing the label of a disk from SMI to EFI, or vice-versa, Symantec recommends that the label be changed on all accessible paths to a disk. That is, use the `format -e` command to stamp the new label on all accessible paths. For Active/Passive (A/P) class of arrays, this should be done only on the active paths. For other arrays, all paths should be labeled.

Symantec also recommends the installation of the patch provided by Oracle for EFI label issues (IDR144101-01 or IDR144249-01 or release kernel patch 142909-17). If this patch is installed, you can run the `format -e` command only on one path. After that, perform a read operation (such as `dd if=/dev/rdsk/<path> of=/dev/null count=1`) on the other accessible paths to propagate the label.

Splitting a mirror from a zpool causes a core dump (2273367)

The following operation to split a mirror from a zpool fails:

```
# zpool split my_pool new_pool mirror
```

This issue is an Oracle issue with zpool. This issue occurs whether DMP is controlling the devices or not. That is, whether the dmp_native_support tunable is on or off.

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.

ZFS pool creation on a DMP device fails when the LUN size is between 1 TB and 2TB (2010919)

Creating a ZFS pool on a DMP device using the whole disk of size > 1TB and < 2TB that contains a SMI SUN label fails. The issue is that zpool create on a whole disk changes the device label from SMI to EFI. This causes confusion between the OS device paths of the same DMP device due to a bug in the Sun SCSI layer. This is due to SUN BugID: 6912703.

After excluding devices managed by PowerPath from VxVM, the devices still show as DMP devices (2494632)

The issue happens after EMC PowerPath is installed and all devices are under PowerPath control. If you want to maintain the devices under PowerPath control, you use the following command to exclude the device that is managed by PowerPath from VxVM:

```
# vxdmpadm exclude dmpnodename=PowerPath_device _name
```

After system reboot, the PowerPath device still shows as a DMP device, although the device is managed by EMC PowerPath.

Workaround:

This issue is seen only during the first bootup discovery after reboot. To resolve the issue, manually trigger DMP device discovery:

```
# vxdisk scandisks
```

Limitation to DMP support for ZFS root in the LDOM guest (3221944)

DMP support for ZFS root is not supported in the LDOM guest. If DMP meta devices are exported to LDOM and used for root pool, then enabling the dmp_native_support tunable fails with the following error:

```
root@swsx39-v05#vxdmpadm settune dmp_native_support=on
VxVM vxdmpadm ERROR V-5-1-15690 Operation failed for one or more
zpools

VxVM vxdmpadm ERROR V-5-1-15686 The following zpool(s) could not
be migrated as failed to obtain root pool information -

        rpool
```

Where *rpool* specifies the root pool name on the LDOM guest:

DMP supports non-root ZFS in the LDOM guest. You can use DMP devices for non-root ZFS.

Symantec Storage Foundation known issues

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

- Symantec Storage Foundation and High Availability Solutions known issues
- Veritas Volume Manager known issues
- Veritas File System known issues
- Replication known issues
- Symantec Storage Foundation for Databases (SFDB) tools known issues
- Virtualization known issues

Symantec Storage Foundation and High Availability Solutions known issues

This section describes the known issues in this release of Symantec Storage Foundation and High Availability Solutions (SFHA Solutions).

Some dbed DST commands do not work correctly in non-POSIX locales (2138030)

Some dbed DST commands do not work correctly in non-POSIX locale settings.

Workaround: Set the environment variable `LANG=C` systemwide in the `/etc/profile` file.

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.

Boot fails after installing or removing SFHA Solutions packages from a Solaris 9 system to a remote Solaris 10 system (1747640)

The following issue occurs if you install or remove a Storage Foundation package or patch from a Sparc Solaris 9 system to a remote Solaris 10 system, using the `-R` *`rootpath`* option of the `pkgadd`, `patchadd`, `pkgrm` or `patchrm` commands.

Generally, when you install or remove a SFHA Solutions package on a Solaris 10 system, the package scripts update the boot archive. However if the local system is Solaris 9 and the remote system is Solaris 10, the scripts fail to update the boot archive on the Solaris 10 system.

Note: The boot archive is synchronized correctly when you upgrade SFHA Solutions using Solaris Live Upgrade.

Workaround: The workaround is to manually clear the boot archive when you boot the alternate. The SUN boot process detects that the boot archive is out sync and displays instructions for how to correct the situation.

For example:

```
WARNING: The following files in / differ from the boot archive:

     stale //kernel/drv/sparcv9/vxportal
     stale //kernel/drv/vxportal.conf
     stale //kernel/fs/sparcv9/vxfs
     ...
     new     /kernel/drv/vxlo.SunOS_5.10
     new     /kernel/drv/vxlo.conf
     changed /kernel/drv/vxspec.SunOS_5.9
     changed /kernel/drv/vxspec.conf

 The recommended action is to reboot to the failsafe archive to correct
 the above inconsistency. To accomplish this, on a GRUB-based platform,
 reboot and select the "Solaris failsafe" option from the boot menu.
 On an OBP-based platform, reboot then type "boot -F failsafe". Then
 follow the prompts to update the boot archive. Alternately, to continue
 booting at your own risk, you may clear the service by running:
 "svcadm clear system/boot-archive"
```

Oracle 11gR1 may not work on pure IPv6 environment (1819585)

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

Tools like `dbca` may hang during database creation.

Workaround: There is no workaround for this, as Oracle 11gR1 does not fully support pure IPv6 environment. Oracle 11gR2 release may work on a pure IPv6 enviroment, but it has not been tested or released yet.

Sybase ASE version 15.0.3 causes segmentation fault on some Solaris version (1819595)

Sybase ASE 15.0.3 produces segmentation fault on Solaris SPARC 10 Update 6 in a pure IPv6 environment. However, Sybase ASE 15.0.3 works on Solaris SPARC 10 Update 5.

When running Sybase ASE 15.0.3 GA on a pure IPv6 environment on Solaris SPARC 10 Update 6, you may receive a segmentation fault message. For example:

```
Building Adaptive Server 'CDGV240AIPV6':
Writing entry into directory services...
Directory services entry complete.
Building master device...
Segmentation Fault - core dumped
Task failed
Server 'CDGV240AIPV6' was not created.
```

This is a Sybase known issue. You should use Sybase Adaptive Server Enterprise Suite version 15.0.3 ESD 1 that supports Solaris 10 Update 6 or later. For details, refer to the Sybase Product Download Center regarding ESD 1.

An error message is received when you perform off-host clone for RAC and the off-host node is not part of the CVM cluster [1922388, 1834860]

There is a known issue when you try to perform an off-host clone for RAC and the off-host node is not part of the CVM cluster. You may receive a similar error message:

```
Cannot open file /etc/vx/vxdba/rac11g1/.DB_NAME
(No such file or directory).
SFORA vxreptadm ERROR V-81-8847 Cannot get filename from sid
for 'rac11g1', rc=-1.
SFORA vxreptadm ERROR V-81-6550 Could not connect to repository
database.
VxVM vxdg ERROR V-5-1-582 Disk group SNAP_rac11dg1: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac11dg1.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac11dg1 failed.
```

Workaround: Currently there is no workaound for this known issue. However, if the off-host node is part of the CVM cluster, then off-host clone for RAC works fine.

Also the `dbed_vmclonedb` command does not support LOCAL_LISTENER and REMOTE_LISTENER in the `init.ora` parameter file of the primary database.

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.

NFS issues with VxFS Storage Checkpoints (2027492)

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

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

```
Stale NFS file handle
```

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

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

Workaround: There is no workaround for this issue.

Not all the objects are visible in the VOM GUI (1791063, 1821803)

After upgrading SF stack from 5.0MP3RP2 to 5.1, the volumes are not visible under the Volumes tab and the shared diskgroup is discovered as Private and Deported under the Diskgroup 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
  ```

A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a dynamic storage tiering 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 Volume Manager known issues

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

The VRTSvxvm and VRTSaslapm packages do not have the info.classification attribute [3565157]

The Solaris 11 Package Manager uses the `info.classification` package attribute to display packages by category. In 6.1.1, the `info.classification` attribute is not defined for VRTSvxvm and VRTSaslapm for Solaris 11 platform. Consequently the attribute does not appear in the output of the `pkg contents` command. Therefore you may not be able to search the package according to the classification.

Workaround: In order to search the VRTSvxvm or VRTSaslapm packages per category, you can use the `pkginfo` command instead of the `pkg` command.

Package verification may fail for VRTSvxvm during the post-installation checks [3544020]

During the post-installation checks, verification of the VRTSvxvm package may fail with the following error messages:

```
PACKAGE                                 STATUS
pkg://Symantec/VRTSvxvm                 ERROR
   file: kernel/drv/vxio.conf
      Group: 'root (0)' should be 'sys (3)'
      Size: 1116 bytes should be 1094
      Hash: 1327b86150cc8366e317276b should be
e6d042db31f325ed420825189199a31212873b3c
```

Workaround: No workaround. You can ignore it since it does not affect any functionalities.

The `vxdisk resize` command does not claim the correct LUN size on Solaris 11 during expansion of the LUN from array side [2858900]

The `vxdisk resize` command does not claim correct LUN size on Solaris 11 during the expansion of the LUN from array side and fails. This is due to Oracle issue -19603615.

On Solaris 11, the `vxdisk resize` command may exit without errors, and return incorrect LUN size or fail with similar error as follows:

```
bash# vxdisk -g testdg resize disk01 length=8g
VxVM vxdisk ERROR V-5-1-8643 /
Device disk01: resize failed:Operation would block
```

Workaround: There is no workaround available which can work in all the configuration.

In some specific configurations, it works if you run the following commands in the specified order after the expansion of the LUN from the array side:

```
# format -d
# vxdisk resize
```

The Dynamic Reconfiguration tool is not supported inside the guest domain for Oracle VM Server for SPARC (3405223)

SFHA Solutions provides a Dynamic Reconfiguration tool to simplify online dynamic reconfiguration of a LUN. The Dynamic Reconfiguration tool is not supported if SFHA Solutions is running inside the guest domain for Oracle VM Server for SPARC.

Creating a disk group with a large number of objects or splitting, joining, or moving such a disk group reports an out of kernel memory error (3069711)

When you create a disk group with an extremely large number of objects (volumes, snapshots, plexes, disks), you may see the following error:

```
ERROR-V-5-1-10128 Out of kernel memory
```

You may also see the error when you perform operations like split, join, move on such a disk group.

Each object has a record which is used for its description and state. These records are stored in the private region of every disk group. The default private region size is 32 MB which can accommodate a sufficient number of objects. If the private region of disk group does not have space to create a new record, the operation fails with the above error message. Typical use cases would not hit this condition.

Workaround:

The best practice is not to have an extremely large number of objects in the disk group. Instead, split the disk group into multiple disk groups.

Refer to the section “Reorganizing the contents of disk groups” in the *Administrator's Guide* for information about splitting disk groups.

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.

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
   ```

Probing vxio with DTrace fails on Sparc machines. (2180635)

This issue exists because of inability of DTrace to load a module whose text size is greater than 2MB on Sparc machines. While trying to load `vxio` with DTrace you may see following warning messages on console:

```
dtrace: WARNING: couldn't allocate SDT table for module vxio
fbt: WARNING: couldn't allocate FBT table for module vxio
```

There is no workaround for this issue.

