Symantec™ Storage Foundation and High Availability Solutions 6.2.1 Release Notes - Solaris

Maintenance Release

Solaris 10, 11 (SPARC)

April 2015

Symantec

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

Product version: 6.2.1

Document version: 6.2.1 Rev 0

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 6.2.1 Maintenance Release (6.2.1 MR).

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

http://www.symantec.com/docs/TECH225259

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 *Symantec Storage Foundation and High Availability Solutions 6.2.1 Installation Guide* 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: 6.2.1 Rev 0" 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)

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
151218-04	VRTSsfmh	ApplicationHA, DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS, VM	94MB
151230-01	VRTSvxfs	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE	39 MB
151231-01	VRTSodm	SF, SF Oracle RAC, SFCFSHA, SFHA	1.1 MB
151232-01	VRTSvxvm	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VM	212 MB
151237-01	VRTSsfcpi62	ApplicationHA, DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS, VM	14 MB
151240-01	VRTSperl	ApplicationHA, DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS, VM	63 MB
151241-01	VRTSfsadv	FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE	15MB
151242-01	VRTSvcs	ApplicationHA, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	249 MB
151243-01	VRTSllt	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	1.8 MB

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

Patch ID	Package Name	Products Affected	Patch Size
151244-01	VRTSgab	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	1.7 MB
151245-01	VRTScps	SF Oracle RAC, SFCFSHA, SFHA, VCS	54 MB
151246-01	VRTSvcsag	ApplicationHA, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	7.6 MB
151247-01	VRTSvcsea	ApplicationHA, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	2.1 MB
151250-01	VRTSvcsvmw	ApplicationHA	13 MB
151251-01	VRTScavf	SF Oracle RAC, SFCFSHA, SFSYBASECE	984 MB
151253-01	VRTSdbed	SF, SF Oracle RAC, SFCFSHA, SFHA	221 MB
151255-01	VRTSvcswiz	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	6.6 MB
151772-01	VRTSspt	ApplicationHA, DMP, FS, SF, SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS, VM	25 MB
151775-01	VRTSvxfen	SF Oracle RAC, SFCFSHA, SFHA, SFSYBASECE, VCS	2.5 MB

Table 1-2 Package for Solaris 11

Package Name	Included Patches	Products Affected	Package Size
VRTSpatches.p5p	VRTSamf VRTSaslapm VRTScps VRTSperl VRTSfsadv VRTScavf VRTSspt VRTSdbed VRTSgab VRTSllt VRTSodm VRTSsfcpi62 VRTSvcs VRTSvcsag VRTSvcsea VRTSvxfen VRTSvxfs VRTSvxvm	DMP, SF, SF Oracle RAC, SFCFSHA, SFHA, VCS	300 M

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/TECH225259
- 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/TECH225258

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 6.2.1

This section lists the changes in Symantec Storage Foundation and High Availability Solutions 6.2.1.

Changes related to VCS

Symantec Cluster Server (VCS) includes the following changes in 6.2.1:

Solaris IPS driver action support for VCS kernel components

On Solaris 11, the LLT, GAB, VxFEN and AMF kernel components of VCS now support using Solaris IPS (Image Packaging system) driver action to add or remove kernel drivers.

Behaviour changes to AMF and SMF because of the driver action support

On Solaris 11, prior to this release, the AMF driver used to be unloaded and removed while disabling the AMF service. From this release, AMF has started using the driver action to add or remove kernel driver. With this change, the AMF kernel driver is installed and loaded during the package installation. When you enable the AMF service by using the `svcadm enable amf` command, the driver gets configured. When you disable the AMF service by using the `svcadm disable amf` command, the driver gets un-configured, and is not un-loaded. A new configuration variable *AMF_DISABLE* is added in the `/etc/default/amf` file. The default value of AMF_DISABLE is [0]. To disable AMF completely, which is to unload and to remove the driver from the system, set the value of AMF_DISABLE to *1* before disabling the AMF service using the `svcadm disable amf` command. Set the value back to *0* to enable the AMF service again. AMF service fails to start if VCS_DISALBE value is *1*.