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 SF 6.1.1, 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.

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

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

To run disk discovery

- Run the following command:

```
# vxdisk scandisks
```

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.

Removing an array node from an IBM Storwize V7000 storage system also removes the controller (2816589)

When using an IBM Storwize V7000 storage system, after removing one array node, the corresponding controller is also removed.

Workaround: The following procedure resolves this issue.

To resolve this issue

1 Set the `iotimeout` tunable to 600:

```
# vxdmpadm setattr enclosure encl1 recoveryoption=throttle \
 iotimeout=600
```

2 After you re-add the SAN VC node, run the `vxdctl enable` command for Dynamic Multi-Pathing (DMP) to detect the added paths:

```
# vxdctl enable
```

vxmirror to SAN destination failing when 5 partition layout is present: for example, root, swap, home, var, usr (2815311)

The `vxmirror` command may fail with following error on a Solaris 10 host, for a thin LUN, if more than one partition excluding root and swap is present.

```
VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
```

Example

```
# /etc/vx/bin/vxmirror" -f -g rootdg_17_23_49 rootdisk01 rootdisk02
        ! vxassist -g rootdg_17_23_49 mirror swapvol rootdisk02
        ! vxassist -g rootdg_17_23_49 mirror rootvol rootdisk02
        ! vxassist -g rootdg_17_23_49 mirror usr rootdisk02
        ! vxassist -g rootdg_17_23_49 mirror var rootdisk02
        ! vxassist -g rootdg_17_23_49 mirror home rootdisk02
        ! vxbootsetup -g rootdg_17_23_49
        VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
        VxVM vxbootsetup ERROR V-5-2-5678 Skipping volume 'home_dcl'
         because no free partitions are available on disk 'disk_0'.
         Either remove the volume or make a partition available
        VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
        VxVM vxbootsetup ERROR V-5-2-5678 Skipping volume 'usr_dcl'
         because no free partitions are available on disk 'disk_0'.
         Either remove the volume or make a partition available
        VxVM vxbootsetup WARNING V-5-2-5667 Max volume count 5 exceeded.
        VxVM vxbootsetup ERROR V-5-2-5678 Skipping volume 'var_dcl' because
         no free partitions are available on disk 'disk_0'.
         Either remove the volume or make a partition available
        /usr/lib/vxvm/bin/vxmksdpart: 3pardata0_2492: is not an identifier
```

Disk group import of BCV LUNs using -o updateid and -ouseclonedev options is not supported if the disk group has mirrored volumes with DCO or has snapshots (2831658)

VxVM uses guid stored in configuration to uniquely identify all objects. The data change object (DCO) volume stores the guid of mirrors and snapshots. If the disk group is imported with `-o updateid` and `-o useclonedev`, it changes the guid of objects in VxVM configuration database and the guids stored in the DCO volume are not updated. The operations involving DCO cannot find objects with the stored guid. This could lead to failure of certain operations involving DCO or could lead to unexpected behavior.

Workaround:

No workaround available.

The DMP EMC CLARiiON ASL does not recognize mirror view not ready LUNs (3272940)

On hosts that have EMC CLARiiON mirror view not ready LUNs, if you enable or disable the switch port and then issue the `vxdisk scandisks` or `vxdctl enable` command, I/O error messages are written continuously in the syslog.

The dynamic multi-pathing (DMP) request for providing information to identify mirror view not ready LUNs through in-band SCSI command is pending with EMC engineering. Not ready LUNs are special kind of LUNs which reject all kinds of I/O requests.

Because DMP does not recognize not ready LUNs, Veritas Volume Manager (VxVM) tries to bring them online. As part of the online process, VxVM issues I/Os to read the disk private region. These I/Os fail and generate error messages in syslog.

Because of events that are generated as part of the online process, the `vxattachd` script triggers the `vxdisk scandisks` command again. This cycle causes continuous I/O error messages. This problem can also other commands to run slowly because the VxVM configuration daemon (`vxconfigd`) is busy servicing `vxdisk scandisks`.

Workaround: Stop the `vxattachd` script and:

1. Disable the `vxattachd` process.

   For more information on how to disable `vxattachd` and what features you lose if `vxattachd` is disabled, see the `vxattachd` man page

2. Set the following EMC CLARiiON values:

   - `recoveryoption=fixedretry`
   - `retrycount=5`

Enter:

```
vxdmpadm setattr enclosure enclosure_name recoveryoption=fixedretry \
retrycount=5
```

The vxdiskadm command cannot exclude or suppress devices (3339195)

The `vxdiskadm` command cannot exclude or suppress devices, because it does not accept the syntax used, which results in syntax errors.

Workaround:

Use the `vxdmpadm exclude` command to exclude or suppress devices.

When all Primary/Optimized paths between the server and the storage array are disconnected, ASM disk group dismounts and the Oracle database may go down (3289311)

The Oracle database shows an I/O error on the control file, but there was no I/O error seen on any DMP device. When all Primary/Optimized paths are disconnected, DMP fails over to other available paths but the failover takes time. In the meantime, the application (ASM/Oracle database) times out the I/O.

The ASM alert log file displays messages such as the following:

```
Errors in file /u01/app/oracle/diag/rdbms/orcl/orcl2/trace/orcl2_ckpt_6955.trc:
ORA-00221: error on write to control file
ORA-00206: error in writing (block 4, # blocks 1) of control file
ORA-00202: control file: '+DATA_P6/ORCL/CONTROLFILE/current.261.826783133'
ORA-15081: failed to submit an I/O operation to a disk
ORA-15081: failed to submit an I/O operation to a disk
Wed Oct 09 14:16:07 2013
WARNING: group 2 dismounted: failed to read virtual extent 0 of file 261
Wed Oct 09 14:16:07 2013
USER (ospid: 6955): terminating the instance due to error 221
Wed Oct 09 14:16:07 2013
WARNING: requested mirror side 2 of virtual extent 0 logical extent 1 offset
16384
is not allocated; I/O request failed
WARNING: requested mirror side 3 of virtual extent 0 logical extent 2 offset
16384
is not allocated; I/O request failed
```

The above issue may occur when the server is configured as follows:

DB: Oracle 12c

Volume Manager: ASM

Multipathing Solutions: DMP

OS: Solaris

Disk Array : HP EVA in ALUA mode

Workaround:

The following workaround can reduce the probability of this issue, and when you see this issue, you could use Oracle commands to start the database manually.

Increase the application time out and make the following changes to reduce the time taken to mark the path as offline:

- In the `/kernel/drv/fp.conf` file, add `fp_offline_ticker=15`.
- In the `/kernel/drv/fcp.conf` file, add `fcp_offline_delay=10`.

Running the vxdisk *disk* set clone=off command on imported clone disk group luns results in a mix of clone and non-clone disks (3338075)

If you do not specify a disk group name, the `vxdisk set` operation works on the `dmname` rather than the `daname`. If a `dmname` is the same as an existing `daname`, the vxdisk set operation reflects on the dm name.

Workaround: Use the following command syntax to set the attributes:

```
vxdisk -g diskgroup_name set dmname clone=off
```

For example:

```
vxdisk -g dg1 set eva4k6k0_12 clone=off
```

Disk greater than 1TB goes into error state (3269099)

If a path of a device having multiple paths is labelled with the EFI format using an operating system command such as `format`, the `vxdisk` list command output shows the device in error state.

Workaround:

This issue is a Solaris OS issue. There is no workaround for this issue.

Importing an exported zpool can fail when DMP native support is on (3133500)

On Solaris, when the tunable `dmp_native_support` is set to `on`, importing an exported zpool using the command `zpool import poolname` can fail with following error:

```
Assertion failed: rn->rn_nozpool == B_FALSE, file
../common/libzfs_import.c,
```

```
line 1084, function zpool_open_func
Abort (core dumped)
```

Workaround:

Import the zpool using the following command, specifying the DMP device directory:

```
# zpool import -d /dev/vx/dmp poolname
```

Server panic after losing connectivity to the voting disk (2787766)

This issue occurs on A/P arrays. If the voting disk loses connectivty to the primary paths, DMP takes some time to analyze the error and fail over the paths. During this time, the `cssd` reports a timeout and panics. When using Oracle ASM over DMP devices, set the `disktimeout` parameter to an appropriate value. This parameter indicates the maximum time allowed for a voting file I/O to complete. If this time is exceeded, the voting disk is marked as offline.

The default of `disktimeout` is 200. If the value of the tunable is less that this value, reset the value to the default value.

Workaround:

To set the `disktimeout` to 200:

```
$CRS_HOME/bin/crsctl set css disktimeout 200 [-force] test
```

Cascaded failure of nodes with ioship enabled may cause the vxconfigd daemon to hang (2865771)

In a shared disk group environment with ioship enabled, the`vxconfigd` daemon may hang in certain cases. When the I/O is initiated from the slave node that has lost connectivity to the disks locally, the I/O is shipped to other nodes. If the node processing the shipped I/O also leaves the cluster shortly after the first node, and tries to rejoin the cluster as a slave, the cascaded failures may cause the `vxconfigd` daemon to hang.

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 muliple enclosures. The rebalancing process takes time, during which time the `vxconfigd` daemon is busy and does not respond to commands.

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

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

Suppressing the primary path of an encapsulated SAN boot disk from Veritas Volume Manager causes the system reboot to fail (1933631)

If you suppress the primary path of an array from VxVM control and then reboot the system, the system boot fails.

If you have an encapsulated SAN boot device with multiple primary paths, the issue occurs when you suppress the first primary path. When you configure a SAN boot device, the primary path is set as a boot device. In general, the first path of the SAN boot device corresponds to the first configured path during SAN boot. Even if another primary path is configured as a boot device, suppressing the first device from VxVM causes the boot to fail.

Workaround:

When the boot device is suppressed from VxVM, change the OS boot device sequencing accordingly.

For Solaris SPARC system, use the `eeprom` *`boot-device`* command to set the boot device sequencing.

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

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

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

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

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

Work-arounds:

Use one of the following work-arounds:

- Remove the faulty slave node out of CVM cluster, restore storage connectivity, and rejoin the node to the cluster.
- Restart `vxconfigd` on the CVM master node.

After changing the preferred path from the array side, the secondary path becomes active (2490012)

For EVA arrays, DMP requires that the prefer bit is static. If the prefer bit is not static, issues like the following may occur. After changing the prefer path of LUN from the array side, and performing a disk discovery (`vxdisk scandisks`) from the host, the secondary path becomes active for the LUN.

Workaround:

To work around this issue

1. Set the pref bit for the LUN.
2. Perform disk discovery again:

```
# vxdisk scandisks
```

Upgrading from Symantec Storage Foundation and High Availability Solutions 5.x to 6.1.1 may fail for IBM XIV Series arrays (2715119)

Starting in the Symantec 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 Symantec Storage Foundation and High Availability Solutions from a release prior to that release to the current 6.1.1 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`.