Support for SAP ASE 16 in single instance mode

In this release, Symantec Cluster Server (VCS) is enhanced to support SAP ASE (previously Sybase) 16 in single instance mode.

Support for Oracle Solaris 11.2 kernel zones

Symantec supports Oracle Solaris 11.2 kernel zones (brand solaris-kz) only with Symantec Cluster Server (VCS) 6.2.1. Symantec does not support Oracle Solaris 11.2 kernel zones with Storage Foundation 6.2.1.

Changes related to SmartIO for solid-state drives

In this release, the SmartIO feature of Storage Foundation and High Availability Solutions (SFHA Solutions) is enhanced to support Flexible Storage Sharing (FSS) with SmartIO for remote caching.

FSS works with SmartIO to provide caching services for nodes that do not have local SSD devices.

In this scenario, Flexible Storage Sharing (FSS) exports SSDs from nodes that have a local SSD. FSS then creates a pool of the exported SSDs in the cluster. From this shared pool, a cache area is created for each node in the cluster. Each cache area is accessible only to that particular node for which it is created. The cache area can be of type, VxVM or VxFS.

The cluster must be a CVM cluster.

The volume layout of the cache area on remote SSDs follows the simple stripe layout, not the default FSS allocation policy of mirroring across host. If the caching operation degrades performance on a particular volume, then caching is disabled for that particular volume. The volumes that are used to create cache areas must be created on disk groups with disk group version 200 or later. However, data volumes that are created on disk groups with disk group version 190 or later can access the cache area created on FSS exported devices.

Note: CFS write-back caching is not supported for cache areas created on remote SSDs.

Apart from the CVM/CFS license, the SmartIO license is required to create cache areas on the exported devices.

Changes related to Veritas Volume Manager

Veritas Volume Manager (VxVM) includes the following changes in 6.2.1:

Remote Caching

Generally, SSDs are not put into every server in cluster, but you may like to use pool of available SSDs across all servers for caching purpose.

In this release, Remote Caching allows you to use SSDs from another node(s) for caching purpose. Flexible Storage Sharing (FSS) feature enables cluster-wide sharing of local storage through the use of a network interconnect among the nodes of the cluster. Integration of SmartIO and FSS feature allows you to utilize the SmartIO feature on node(s) in cluster, which doesn't have locally attached SSDs.

Solaris 11.2 Feature Support

In this release, SFHA Solutions 6.2.1 supports Solaris 11.2 Feature.

Bioclone and ppmapin changes to enable Solaris 11.2 feature

In this release, we use Oracle's implementation of bioclone function instead of vxvm_bioclone function.

On Soalris11, Solaris kernel used to do paging twice, which affected VxVM performance badly, especially on vxio.

Oracle provides new biclone() interface function on Solaris11U2 which addresses all paging related issues ppmapin and bioclone.

So VxVM implements its own vxvm_bioclone() function which does the similar functionality of Solaris's bioclone and additionally addresses the repeated paging problem on VxVM.

Changes related to Dynamic Multi-Pathing

Dynamic Multi-Pathing (DMP) includes the following changes in 6.2.1:

DMP support for NetApp Data On TAP 8.x cluster

In this release, DMP is enhanced to optimally support multi-controller (> 2 controllers) NetApp Data On TAP 8.x cluster mode configurations. This release addresses DMP limitations with supporting NetApp Data On TAP 8.x cluster mode configurations.

System requirements

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 (SRU) 11.1.21.4.1 Solaris 11.2 and up to Support Repository Updates (SRU) 11.2.8.4	SPARC

This release (version 6.2.1) is not supported on the x86-64 architecture.

This release (version 6.2.1) supports native zones 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 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, and OVM 3.2.

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.

Supported database software

For the latest information on supported database, see the following TechNote:http://www.symantec.com/docs/DOC4039

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/TECH211575

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*.