Cannot grow Veritas Volume Manager (VxVM) disk using the vxdisk resize command during Dynamic LUN Expansion operation (2064510)

The following error message is displayed during the Dynamic LUN Expansion operation of a LUN with the SIMPLE format:

```
VxVM vxdisk ERROR V-5-1-8643 Device <device name>: resize failed:
Invalid data in request
```

The `vxdisk resize` command keeps the cylinder size (number of the heads * total number of the sectors per track) constant before and after the resize operation, unless the number of cylinders go beyond 2^16-1 (65535) . Because of the VTOC limitation of storing geometry values only till 2^16 -1, if the number of cylinders increases beyond the limit, `vxdisk resize` increases the cylinder size. If this happens, the private region will overlap with the public region data and corrupt the user data.

As a result of this LUN geometry change, VxVM is unable to complete `vxdisk resize` on simple format disks. VxVM was not designed to handle such geometry changes during Dynamic LUN Expansion operations on simple disks.

Workaround:

The VxVM `vxdisk resize` command behaves differently depending on whether the disk is simple, sliced, or CDS format.

The problem shown above only occurs on simple disk configurations. As a result of this difference in behavior, if the geometry changes during a Dynamic LUN Expansion operation at the LUN level, you can convert the disk to a CDS format disk. Use the `vxcdsconvert` command on the disk. Then you can issue the `vxdisk resize` command.

See http://www.symantec.com/docs/TECH136240 for more information.

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
```

A master node is not capable of doing recovery if it cannot access the disks belonging to any of the plexes of a volume (2764153)

A master node with missing disks is not capable of doing recovery, as it does not have access to the disks belonging to any of the plexes of a volume.

Workaround:

If other nodes have access to the storage, they can do the recovery. Switch the master role to some other node with better storage connectivity.

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

In an enviroment 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 to each LUN path to check the health of path, which are delayed by the presence of EMC PowerPath.

CVMVolDg agent may fail to deport CVM disk group

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

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

vxresize does not work with layered volumes that have multiple plexes at the top level (3301991)

If a layered volume has multiple plexes at the top level, `vxresize` does not work. For example, if you add a mirror to a concat-mirror volume for a third mirror snapshot. The `vxresize` operation fails with the following message:

```
VxVM vxassist ERROR V-5-1-2528 Volume volname built on layered volumes
have multiple plexes
VxVM vxresize ERROR V-5-1-4703 Problem running vxassist command for
volume volname, in diskgroup dgname
```

Workaround:

To resize the volume

1. After adding the mirror to the volume, take a snapshot using the plex.
2. Grow the volume and snapshot volume with `vxresize`.
3. Reattach the snapshot volume to the source volume.

In a clustered configuration with Oracle ASM and DMP and AP/F array, when all the storage is removed from one node in the cluster, the Oracle DB is unmounted from other nodes of the cluster (3237696)

In a clustered configuration with Oracle ASM and DMP and AP/F array, when you remove all the storage from one node in the cluster, I/O is expected to fail on this node. Due to an issue with the Oracle ASM configuration, the Oracle database is unmounted from other nodes of the cluster. This issue is not seen if you delay the I/O failure from DMP. The Oracle database works fine on other node.

Workaround:

Increase the `dmp_lun_retry_timeout` tunable value to 300 with following command.

```
# vxdmpadm settune dmp_lun_retry_timeout=300
```

Disk group deport operation reports messages in the syslog for remote disks (3283518)

During the **vxdg deport** operation, the following messages may be seen in the syslog for remote disks:

Aug 12 14:51:57 swlx87 vxvm:vxconfigd: V-5-1-12708 vold_pgr_unregister(): failed to get key (Error 9).

Workaround:

These messages can be ignored for all the remote disks.

Changes in enclosure attributes are not persistent after an upgrade from release prior to VxVM 5.1SP1 (2082414)

The Veritas Volume Manager (VxVM) 6.1.1 includes several array names that differ from the array names in releases 5.1SP1 or prior. Therefore, if you upgrade to VxVM 6.1.1 from a release 5.1SP1 or earlier, 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 6.1.1. Manually reconfigure the enclosure attributes to resolve the issue.

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

Table 1-28 Hitachi arrays with new array names

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

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 6.1.1. 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

MPxIO device names shown in error state (3169587)

In this release, DMP does not support extended attributes like AVID for Solaris MPxIO devices. Up until the 5.1SP1 release, DMP used to support AVID for the MPxIO devices. When you upgrade from 5.1SP1 or prior release to 6.0 or later release, DMP assigns new names to the MPxIO devices.

The MPxIO device may go into an error state after the upgrade, if a persistent disk access record (entry in `/etc/vx/darecs`) exists with the old name, and the device was assigned a new name.

The same issue may occur if the MPxIO device name changes for another reason, such as the changed cabinet serial numbers for 3PAR or XIV devices from 6.0.

Workaround:

Use the following procedure to remove the persistent disk access record and resolve the issue.

To resolve the issue with MPxIO devices in error state:

1 Remove the following file:

```
# rm /etc/vx/darecs
```

2 Reset the `vxconfigd` daemon:

```
# vxconfigd -kr reset
```

The administrator must explicitly enable and disable support for a clone device created from an existing root pool (3110589)

A non-rpool is a clone of the existing root pool. When native support is enabled, DMP does not touch the clone root pool because the clone may or may not have the VxVM package.

Workaround: To add or remove DMP support for a clone boot device, the administrator must boot through the clone and turn on/off `dmp_native_support`.

For Solaris 11.1 or later, the system can panic when system is rebooted after turning dmp_native_support to on (3341674)

For Solaris 11.1 or later when more than 512 LUNs are configured, the system can panic when the system is rebooted after setting the tunable parameter `dmp_native_support` to `on`.

Workaround:

For Solaris 11.1 or later, DMP native support for ZFS is restricted to set-ups with no more than 512 LUNs.

Disks on the LDOM guest are claimed under other_disks category (2354005)

The disks on the LDOM guest are claimed under "other_disks" enclosure, because these disks are not capable of being multi-pathed by DMP. This is expected because these devices represent VxVM volumes in the host. By design, devices under other_disks enclosure have their name based on underlying OS path regardless of the DDL naming scheme.

DMP uses OS device physical path to maintain persistence of path attributes from 6.0 (2410716)

From release 6.0, DMP uses OS device physical path instead of logical name to maintain persistence of path attributes. Hence after upgrading to DMP 6.0 or later releases, path attributes are reset to the default values. You must reconfigure the path-level attributes that were defined in the /etc/vx/dmppolicy.info file.

Workaround:

To configure path-level attributes:

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

2 Reset the path attributes.

1 TB luns goes in error state with Solaris x86 (2706776)

If you label a disk device as EFI using format on a subset of the paths or on the DMP device, Solaris will not be able to propagate the label to all the other paths of the LUN. This will lead the device to appear in the error state under 'vxdisk list'.

Workaround: There is no workaround for this issue.

For Solaris 11.1 or later, system hangs when both QLogic and Emulex HBAs are present and dmp_native_support is turned on (3138703)

For Solaris 11.1 or later, the system may hang when both QLogic and Emulex HBAs are present, and **dmp_native_support** is turned on.

Workaround:

The system hang is not seen if all of the HBAs are either from Emulex or from QLogic. Do not combine both HBAs on the same system.

For Solaris 11.1 or later, enabling DMP native support requires steps to enable booting from alternate root pools (3133514)

For Solaris 11.1 or later, if the tunable parameter dmp_native_support is set to on, using the following command causes alternate root pools on OS devices to migrate to DMP devices:

```
# zpool import -d /dev/vx/dmp
```

After the above command is run, the system cannot boot using these alternate root pools because the DMP driver is not configured for these root pools. This scenario is shown by the following output.

```
# zpool status

pool: crpool
 state: ONLINE
  scan: none requested
config:

      NAME                    STATE   READ WRITE CKSUM
      crpool                  ONLINE  0     0     0
        /dev/vx/dmp/disk_0s0  ONLINE  0     0     0
```

Workaround:

To boot using the alternate root pools, export and re-import the root pools using the OS device.

To boot using the alternate root pools

1 Export the root pool:

```
# zpool export crpool
```

2 Display the OS path name for the device:

```
# vxdmpadm getsubpaths dmpnodename=disk_0
NAME      STATE[A]   PATH-TYPE[M] CTLR-NAME ENCLR-TYPE ENCLR-NAME ATTRS
==========================================================================
c3t2d0s2 ENABLED(A) -            c3        Disk       disk       -
```

3 Re-import the root pools using the OS device.

```
# zpool import crpool -d /dev/dsk/c3t2d0s0
```

The system is now bootable using the alternate root pools.

For Solaris 11.1 or later, uninstalling DMP or disabling DMP native support requires steps to enable booting from alternate root pools (3178642)

For Solaris 11.1 or later, after you uninstall the VxVM package or after you turn off DMP native support, you may see this issue. After reboot, the root pool containing the active boot environment is migrated to the OS device but alternate root pools continue to show DMP device. The status of the alternate root pools and their DMP devices is shown as "UNAVAIL".

```
pool: crpool
 state: UNAVAIL
status: One or more devices are unavailable in response to persistent
        errors. There are insufficient replicas for the pool to continue
        functioning.
action: Destroy and re-create the pool from a backup source. Manually
        marking the device repaired using 'zpool clear' or 'fmadm repaired'
        may allow some data to be recovered.
        Run 'zpool status -v' to see device specific details.
  scan: none requested
config:

        NAME                  STATE    READ WRITE CKSUM
```

```
crpool                  UNAVAIL      0      0      0
  emc_clariion1_82s0    UNAVAIL      0      0      0
```

The tunable parameter `dmp_native_support` only unconfigures DMP for the single root pool containing the active boot environment. If the setup has any alternate root pools, for which DMP native support was enabled, then the alternate root pools continue to show the DMP device. If the alternate root pool is configured in the current boot environment and DMP support is removed, the DMP devices required for ZFS are not found. The DMP devices and the root pools display the state as "UNAVAIL".

Workaround:

Even though the status of alternate root pool is "UNAVAIL", the system is bootable using the disk containing the alternate root pool. Reboot the system with the disk containing the alternate root pool. The system comes up with the root pool using the DMP device.

For Solaris 11.1 or later, after enabling DMP native support for ZFS, only the current boot environment is bootable (3157394)

After enabling DMP native support for ZFS on Solaris 11.1 or later, only the current boot environment (BE) is bootable. Any alternate BEs in the same root pool are not bootable. This situation occurs because the DMP native support configures the ZFS root pool so that only DMP can import the root pool. If you attempt to boot the system from the alternate BE, the system panics with the following message:

```
NOTICE: zfs_parse_bootfs: error 19
Cannot mount root on rpool/193 fstype zfs

panic[cpu0]/thread=10012000: vfs_mountroot: cannot mount root

Warning - stack not written to the dumpbuf
000000001000fa00 genunix:main+17c (1, 100dc958, 12d5c00, 124702c, 0, 10828000)
  %l0-3: 0000000010010000 0000000000000000 00000000100dc800 0000000000000000
  %l4-7: 0000000010012000 0000000000000000 000000001038f7c0 000000000104c800
```

Workaround:

To enable booting from another BE, configure the ZFS root pool so that it can be imported without DMP.

To configure ZFS root pool to enable booting from all the BEs

1 At the OBP PROM, run the following command to list all the BEs:

```
ok> boot -L
```

2 Use the following command to boot from the BE for which DMP native support for ZFS is enabled.

```
ok> boot -Z rpool/ROOT/BE_name
```

3 After booting through new BE, disable the DMP native support using the following command:

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

The system is now bootable from any BEs in the ZFS root pool.

When dmp_native_support is set to on, commands hang for a long time on SAN failures (3084656)

When dmp_native_support is set to on, on SAN failures, commands that do I/O operations to the root file system or I/O to disks that contain the root pool may hang for about 1-5 minutes. The commands include commands like "zpool status", or telnet initiated to connect the system. The hang is seen because the drivers below the DMP layer take more time to report the I/O failure when some of the paths to the disk containing the root pool are disconnected. This situation should not lead to any root pool data corruption.

Workaround:

This hang cannot be avoided but the hang time can be reduced by tuning the following parameters

To tune the parameters

1 In the `/kernel/drv/fp.conf` file, set

```
fp_offline_ticker=15
```

2 In the `/kernel/drv/fcp.conf` file, set

```
fcp_offline_dely=10
```

3 Reboot the sytem to apply the changes.

These steps reduce the hang time to a maximum of 1 minute.

Veritas File System known issues

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

Write operation fails with error [3558675]

The write operation may fail in the presence of checkpoints that have quota limits assigned for checkpoint usage. For example, if a checkpoint is in the process of being removed when a write operation is in progress, the space needed by the write operation may not be available as the removal is not complete. As a result, the write operation fails with the following error:

```
disk quota exceeded
```

Workaround: Retry the write operation.

Command `fsppadm` (1m) for assigning placement policy gives unrecognizable warnings [3538882]

When you run the `fsppadm` (1m) command for assigning the placement policy, the placement policy doesn't take effect and gives the following warnings:

```
# fsppadm assign placement_policy_xml UX:vxfs fsppadm:
WARNING: V-3-26733: One or more of the PERIODs using in the IOTEMP
and/or ACCESSTEMP criteria are too small to collect reasonable amount
of IO statistics for mount_point. The default fcl_winterval is one
hour. So please reduce fcl_winterval with vxtunefs and persist the
new setting with the help of /etc/vx/tunefstab.
And then assign the policy again UX:vxfs fsppadm:
WARNING: V-3-26712: The fcl_keeptime for mount_point is 0 hours,
policy requires 1 hours of filestats. Some IOTEMP and ACCESSTEMP based
relocation and deletion may not occur
```

Workaround:

If `UX:vxfs fsppadm: WARNING: V-3-26733` is displayed, set the value of the `fcl_keeptime` tunable larger or equal to the value (in seconds) of the hours field in the policy.

If `UX:vxfs fsppadm: WARNING: V-3-26712` is displayed, set the value of the hours field in the placement policy larger than 3 hours.

Full fsck flag may be set on the file system during the write operation [3451284]

If the data of summary and bitmap for the file system allocation unit gets mismatched, when VxFS allocates extents during the write operation, the full fsck flag might be set on the file system.

Workaround: Run `fsck` on the file system to correct the data of summary and bitmap for the file system allocation unit.

Some tunables go to default value when zero value is tried [3449606]

The VxFS kernel performs sanity check for the tunables when their values are changed. Some of the tunables have minimum and maximum. If the value is less than minimum or greater than maximum, the default is set without notification.

Workaround: Set the tunable to desired value using the `vxtunefs -o` *`tunable=value mntpt`* command.

The vxtunefs(1M) command accepts garbage values [3449150]

When garbage values are given for some tunables to the `vxtunefs(1M)` command, the command accepts the values and gives a successful update message. However, the values don't get reflected in the system. Consequently the garbage values for the tunable are accepted.

The issue can occur to the following tunables:

```
max_direct_iosz
default_indir_size
qio_cache_enable
odm_cache_enable
max_diskq
initial_extent_size
max_seqio_extent_size
hsm_write_prealloc
read_ahead
inode_aging_size
oltp_load
inode_aging_count
discovered_direct_iosz
write_throttle
```

Workaround: No workaround.

The `vxtunefs` commandreports error when it attempts to set the tunable values from the tunefstab file [3417076]

The `vxtunefs` command sets the tunable values that are mentioned in the `tunefstab` file (`/etc/vx/tunefstab` by default). When the file unusually contains blank lines or white spaces, the device name buffer contains the name of the previous (stale) device, and the tunable values are null. Consequently the command fails to set tunables for that device.

Workaround: Remove any unnecessary white spaces (such as blank lines, extra spaces or tabs) from the tunefstab file. Then set the tunables using the `vxtunefs -s` *`mntpt`* command.

spfile created on VxFS and ODM may contain uninitialized blocks at the end (3320906)

spfile created on VxFS and ODM may contain uninitialized blocks at the end due to space allocation with file system block size alignment. This is harmless and does not cause any problem to Oracle startup.

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

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

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

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

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.

Deduplication can fail with error 110 (2591473)

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

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

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

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

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

Workaround: Make more space available on the file system.

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

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

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

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

Warning message sometimes appear in the console during system startup (2354829)

During system startup, following messages sometimes appear in system console:

```
WARNING: couldn't allocate SDT table for module vxfs
WARNING: couldn't allocate FBT table for module vxfs
Loading smf(5) service descriptions: 2/2
```

These warnings indicate that the SDT and FBT DTrace probes might not be available for the VxFS module. The VxFS module still loads and works correctly. Dtrace SDT/FBT has limits on the size of module that it can support. Since the VxFS module exceeds the size that Dtrace can support, SDT and FBT Dtrace probes might not work for VxFS.

Workaround: There is no workaround for this issue.

In a VxFS file system that has compression enabled, the file system may get disabled due to the ENOSPC error (3301716)

In a VxFS file system that has compression enabled, the file system may get disabled due to the ENOSPC error. This occurs because of a defect in the delayed allocation feature.

Workaround: Turn off the delayed allocation.

Two VxFS parameters are deprecated (3260671)

The `read_unit_io` and `write_unit_io` parameters have been deprecated. If you specify values for these parameters, they do not affect system behavior. The parameters will be removed in a future release.

Workaround: Not applicable.

When in-place and relocate compression rules are in the same policy file, file relocation is unpredictable (3278193)

You cannot have in-place compress/uncompress rules and relocate compress/uncompress rules in the same policy file. If they are in the same file, file relocation is unpredictable.

Workaround: Create a different policy file for each policy, and enforce the policy as per the required sequence.

The file system may hang when it has compression enabled (3331276)

In a VxFS file system that has compression enabled, a deadlock in page fault handler can lead to the file system hang.

Workaround:

There is no workaround for this issue.

The file system may hang due to file system full conditions when file level snapshots are present (2746259)

In the presence of file level snapshots, file system full conditions may lead to the file system hang. Following a reboot, a mount may hang as well.

Workaround:

There is no workaround for this issue.

The file system may be marked for full fsck during a clone removal (2977828)

Under low memory conditions, a clone removal may lead to file system being marked for full fsck.

Workaround:

A full fsck of the file system will be required to recover the file system.

NFSv4 server panics in unlock path (3228646)

In a CFS configuration, if `fcntl(1m)` fails, some NFS specific structures `(I_pid)` are not updated correctly and may point to stale information. This causes the NFSv4 server to panic.

Workaround:

There is no workaround for this issue.

I/O errors on the file system may lead to data inconsistency (3331282)

If there are writable clones on the file system, I/O errors may lead to data inconsistency.

Workaround:

Run a full `fsck` to recover the file system.

Forcing the system to unmount during heavy I/O load may result in system panic in vx_is_fs_disabled_impl (3331284)

Forcing the system to unmount during heavy I/O load may result in system panic in `vx_is_fs_disabled_impl` .

Workaround:

There is no workaround for this issue.

The file system deduplication operation fails with the error message "DEDUP_ERROR Error renaming X checkpoint to Y checkpoint on filesystem Z error 16" (3348534)

The file system deduplication operation fails with the error message "DEDUP_ERROR Error renaming X checkpoint to Y checkpoint on filesystem Z error 16", due to the failure in unmounting the checkpoint.

Workaround: Retry the deduplication operation to resolve the problem.

The file system operations that need a file system freeze may take long to execute in the presence of file-level snapshots (3317368)

The file system operations that need a file system freeze may take long to execute in the presence of file-level snapshots, when there is heavy I/O load.

Workaround: There is no workaround for this issue.

On a system that has Solaris 11 Update 1, certain driver modules such as "fdd" may not be removed properly (3348829)

On a system that has Solaris 11 Update 1, certain driver modules such as “fdd" may not be removed properly during the uninstallation of the SF or SFCFS stack.

Workaround: Prior to uninstallation of the stack, this can be mitigated by following the workaround indicated below:

```
# rm /usr/kernel/drv/sparcv9/fdd
```

Replication known issues

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

vradmin syncvol command compatibility with IPv6 addresses (2075307)

The `vradmin 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 `vradmin 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 (2036605)

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" (2036644)

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

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

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

Workaround:

To resolve this issue:

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

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

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

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

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

Issue 1:

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

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

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

Issue 2:

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

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

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

Workaround: The following workarounds resolve these issues.

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

For example:

```
# fsck -F 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 -F 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 release 6.0 or later, `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.

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

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

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

Workaround:

To resolve this issue:

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

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

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

vxassist relayout removes the DCM (145413)

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

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

vradmin functionality may not work after a master switch operation (2158679, 2158679)

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

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

Workaround:

To restore vradmin functionality after a master switch operation

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

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

2 Re-enter the command that failed.

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:

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

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

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

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

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

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

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

7. Start the RVG. Enter the following:

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

8. Resume or start the applications.

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 *Symantec 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.

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.

vradmin verifydata operation fails if the RVG contains a volume set (2808902)

In a VVR environment, the `vradmin verifydata` command fails with the following error if the replicated volume group (RVG) contains any volume set:

```
Message from Primary:
VxVM VVR vxrsync ERROR V-5-52-2009 Could not open device
/dev/vx/dsk/vvrdg/<volname> due to: stat of raw character volume path
failed
```

RLINK name cannot exceed 31 characters

The `vradmin` utility truncates the RLINK name to 31 characters, as the `vxmake` utility does not support the creation of RLINK names that are longer than 31 characters.

Workarounds:

- Specify the `prlink` and `srlink` attributes using the `vradmin addsec` command, so you can choose the RLINK name in the `addsec` command line.
- If using IPv6 addresses, create host name aliases for the IPv6 addresses and specify the aliases in the `addsec` command line.

Bunker replay does not occur with volume sets (3329970)

There are issues with bunker replication using Volume Replicator (VVR) with volume sets. Do not upgrade to Symantec Storage Foundation HA 6.1.1 if you have configured or plan to configure bunker replication using VVR with volume sets.

Workaround: Contact Symantec Technical Support for a patch that enables you to use this configuration.