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

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

Installation and upgrade fixed issues in 6.2.1

Table 1-5 covers the incidents that are fixed related to installation and upgrade in 6.2.1.

Table 1-5 Installation and upgrade 6.2.1 fixed issues

Incident	Description
3638757	SMF of VCS is disabled after rolling upgrade to 6.0.5 on Solaris11.

Table 1-5 Installation and upgrade 6.2.1 fixed issues *(continued)*

Incident	Description
3661652	When VCS is upgraded from version 6.0.5 to 6.2, the new types of agents such as types.cf are not detected.
3719159	In Oracle RAC version 12.1.0.2, warning messages are reported during the post configuration checks.
3739179	The module OpenSSL 1.0.1i in the VRTSperl package has security issues.

Installation and upgrade fixed issues in 6.2

Table 1-6 covers the incidents that are fixed related to installation and upgrade in 6.2.

Table 1-6 Installation and upgrade 6.2 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.
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.
3325954	On Solaris 10 `xprtld` will not be started if user use jumpstart to install product
3326196	Rolling upgrade may encounter a problem if open volumes from different disk groups have the same name.

Table 1-6 Installation and upgrade 6.2 fixed issues *(continued)*

Incident	Description
3326639	CP server service group fails to come online with the default database path after the CP server is upgraded from 6.0 to 6.2 on a multi-node cluster.
3341674	For Solaris 11.1 or later, the system can panic when system is rebooted after turning dmp_native_support to on.
3343592	`installer -license sys1` command does not show the check result if the system already has the SF Basic license key registered.
3442070	If you select rolling upgrade task from the Install Bundles menu, the Installer exits with an error.

Symantec Storage Foundation and High Availability fixed issues

This section includes issues fixed for Symantec Cluster Server, Veritas File System, and Veritas Volume Manager.

See "Symantec Cluster Server fixed issues" on page 21.

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

See "Veritas Volume Manager fixed issues" on page 33.

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 6.2.1

Table 1-7 describes the Symantec Storage Foundation for Databases (SFDB) tools issues fixed in 6.2.1.

Table 1-7 SFDB tools 6.2.1 fixed issues

Incident	Description
3676518	For Oracle database 12.1.02 the new odm path is not being considered by dbed_codm_adm

Table 1-7 SFDB tools 6.2.1 fixed issues *(continued)*

Incident	Description
3676521	dbed_codm_adm commands fail with an error when run as an Oracle user.

Symantec Storage Foundation for Databases (SFDB) tools fixed issues in 6.2

Table 1-8 describes the Symantec Storage Foundation for Databases (SFDB) tools issues fixed in 6.2.

Table 1-8 SFDB tools 6.2 fixed issues

Incident	Description
2869266	Checkpoint clone fails if the archive log destination is same as the datafiles destination.
3313775	SmartIO options are not restored after Reverse Resync Commit operation is performed.
3615735	During a Reverse Resync Begin operation, a mismatch in database control file version is observed.
3615745	For thin storage setups, the snapshot operation reports that the diskgroup cannot be split.
3615764	The flashSnap operation fails to create a symlink on a Symantec Volume Replicator (VVR) secondary site.

Symantec Cluster Server fixed issues

This section describes Symantec Cluster Server fixed issues.

Symantec Cluster Server fixed issues in 6.2.1

Table 1-9 covers the fixed issues of Symantec Cluster Server in 6.2.1.

Table 1-9 Symantc Cluster Server 6.2.1 fixed issues

Incident	Description
3520211	The HostMonitor agent reports incorrect memory usage.
3662501	The notifier process fails to clearly distinguish whether the CPU, Memory, or Swap has exceeded the warning or critical threshold level.