During moderate to heavy I/O, the vradmin verifydata command may falsely report differences in data (3270067)

While an application is online at the Volume Replicator primary site, the `vradmin verifydata` command may fail. The command output shows the differences between the source data volume and the target data volume.

Workaround: The reason for this error is that the cache object that is used for the verification might be under allocated. You might need to allocate more space for the shared cache object. For guidelines on shared cache object allocation, see the section "Creating a shared cache object" in the *Symantec Storage Foundation Administrator's Guide*.

The vradmin repstatus command does not show that the SmartSync feature is running (3345984)

In a Volume Replicator (VVR) environment, after you start the initial synchronization with the `vradmin -a startrep` command with file system mounted on the primary data volumes, the `vradmin repstatus` command does not show that the SmartSync feature is running. This is an only issue with the output of the `vradmin repstatus` command.

Workaround: To confirm that SmartSync is running, enter:

```
vxrlink status rlink
```

While vradmin commands are running, vradmind may temporarily lose heartbeats (3347656)

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

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

Workaround: To resolve this issue:

1 Depending on the application I/O workload and the 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 replicated data set (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
```

Write I/Os on the primary logowner may take a long time to complete (2622536)

Under a heavy I/O load, write I/Os on the Volume Replicator (VVR) primary logowner take a long time to complete.

Workaround: None

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

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

Workaround:

To resize layered volumes that are associated to an RVG:

1 Pause or stop the applications.

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

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

3 Stop the affected RVG. Enter the following:

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

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

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

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

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

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

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

7 Start the RVG. Enter the following:

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

8 Resume or start the applications.

Symantec Storage Foundation for Databases (SFDB) tools known issues

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

ReverseResyncBegin operation may fail to create checkpoint [3079910]

Once the specified operations are performed in the following order, they may cause error:

1 `ReverseResyncBegin`

2 `ReverseResyncAbort`

3 `ReverseResyncBegin`

The last `ReverseResyncBegin` operation may fail with the following error messages:

```
SFDB vxsfadm ERROR V-81-0425 Checkpoint Creation(all) failed

Reason: SFDB vxsfadm ERROR V-81-0560 UX:vxfs fsckptadm: ERROR: V-3-24643:
storage checkpoint createall failed on snapha31360308368, A file or\
directory in the path name does not exist. (2)
```

Workaround: There is no workaround.

Instant mode clone fails in RAC environment for all FSMs with data loading (3517782)

When you use the instant clone mode for RAC databases, the clone operation may fail during Oracle recovery. The issue is more likely to be seen when there is load activity on some of the RAC nodes.

Workaround: Use either online or offline snapshot mode.

SFDB commands do not work in IPV6 environment (2619958)

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

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 Checkpoint, 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.

Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.1.1 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.1.1.

When upgrading from SFHA Solutions version 5.0 to SFHA Solutions 6.1.1 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

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

Workaround

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

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

If you have a parameter, such as `log_archive_dest_1`, in single line in the `init.ora` file, then `dbed_vmclonedb` works but `dbed_vmcloneb` fails if you put in multiple lines for parameter.

Workaround: Edit the PFILE to arrange the text so that the parameter values are on a single line. If the database uses a spfile and some parameter values are spread across multiple lines, then use the Oracle commands to edit the parameter values such as they fit in a single line.

Workaround

There is no workaround for this issue.

Data population fails after datafile corruption, rollback, and restore of offline 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 workround at this point of time.

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

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 6.1.1 release, create distinct archive and datafile mounts for the 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.

'vxdbd' process is online after Flash archive installation (2869269)

After a Flash archive installation of the SF stack, the `vxdbd` process is up, even if the stack is not configured.

Workaround: You can ignore, or stop the `vxdbd` process using the `/opt/VRTSdbed/common/bin/vxdbdctrl stop`command.

`sfua_rept_migrate` fails after phased SF Oracle RAC 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.

The ReverseResyncBegin (RRBegin) operation fails when performed on multiple snapshot configurations (3066532)

When you perform a Reverse Resync operation on multiple snapshot configurations, SFDB reports the following error message:

```
[oracle@dblxx64-3-vip3 ~]$ vxsfadm -a oracle -s flashsnap --name \
man -o rrbegin

SFDB vxsfadm ERROR V-81-0943 Repository already relocated to alternate
location.
```

As per the Reverse Resync design, the first RRBegin operation relocates the SFDB repository to a backup location, and the ReverseResyncAbort and ReverseResyncCommit operations restore it to the original location. When the second RRBegin operation attempts to relocate the same repository which is already relocated, SFDB reports the error message.

Workaround: Make sure to perform the RRAbort or RRCommit operation using the snapshot configuration that is in the RRBegin state.

Note: You must complete Reverse Resync operations for a particular configuration before you start with another configuration.

The ReverseResyncBegin (RRBegin) operation with recovery option as AUTO fails (3076583)

The RRBegin operation with the recovery option as AUTO fails when you perform the following sequence of operations:

1. Validate the FlashSnap setup using the validate operation.
2. In the database, take the tablespace offline.
3. Perform a snapshot operation.
4. Bring the tablespace online which was taken offline in 2.
5. Perform the Reverse Resync Begin operation.

Note: This issue is encountered only with Oracle version 10gR2.

Workaround: Perform one of the following:

- Make sure to bring the tablespace online only after performing the RRBegin and RRCommit operations. Otherwise, perform the Reverse Resync Begin operation while the tablespace is in the offline mode.
- To recover a database, specify the recovery option as **AUTO_UNTIL_SCN** in the RRBegin operation.

The ReverseResyncBegin (RRBegin) operation fails and reports an error message due to a missing binary control file (3157314)

When the RRBegin operation cannot find the binary control file that is used to recover a database instance, it reports the following error message:

```
[oracle@testbox ~]$ vxsfadm -a oracle -s flashsnap -name man -o rrbegin

SFDB vxsfadm ERROR V-81-0949 Binary Control file is not available for
recovery purposes
```

This issue is observed in the third-mirror break-off type (FlashSnap) snapshots that are created using the older SFDB version, which did not include the binary control file in the snapshot images.

Workaround:

There is no workaround for this issue.

The dbdst_obj_move(1M) command moves all the extents of a database table (3277003)

The `dbdst_obj_move(1M)` command moves all the extents of a database table when:

- The `dbdst_obj_move(1M)` command is run from the CFS secondary node.
- The object is an Oracle database table (-t option)
- A range of extents is specified for movement to a target tier (-s and -e options). The `dbdst_obj_move(1M)` command moves all extents of the specified table to a target tier when the extent size is greater than or equal to 32768. However, the expectation is to move only a specified range of extents.

Workaround: Run the `dbdst_obj_move(1M)` command from the CFS primary node.

Use the `fsclustadm showprimary` *<mountpoint>* and `fsclustadm idtoname` *<nodeid>*commands to determine the mode of a CFS node.

When you attempt to move all the extents of a table, the `dbdst_obj_move`(1M) command fails with an error (3260289)

When you attempt to move all the extents of a database table, which is spread across multiple mount-points in a single operation, the `dbdst_obj_move(1M)` command fails. The following error is reported:

```
bash-2.05b$ dbdst_obj_move -S sdb -H $ORACLE_HOME -t test3 -c MEDIUM
FSPPADM err : UX:vxfs fsppadm: WARNING: V-3-26543: File handling failure
on /snap_datadb/test03.dbf with message -
SFORA dst_obj_adm ERROR V-81-6414 Internal Error at fsppadm_err
```

Note: To determine if the table is spread across multiple mount-points, run the dbdst_obj_view(1M) command

---

Workaround: In the `dbdst_obj_move(1M)` command, specify the range of extents that belong to a common mount-point. Additionally, if your table is spread across "n" mount-points, then you need to run the `dbdst_obj_move(1M)` command "n" times with a different range of extents.

Sometimes SFDB may report the following error message: SFDB remote or privileged command error (2869262)

While using SFDB tools, if you attempt to run commands, such as `dbed_update` then you may observe the following error:

```
$ /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 swpa04

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.
```

Workaround: There is no workaround for this issue.

The information file that is generated after a DBED data collector operation reports an error (2795490)

When the VRTSexplorer DBED scripts use the old VRTSdbms3-specific scripts that are removed from the products, the information file reports the following error:

```
/opt/VRTSdbms3/vxdbms_env.sh: cannot open [No such file or directory]
```

Workaround:

1. Run the `cd /opt/VRTSspt/DataCollector/sort` command. If this directory does not exist, run `sh /opt/VRTSspt/DataCollector/*.sh`.
2. Run the `cd advanced/lib/VOS/v10/Collector/VxExpCollector/explorer_scripts` command.
3. In `dbed_rept_sql`, comment

   `$`**`VXDBMS_DIR/vxdbms_env.sh`**

   Or

   Replace **$VXDBMS_DIR/vxdbms_env.sh** with

   ```
   [[ -f $VXDBMS_DIR/vxdbms_env.sh ]] &&
       {
            . $VXDBMS_DIR/vxdbms_env.sh
       }
   ```

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

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

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

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

Workaround: To use FlashSnap, modify the above configuration to *.log_archive_dest_1='location=/tpcc_arch'. For example,

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

Virtualization known issues

There are no new virtualization known issues in this release of Symantec Storage Foundation and High Availability Solutions (SFHA Solutions).

Symantec Storage Foundation Cluster File System High Availability known issues

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

See "Veritas File System known issues" on page 121.

See "Veritas Volume Manager known issues" on page 101.

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

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

Workaround

To resolve this issue

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

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

The svsiscsiadm create lun command fails if you create a LUN greater than the available space on the file system (2567517)

The `svsiscsiadm create lun` command fails if you create a LUN of a size greater than the total amount of space available on the file system. The underlying `iscsitadm` command fails with the following error message:

```
iscsitadm: Error Requested size is too large for system
```

The report of this error is logged in the `/var/VRTSvcs/log/engine_A.log` file.

If you then try to create a LUN on the same target, the LUN creation call fails again with the following error message:

```
iscsitadm: Error Failed to create a symbolic link to the backing store
```

The report of this error is logged in the `/var/VRTSvcs/log/engine_A.log` file.

This makes the target unusable.

Workaround

To resolve this issue

1. Note the TargetID and LunID on which the `svsiscsiadm create lun` command failed. To find the failed LunID, note the last LunID for the target on which `svsiscsiadm create lun` command failed with the use of the `svsiscsiadm list` command. To calculate the failed LunID, add 1 to last LunID seen by the `svsiscsiadm list` command.
2. Go to the configuration directory for the TargetID:

   ```
   # cd /etc/iscsi/TargetID .
   ```

3. Delete the symlink pointing to path of LUN backing file which failed to get added. The below LunID is the failed LunID, which is a result of calculation in point 1:

   ```
   # rm -f /etc/iscsi/TargetID/lun.($lunid + 1)
   ```

After removal of the symlink you should be able to add LUNs on the unusable target.

Miscalculated file set usage (2123429)

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

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

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

Workaround:

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

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

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

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

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

The fsappadm subfilemove command moves all extents of a file (3258678)

This issue occurs under following conditions:

- You run the `fsppadm subfilemove` command from a cluster file system (CFS) secondary node.
- You specify a range of extents for relocation to a target tier.

If the extent size is greater than or equal to 32768, the `fsppadm subfilemove` command moves all extents of the specified table to the target tier. The expectation is to move a specified range of extents.

Workaround:

- On the CFS primary node, determine the primary node using one of the following commands:

  ```
  # fsclustadm showprimary mountpoint
  ```

  ```
  # fsclustadm idtoname nodeid
  ```

The cluster may hang due to a known lock hierarchy violation defect (2919310)

If VxFS File Change Log (FCL) is turned ON in Cluster File System (CFS) environments, a known lock hierarchy violation defect may lead to the cluster hang.

Workaround:

There is no workaround for this issue.

Certain I/O errors during clone deletion may lead to system panic. (3331273)

Certain I/O errors during clone deletion may lead to system panic.

Workaround:

There is no workaround for this issue.

Symantec Storage Foundation for Oracle RAC known issues

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

Oracle RAC issues

This section lists the known issues in Oracle RAC.

Oracle Universal Installer (OUI) abruptly closes without responding [3566625]

During the installation of Oracle Grid Infrastructure 11.2.0.4 on SF Oracle RAC nodes running Solaris 10/11, the Oracle Universal Installer (OUI) abruptly closes without responding if the “Shared File System” option is selected on the “Oracle Cluster Registry” page of the installer.

Workaround:

- For “Typical Installation” of Oracle Grid Infrastructure 11.2.0.4, enter:

```
# ./runInstaller -J-Doracle.install.cvu.getSharedPartitionList=false \
-J-Doracle.install.cvu.checkSharedStorageFileSystemType=false \
-J-Doracle.install.grid.validate.QuickInstallUI=false
```

Note: The option "-J-Doracle.install.grid.validate.QuickInstallUI=false" skips validations on the "Specify Install Locations" page.

Ensure that you provide correct inputs for all the fields on this page, that is, manually verify the inputs before proceeding with the installation.

- For “Advanced Installation” of Oracle Grid Infrastructure 11.2.0.4, enter:

```
# ./runInstaller -J-Doracle.install.cvu.getSharedPartitionList=false \
-J-Doracle.install.cvu.checkSharedStorageFileSystemType=false \
-J-Doracle.install.grid.validate.OCRStorageUI=false \
-J-Doracle.install.grid.validate.VDSKStorageUI=false
```

Note: The options skip all the checks for Oracle Cluster Registry and voting disk storage. Ensure that you manually verify the location that is provided before proceeding with the installation.

See http://www.symantec.com/docs/TECH213369 for more details.

Oracle Grid Infrastructure installation may fail with internal driver error

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

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

Workaround::

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

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

  ```
  export OUI_ARGS=-ignoreInternalDriverError
  ```

  For more information, see the Oracle Metalink document: 970166.1

- Web-based installer
  When you run the Web-based installer, in the **Enter the arguments to be passed to the Oracle installer** text box, enter the value `-ignoreInternalDriverError`.
  For more information, see the *Symantec Storage Foundation for Oracle RAC Installation and Configuration Guide*.

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

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

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

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

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

Enabling ODM in Oracle RAC 11 Release 2 installations causes errors (1913013)

Enabling ODM in Oracle RAC 11 Release 2 installations causes the following error:

```
'ODM ERROR V-41-4-1-253-12 Not enough space'
Oracle instance may also crash with same error.
```

The error is observed if the DISM (Dynamic Intimate Shared memory) feature is enabled. In Solaris, the Oracle database uses DISM if it is available on the system, and if the value of the `sga_max_size` initialization parameter is larger than the size required for all SGA components combined.

Workaround: Make sure that the file `ORACLE_HOME/bin/oradism` is owned by the root user with "execute" and "setuid" permissions. If the problem persists after correcting the permissions, uncomment the `sga_max_size` and `memory_target` `init.ora` parameters.

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), "net0" 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.

Oracle Database Configuration Assistant displays an error

The Database Configuration Assistant utility displays the following error:

```
SGA size cannot be greater than maximum shared memory
segment size (0).
```

Workaround: Ignore this message and manually configure the database memory parameters for Oracle. In the "Memory" tab of the Oracle Database Creation Assistant (DBCA), select a Custom and Manual shared memory management configuration and enter the appropriate values.

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

CSSD agent forcibly stops Oracle Clusterware if Oracle Clusterware fails to respond (3352269)

On nodes with heavy load, the CSSD agent attempts to check the status of Oracle Clusterware till it reaches the`FaultOnMonitorTimeouts` value. However, Oracle Clusterware fails to respond and the CSSD agent forcibly stops Oracle Clusterware. To prevent the CSSD agent from forcibly stopping Oracle Clusterware, set the value of the `FaultOnMonitorTimeouts` attribute to 0 and use the `AlertOnMonitorTimeouts` attribute as described in the following procedure.

Perform the following steps to prevent the CSSD agent from forcibly stopping Oracle Clusterware:

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

```
# haconf -makerw
```

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

```
# hatype -display CSSD | grep AlertOnMonitorTimeouts
CSSD  AlertOnMonitorTimeouts 0
# hares -override cssd_resname AlertOnMonitorTimeouts
# hatype -modify CSSD AlertOnMonitorTimeouts 4
```

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

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

4 Verify the `AlertOnMonitorTimeouts` and `FaultOnMonitorTimeouts` settings:

```
# hatype -display CSSD | egrep "AlertOnMonitorTimeouts|FaultOnMonitorTime
CSSD  AlertOnMonitorTimeouts 4
CSSD  FaultOnMonitorTimeouts 0
```

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

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

Unable to configure Highly Available IP (HAIP) using the web installer (3348812)

During HAIP configuration, the SFRAC web installer fails to update the `/etc/hosts` file with the HAIP alias.

Workaround:

Use one of the following options:

- Use the SFRAC script installer to configure HAIP.
- Use the SFRAC web installer, but add the IP addresses and aliases in the `/etc/hosts` file manually.

Intelligent Monitoring Framework (IMF) entry point may fail when IMF detects resource state transition from online to offline for CSSD resource type (3287719)

When IMF detects a state transition from ONLINE to OFFLINE state for a registered online resource, it sends a notification to the CSSD agent. The CSSD agent schedules a monitor to confirm the state transition of the resource. The resources of type CSSD takes more time to go online or offline fully. Therefore, if this immediate monitor finds the resource still in online state, it assumes that the IMF notification is false and attempts to register the resource in online state again.

In such partial state transitions, the agent repeatedly attempts to register the resource until the`RegisterRetryLimit` is reached (default value is 3) or the resource registration is successful. After the resource is completely offline, the next resource registration with IMF will be successful.

Workaround: Increase the value of the `RegisterRetryLimit` attribute if multiple registration attempts fail.

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.

Warning message displayed on taking cssd resource offline if LANG attribute is set to "eucJP" (2123122)

When you take the cssd resource offline using the `hares -offline cssd` command and the LANG attribute is set to “eucJP”, the following message may be observed in the `hamsg engine_A` command output:

```
VCS INFO V-16-2-13716 Could not find message V-16-2-13716
```

You may ignore the message.

PrivNIC resource faults in IPMP environments on Solaris 11 systems (2838745)

The PrivNIC resource faults on Solaris 11 systems when private interfaces used by IPMP are configured under PrivNIC resource.

Workaround: Avoid using PrivNIC or MultiPrivNIC agents in IPMP environments.

Error displayed on removal of VRTSjadba language package (2569224)

Removal of the VRTSjadba language package displays the following error on the screen:

```
Executing postremove script.
Generating BMC map file...
bmcmap ERROR V-33-1000-10001 Unable to create BMC map
```

You may ignore the error.

Veritas Volume Manager can not identify Oracle Automatic Storage Management (ASM) disks (2771637)

Veritas Volume Manager (VxVM) commands can not identify disks that are initialized by ASM. Administrators must use caution when using the VxVM commands to avoid accidental overwriting of the ASM disk data.

Oracle Universal Installer fails to start on Solaris 11 systems (2784560)

The Oracle Universal Installer (OUI) fails to start when the SF Oracle RAC installer invokes the OUI for the installation of Oracle Clusterware/Grid Infrastructure software.

Workaround: Install the following packages before installing Oracle Clusterware/Grid Infrastructure.

```
SUNWxwplt
SUNWmfrun
```

For instructions, see the Oracle documentation.

Symantec Storage Foundation for Sybase ASE CE known issues

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

Sybase Agent Monitor times out (1592996)

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

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

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

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

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

Installer warning (1515503)

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

```
Error: CPI WARNING V-9-40-5460 The quorum file /qrmmnt/qfile
cannot be accessed now. This may be due to a file system not being mounted.
```

The above warning may be safely ignored.

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

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

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

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

Unexpected node reboot when invalid attribute is given (2567507)

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

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

Bus error while stopping the ports (2358568)

Problem: When the `hastop -local` command or the `hastop -all` command is issued, the `fuser -kill` command is issued on the mounted Sybase mount point. This results in bus error and a core dump, though the ports stop cleanly.

Resolution: Before issuing the `hastop -local` command or the `hastop -all` command, ensure that the `uafstartup.sh` script is stopped, so that the `fuser -kill` command is not issued on the mounted Sybase mount point.

Symantec ApplicationHA known issues

App.RestartAttempts setting does not take effect if value is set to 2 or more [2508392]

App.RestartAttempts configuration option defines the number of times Symantec ApplicationHA tries to restart a failed application or its component. Its value can range from 1 to 6.

For certain application configurations, this setting fails to take effect if its value is set to 2 or more. After successfully configuring an application, if there is a fault in the application or its dependent component, ApplicationHA attempts to restart it once. If the application fails to start, ApplicationHA reports the application state as faulted.

This issue is applicable only for the following applications/components:

On Solaris

- Custom Application
- Apache HTTP Server

Workaround

Currently there is no workaround to resolve this issue.

Symantec recommends that for applications mentioned earlier, you set the App.RestartAttempts value to 1.

This ensures that ApplicationHA makes at least one attempt to restart the failed component. If the component still fails to start, ApplicationHA then declares it as faulted and takes further action as per the configuration settings.

ApplicationHA fails to work if VRTSsfmh is uninstalled [2361128]

The VRTSsfmh package (Managed Host component for Veritas Operations Manager) is installed on the managed domain as part of ApplicationHA installation. VRTSsfmh contains the 'Veritas Storage Foundation Messaging Service' (xprtld) that is used by both, ApplicationHA and VOM. If VRTSsfmh is uninstalled, ApplicationHA functionality fails.

Workaround

Perform the following steps

1. Insert the ApplicationHA software disc into your system drive and navigate to the directory that contains the package for the Solaris SPARC 10 operating system:

   ```
   # cd cdrom_root/applicationha/sol10_sparc/pkgs
   ```

2. Run the following command:

   ```
   # pkgadd -a VRTSsfmh.pkg
   ```

3. Stop the `xprtld` service.

   ```
   # svcadm disable svc:/system/xprtld:default
   ```

4. Ensure that the file `/etc/opt/VRTSsfmh/xprtld.conf` contains the following text:

   ```
   namespaces vcs=/opt/VRTSvcs/portal
   ```

5. Start the `xprtld` service.

   ```
   # svcadm enable svc:/system/xprtld:default
   ```

Refreshing the Symantec High Availability view multiple times displays a network connectivity error [2379946, 2379707]

This issue is typically observed in case of IE7 browser.

Symantec High Availability view refreshes the application status every 60 seconds. However, in case of network failure if you manually refresh the ApplicationHA view multiple times, IE displays a network connectivity error.

If you click **Ok** on the error message and then click another virtual machine on the VOM Management Server or vSphere Web Client (in VMware environments), then the Symantec High Availability view displays the application status of an unknown application.

This issue also occurs if you refresh the Symantec High Availability view and simultaneously reset the virtual machine.

Workaround

For details, refer to the following knowledge base article from Microsoft.

http://support.microsoft.com/kb/927917#more_information

VCS configuration incorrectly retains read-write mode [2607134]

When you execute the enable_applicationha script on the control domain, if an error occurs, the script exits. However, the VCS configuration remains in the read-write mode. In this mode, the configuration is vulnerable to unintentional editing.

Workaround

Revert the VCS configuration to the read-only mode by using the following command:

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

Configuration option of ApplicationHA installer malfunctions [2621468]

When you run the Symantec ApplicationHA installer, it displays the following option to configure ApplicationHA: **Configure an Installed Product**.

If you specify this option, the installer fails to configure ApplicationHA. Instead, the installer starts stopping certain ApplicationHA processes.

Workaround

Do not use the installer option to configure an application. Instead, to configure Symantec ApplicationHA for monitoring an application, use one of the following methods:

- If you have already installed ApplicationHA, navigate to the following URL, and use the **Configure Application Monitoring** link to launch the Symantec High Availability Configuration Wizard:

  ```
  https://<logicalDomainNameorIPaddress>:5634/vcs/admin/
  application_health.html?priv=ADMIN
  ```

Where logicalDomainNameorIPaddress is the system name or IP address of the logical domain (Oracle VM) where you want to configure application monitoring.

- You can launch the wizard from the Symantec High Availability view of the Veritas Operations Manager Management Server Console.
  For more information on working with VOM and accessing the Symantec High Availability view, see the *Symantec ApplicationHA User Guide* or the *Veritas Operations Manager User Guide*.

Heartbeat service group may fail to come online [2605506]

If the high availability daemon (HAD) on the guest domain is restarted, the configured heartbeat service group (VCSAppMonHBSG) does not automatically come online.

Workaround

To continue application monitoring, you must manually bring the VCSAppMonHBSG online by using the following command:

```
# /opt/VRTSvcs/bin/hagrp -online VCSAppMonHBSG -sys System
```

Where *System* is name of the guest domain.

Attributes of a virtual machine may retain stale values [2611726]

If the physical host crashes, the virtual machines may indicate stale values for attributes such as ConnectionState and SysState. The settings are updated after a virtual machine fails over to a new physical host.

Attributes of a guest domain may retain stale values [2611726]

If the physical host crashes, the guest domains may indicate stale values for attributes such as ConnectionState and SysState. The settings are updated after a guest domain fails over to a new physical host.

Install program does not provide keyless licensing option during upgrade [3335745, 3336308]

If you try to upgrade to ApplicationHA 6.1 but do not have a valid license key already installed, the install program prompts you to specify a valid license key. In this step, the install program presently does not provide you with the option to specify keyless licensing.

The issue is also seen when you upgrade to ApplicationHA 6.1.1 using Install Bundles. Other upgrade methods do not trigger the issue.

Workaround:

Perform the following steps:

1. When the installer prompts you to specify a license key, if you want to specify keyless licensing, enter 'q' to quit the install program.
2. From the command line, execute the following command:

   # /opt/VRTS/install/installapplicationha61 -license *sys1*

   Where sys 1 is the name of the system where you want to enable keyless licensing.

ApplicationHA installer displays incorrect EULA path for non-English locales [3344863]

During ApplicationHA 6.1 installation on guests running Linux, AIX, or Solaris SPARC operating systems, the installer prompts the user to accept the End User's License Agreement (EULA) at following locations for Japanese and Chinese locales, respectively:

`applicationha/EULA/ja/`

`applicationha/EULA/zh/`

No EULA exists at the said locations.

Workaround

ApplicationHA 6.1 users must read and accept the following EULA:

applicationha/EULA/en/EULA_ApplicationHA_Ux_6.1.pdf

LLT known issues

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

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.

Cannot configure LLT if full device path is not used in the lltab file (2858159)

(Oracle Solaris 11) On virtual machines ensure that you use the full path of the devices corresponding to the links in lltab. For example, use /dev/net/net1 instead of /dev/net/net:1 in the lltab file, otherwise you cannot configure LLT.

Fast link failure detection is not supported on Solaris 11 (2954267)

Fast link failure detection is not supported on Solaris 11 operating system because the operating system cannot provide notification calls to LLT when a link failure occurs. If the operating system kernel notifies LLT about the link failure, LLT can detect a link failure much earlier than the regular link failure detection cycle. As Solaris 11 does not notify LLT about link failures, failure detection cannot happen before the regular detection cycle.

Workaround: None

GAB known issues

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

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

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

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

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

Workaround: There is no workaround for this issue.

Cluster panics during reconfiguration (2590413)

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

Workaround: There is no workaround for this issue.

GAB may fail to stop during a phased upgrade on Oracle Solaris 11 (2858157)

While performing a phased upgrade on Oracle Solaris 11 systems, GAB may fail to stop. However, CPI gives a warning and continues with stopping the stack.

Workaround: Reboot the node after the installer completes the upgrade.

Cannot run pfiles or truss files on gablogd (2292294)

When pfiles or truss is run on gablogd, a signal is issued to gablogd. gablogd is blocked since it has called an gab ioctl and is waiting for events. As a result, the pfiles command hangs.

Workaround: None.

(Oracle Solaris 11) On virtual machines, sometimes the common product installer (CPI) may report that GAB failed to start and may exit (2879262)

GAB startup script may take longer than expected to start up. The delay in start up can cause the CPI to report that GAB failed and exits.

Workaround: Manually start GAB and all dependent services.

I/O fencing known issues

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

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 cooridnator disk group

Workaround: There is no workaround for this issue.

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

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

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

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

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

Workaround: Start VxFEN again after some time.

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

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

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

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

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

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

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

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

```
operation failed.
CPS ERROR V-97-1400-446 Un-authorized user cpsclient@sys1,
domaintype vx; not allowing action
```

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

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

When I/O fencing is not up, the svcs command shows VxFEN as online (2492874)

Solaris 10 SMF marks the service status based on the exit code of the start method for that service. The VxFEN start method executes the vxfen-startup script in the background and exits with code 0. Hence, if the vxfen-startup script subsequently exits with failure then this change is not propagated to SMF. This behavior causes the `svcs` command to show incorrect status for VxFEN.

Workaround: Use the `vxfenadm` command to verify that I/O fencing is running.

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

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

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

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

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

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

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

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

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

CP server repetitively logs unavailable IP addresses (2530864)

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

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

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

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

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

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

If user configures I/O fencing in Sybase mode and an old `/opt/VRTSvcs/lock/.vxfendargs` file is present in the system from an earlier configuration of I/O fencing in customized mode, then VCS attempts to restart the vxfend daemon every time it is killed. This interferes with the functioning of I/O fencing in the Sybase mode.

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

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.

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 [3226290]

The issue exists because the VCS 5.1SP1 release version does not support separate directories for truststores. However, VCS version 6.0 and later support separate directories for truststores. 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 so that the servers and client systems can communicate in a secure mode.

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

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

Workaround:

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

The vxfenswap utility deletes comment lines from the `/etc/vxfemode` file, if you run the utility with hacli option (3318449)

The vxfenswap utility uses RSH, SSH, or hacli protocol to communicate with peer nodes in the cluster. When you use vxfenswap to replace coordination disk(s) in disk-based fencing, vxfenswap copies `/etc/vxfenmode` (local node) to `/etc/vxfenmode` (remote node).

With the hacli option, the utility removes the comment lines from the remote `/etc/vxfenmode` file, but, it retains comments in the local `/etc/vxfenmode` file.

Workaround: Copy the comments manually from local `/etc/vxfenmode` to remote nodes.

When you configure CP server only for HTTPS-based communication, the `engine_A.log` displays a misleading message (3321101)

The `engine_A.log` file displays the following message when you configure CP server only for HTTPS-based communication but not for IPM-based communication.

```
No VIP for IPM specified in /etc/vxcps.conf
```

Workaround: Ignore the message.

The `vxfentsthdw` utility may not run on systems installed with partial SFHA stack [3333914]

The `vxfentsthdw` utility runs if the SFHA stack and VCS are fully installed with properly configured SF and VxVM. It also runs if the entire SFHA stack and VCS are not installed. However, partial installs where SF is installed and configured but VCS is not installed is not supported. The utility will display an error with the `-g` or `-c` options.

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

Fencing configuration fails if SysDownPolicy is set to AutoDisableNoOffline in online service groups [3335137]

If SysDownPolicy of one or more online service groups is configured to AutoDisableNoOffline, fencing configurations such as server-based, disk-based and disable mode fail. Since the service groups is configured with `SysDownPolicy = { AutoDisableNoOffline }`, stopping VCS fails which leads to the failure of fencing configuration.

Workaround: When configuring fencing and before stopping VCS, you must offline the service groups configured with `SysDownPolicy = { AutoDisableNoOffline }` manually.

When a client node goes down, for reasons such as node panic, I/O fencing does not come up on that client node after node restart (3341322)

This issue happens when one of the following conditions is true:

- Any of the CP servers configured for HTTPS communication goes down.
- The CP server service group in any of the CP servers configured for HTTPS communication goes down.
- Any of the VIPs in any of the CP servers configured for HTTPS communication goes down.

When you restart the client node, fencing configuration starts on the node. The fencing daemon, vxfend, invokes some of the fencing scripts on the node. Each of these scripts has a timeout value of 120 seconds. If any of these scripts fails, fencing configuration fails on that node.

Some of these scripts use `cpsadm` commands to communicate with CP servers. When the node comes up, `cpsadm` commands try to connect to the CP server using VIPs for a timeout value of 60 seconds. So, if the multiple `cpsadm` commands that are run within a single script exceed the timeout value, then the total timeout value exceeds 120 seconds, which causes one of the scripts to time out. Hence, I/O fencing does not come up on the client node.

Note that this issue does not occur with IPM-based communication between CP server and client clusters.

Workaround: Fix the CP server.

Delay in rebooting Solaris 10 nodes due to vxfen service timeout issues (1897449)

When you reboot the nodes using the `shutdown -i6 -g0 -y` command, the following error messages may appear:

```
svc:/system/vxfen:default:Method or service exit
timed out. Killing contract 142
svc:/system/vxfen:default:Method "/lib/svc/method/vxfen stop"
failed due to signal Kill.
```

This error occurs because the vxfen client is still active when VCS attempts to stop I/O fencing. As a result, the vxfen stop service times out and delays the system reboot.

Workaround: Perform the following steps to avoid this vxfen stop service timeout error.

To avoid the vxfen stop service timeout error

1 Stop VCS. On any node in the cluster, run the following command:

```
# hastop -all
```

2 Reboot the systems:

```
# shutdown -i6 -g0 -y
```

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

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

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

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

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

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

Unable to customize the 30-second duration (2551621)

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

Workaround: There is no workaround for this issue.

CP server configuration fails while setting up secure credentials for CP server hosted on an SFHA cluster (2621029)

When you configure CP server using the `configure_cps.pl` utility, the configuration fails while setting up secure credentials for CP server that is hosted on an SFHA cluster. You may see the following error:

```
Creating softlink to credential directory /etc/VRTScps/db/CPSERVER
on node nodename.
Unable to connect to node nodename using /usr/bin/ssh.
Please configure ssh communication and retry. Exiting.
```

Workaround: You can use any of the following options:

- Before running the `configure_cps.pl` utility, change the default shell for root user to either KSH or bash.
- Perform the following steps after running the `configure_cps.pl` utility on each node of the cluster:
  - Manually remove the old credential directory or softlink. For example:

    ```
    # rm -rf /var/VRTSvcs/vcsauth/data/CPSERVER
    ```

  - Create a new soft-link to the shared location of the credential directory:

    ```
    # ln -s path_of_CP_server_credential_directory \
     /var/VRTSvcs/vcsauth/data/CPSERVER
    ```

  - Start the CPSSG service group:

    ```
    # hagrp -online CPSSG -any
    ```

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

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

Workaround: To resolve this issue, perform the following procedure on all of the nodes of the CP server:

1 Rename `cpsadm` to `cpsadmbin`:

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

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

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

3 Change the permissions of the new file to 775:

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

Coordination point server-based fencing may fail if it is configured on 5.1SP1RP1 using 6.0.1 coordination point servers (2824472)

The 5.1SP1 installer (CPI) cannot set up trust between a 5.1SP1 client and a 6.0 or later server, because there are no separate directories for truststores in the 5.1SP1. When trust cannot be setup, the 5.1SP1 installer cannot configure 5.1SP1 clients to work with 6.0 or later CPS in secure mode.

Workaround:

Set up trust manually between the CPS and clients using the cpsat or the vcsat command. After that, CPS and client will be able to communicate properly in the secure mode.

Software limitations

This section covers the software limitations of this release.

Limitations related to installation

This is the limitations related to installation in the 6.1.1 release.

Limitations related to web-based installer for SF Oracle RAC

- Web-based installer on local disk is not supported.
- If SF Oracle RAC is not configured before upgrade, the web-based installer does not support to upgrade SF Oracle RAC to 6.1.1.

Veritas Volume Manager software limitations

The following are software limitations in this release of Veritas Volume Manager.

Converting a multi-pathed disk

When converting a multi-pathed disk that is smaller than 1 TB from a VTOC label to an EFI label, you must issue the `format -e` command for each path. For example, if a node has two paths, `c1t2d0s2` and `c2tsd0s2`, you must run the `format -e` command on each of the two paths.

SFHA Solutions does not support thin reclamation of space on a linked mirror volume (2729563)

The thin reclamation feature does not support thin reclamation for a linked mirror volume.

Thin reclamation requests are not redirected even when the ioship policy is enabled (2755982)

Reclamation requests fail from nodes that do not have local connectivity to the disks, even when the ioship policy is enabled. Reclamation I/Os are not redirected to another node.

Veritas Operations Manager does not support disk, disk group, and volume state information related to CVM I/O shipping feature (2781126)

The Veritas Operations Manager (VOM) does not support disk, disk group, and volume state information related to the I/O shipping feature introduced in this release of Cluster Volume Manager. New states such as lfailed, lmissing or LDISABLED are introduced when I/O shipping is active because of storage disconnectvity.

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment,change the default values for the DMP tunable parameters.

Table 1-29 describes the DMP tunable parameters and the new values.

Table 1-29 DMP settings for NetApp storage attached environment

Parameter name	Definition	New value	Default value
dmp_restore_interval	DMP restore daemon cycle	60 seconds.	300 seconds.
dmp_path_age	DMP path aging tunable	120 seconds.	300 seconds.

The change is persistent across reboots.

To change the tunable parameters

1 Issue the following commands:

```
# vxdmpadm settune dmp_restore_interval=60

# vxdmpadm settune dmp_path_age=120
```

2 To verify the new settings, use the following commands:

```
# vxdmpadm gettune dmp_restore_interval

# vxdmpadm gettune dmp_path_age
```

Snapshot configuration with volumes in shared disk groups and private disk groups is not supported

A snapshot configuration with volumes in the shared disk groups and private disk groups is not a recommended configuration. In this release, this configuration is not supported.

A 1 TB disk that is not labeled using operating system commands goes into an error state after the vxconfigd daemon is restarted

A 1 TB disk that is not labeled using operating system commands goes into an error state after the vxconfigd daemon is restarted.`vxconfigd` daemon is restarted.

Currently, a solution from the vendor is not available.

Veritas File System software limitations

The following are software limitations in this release of Veritas File System.

Recommended limit of number of files in a directory

To maximize VxFS performance, do not exceed 100,000 files in the same directory. Use multiple directories instead.

The vxlist command cannot correctly display numbers greater than or equal to 1 EB

The `vxlist` command and all of the other commands that use the same library as the `vxlist` command cannot correctly display numbers greater than or equal to 1 EB.

Limitations with delayed allocation for extending writes feature

The following limitations apply to the delayed allocation for extending writes feature:

- In the cases where the file data must be written to disk immediately, delayed allocation is disabled on that file. Examples of such cases include Direct I/O, concurrent I/O, FDD/ODM access, and synchronous I/O.
- Delayed allocation is not supported on memory mapped files.
- Delayed allocation is not supported with BSD quotas. When BSD quotas are enabled on a file system, delayed allocation is turned off automatically for that file system.
- Delayed allocation is not supported for shared mounts in a cluster file system.

FlashBackup feature of NetBackup 7.5 (or earlier) does not support disk layout Version 8, 9, or 10

The FlashBackup feature of NetBackup 7.5 (or earlier) does not support disk layout Version 8, 9, or 10.

Compressed files that are backed up using NetBackup 7.1 or prior become uncompressed when you restore the files

The NetBackup 7.1 release and prior does not support the file compression feature. If you back up compressed files using NetBackup 7.1 or a prior release, the files become uncompressed when you restore the files.

When an I/O domain fails, the vxdisk scandisks or vxdctl enable command take a long time to complete (2791127)

When an I/O domain fails, the vxdisk scandisks or vxdctl enable from the logical domain (LDOM) guest take a long time to complete. `vdc_ioctls` like `DKIOCGGEOM`

and DKIOCINFO also take more time to return. These issues seem to be due to retry operations performed at the Solaris operating system layer.

Reducing the vdc_timeout value to lower value might help to bring down time. Dynamic multi-pathing (DMP) code is optimized to avoid making such vdc_ioctl calls in an LDOM guest environment as much possible. This change considerably reduces delays.

A complete resolution to this issue may require changes at the Solaris operating system level.

Replication software limitations

The following are replication software limitations in this release of Symantec Storage Foundation and High Availability Solutions.

VVR Replication in a shared environment

Currently, replication support is limited to 8-node cluster applications.

VVR IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

- A replication configuration from an IPv4-only node to an IPv6-only node and from an IPv6-only node to an IPv4-only node is not supported, because the IPv6-only node has no IPv4 address configured on it and therefore VVR cannot establish communication between the two nodes.
- A replication configuration in which an IPv4 address is specified for the local_host attribute of a primary RLINK and an IPv6 address is specified for the remote_host attribute of the same RLINK.
- A replication configuration in which an IPv6 address is specified for the local_host attribute of a primary RLINK and an IPv4 address is specified for the remote_host attribute of the same RLINK.
- IPv6 is not supported in a CVM and VVR cluster where some nodes in the cluster are IPv4-only and other nodes in the same cluster are IPv6-only, or all nodes of a cluster are IPv4-only and all nodes of a remote cluster are IPv6-only.
- VVR does not support Edge and NAT-PT routers that facilitate IPv4 and IPv6 address translation.

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 6.1 and the prior major releases of Storage Foundation (6.0 and 6.0.1). Replication between versions is

supported for disk group versions 170, 180, and 190 only. Both the Primary and Secondary hosts must be using a supported disk group version.

Limitations related to I/O fencing

This section covers I/O fencing-related software limitations.

Preferred fencing limitation when VxFEN activates RACER node re-election

The preferred fencing feature gives preference to more weighted or larger subclusters by delaying the smaller subcluster. This smaller subcluster delay is effective only if the initial RACER node in the larger subcluster is able to complete the race. If due to some reason the initial RACER node is not able to complete the race and the VxFEN driver activates the racer re-election algorithm, then the smaller subcluster delay is offset by the time taken for the racer re-election and the less weighted or smaller subcluster could win the race. This limitation though not desirable can be tolerated.

Stopping systems in clusters with I/O fencing configured

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect, or “split brain.” See the *Symantec Cluster Server Administrator's Guide* for a description of the problems a failed interconnect can create and the protection I/O fencing provides.

In a cluster using SCSI-3 based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on both the data disks and coordinator disks. In a cluster using CP server-based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on data disks and similar registrations on CP server. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordination points and data disks to prevent possible difficulties with subsequent cluster startup.

Using the reboot command rather than the shutdown command bypasses shutdown scripts and can leave keys on the coordination points and data disks. Depending on the order of reboot and subsequent startup events, the cluster may warn of a possible split brain condition and fail to start up.

Workaround: Use the shutdown -r command on one node at a time and wait for each node to complete shutdown.

Uninstalling VRTSvxvm causes issues when VxFEN is configured in SCSI3 mode with dmp disk policy (2522069)

When VxFEN is configured in SCSI3 mode with dmp disk policy, the DMP nodes for the coordinator disks can be accessed during system shutdown or fencing arbitration. After uninstalling VRTSvxvm package, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm package is uninstalled, if VxFEN attempts to access DMP devices during shutdown or fencing arbitration, the system panics.

Symantec Storage Foundation for Databases (SFDB) tools software limitations

The following are the SFDB tools software limitations in this release.

Oracle Data Guard in an Oracle RAC environment

SFDB tools cannot be used with RAC standby databases. SFDB tools can still be used with the primary database, even in a Data Guard Oracle RAC environment.

Upgrading to Oracle 10.2.0.5 is required if using SFDB tools

If you are running Oracle version 10.2.0.4 and upgrading a Storage Foundation product with SFDB tools to 6.1.1, you must upgrade the Oracle binaries and database to version 10.2.0.5, before upgrading to 6.1.1.

Parallel execution of `vxsfadm` is not supported (2515442)

Only one instance of the `vxsfadm` command can be run at a time. Running multiple instances of `vxsfadm` at a time is not supported.

Creating point-in-time copies during database structural changes is not supported (2496178)

SFDB tools do not support creating point-in-time copies while structural changes to the database are in progress, such as adding or dropping tablespaces and adding or dropping data files.

However, once a point-in-time copy is taken, you can create a clone at any time, regardless of the status of the database.