Table 1-9 Symantc Cluster Server 6.2.1 fixed issues *(continued)*

Incident	Description
3662508	The CmdServer log incorrectly reports warning messages relevant to licensing even when keyless licensing is enabled and the host is configured under Veritas Operations Manager (VOM).
3666049	VCS does not support SAP ASE 16 Sybase agent.
3668853	The halog(1M) command becomes unresponsive when VCS is in the REMOTE_BUILD state.
3683648	The clean entry point of Zpool agent reports incorrect status.
3753194	The AMF driver remains loaded after disabling AMF service on Solaris 11.

Symantec Cluster Server fixed issues in 6.2

This section describes Symantec Cluster Server fixed issues in 6.2.

VCS engine fixed issues

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

Table 1-10 VCS engine fixed issues

Incident	Description
3381042	The checkboot utility core dump or time difference between a system and Network Time Protocol (NTP) time leads to unexpected deletion of the temporary files. The deletion causes the VCS agents to report an incorrect state.
3448510	The `hastatus` command fails and dumps core when it is run from a local zone.
3468891	Inconsistencies in `/etc/VRTSvcs/conf/attributes/cluster_attrs.xml` file and hide resource-level ContainerInfo attribute from `hares -display` command.
3211834	CurrentLimits attribute value is not updated correctly when a service group faults.
3385820	Sometimes the high availability daemon (HAD) crashes if it runs for a long duration.

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

Incident	Description
3436617	When invoking triggers if some arguments are left unassigned, the `hatrigger` command fails due to compilation errors.
3471819	The service group fails to go online if the CurrentCount attribute value is incorrect.
3498072	The hazonesetup (1M) utility reports a Perl warning message when the locale is set to non-English.
3464981	The file size of Engine_A.log file could not be increased beyond 32 MB.
3580940	VCS configuration becomes unclean when incorrect filename is provided with the `ha` command for SourceFile attribute.
3603275	HAD and other nodes abort while shutting down two nodes simultaneously.

Bundled agents fixed issues

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

Table 1-11 Bundled agents fixed issues

Incident	Description
2490296	Application agent cannot handle a case with user as root, envfile set and shell as `csh`.
2618482	Some agents may fail to come online after full upgrade to VCS 6.0 if they were online before the upgrade.
3536195	The Online operation of LDom resource fails when Disaster Recovery options are configured. This is because the online entry point fails to get the definition of a function.
3326591	The IPMultiNICB agent delays bringing IPv4 VIP online on Solaris 10 by 5 seconds.
3621042	NIC agent might incorrectly report NIC resource state as offline.
3590419	On Solaris, the MultiNICB agent does not add the default route when the resource state is ONLINE in the first probe.
3593137	Mount agent unmounts the incorrect file system due to incorrect pattern matching.

Table 1-11 Bundled agents fixed issues *(continued)*

Incident	Description
3422904	Zone agent fails to bring the Zone resource to online state when the state of the zone is unavailable.
3576701	The Share resource fails to unshare the path name when file system is unmounted.
3505202	VCS does not support non-default value for VCS_LOG environment variable.

Fixed issues related to AMF

Table 1-12 AMF fixed issues

Incident	Description
2848007	The libvxamf library encounters an error condition while doing a process table scan.
3333913	AMF may panic the system if it receives a request to unregister an already unregistered resource.
3407338	If one of the events of a group is in triggered state, then the group fails to unregister because of the triggered event.
3606494	During a branded zone boot operation, under certain circumstances the system panics because of memory corruption in Asynchronous Monitoring Framework.
3338946	Sometimes when a process offline registration is requested and the system is under heavy load, AMF library fails to verify whether the resource is actually offline. As a result, registration fails.

Cluster configuration wizard fixed issues

Table 1-13 lists the cluster configuration wizard fixed issues.

Table 1-13 Wizard fixed issues

Incident	Description
3593390	In a Solaris 11 setup, if you execute the deprecated ‘pkginfo’ command for the VRTSwiz package, an incorrect error message appears, stating that VRTSwiz information cannot be found.

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 6.2.1

There is no fixed issue for Symantec Storage Foundation for Oracle RAC in 6.2.1.

Symantec Storage Foundation for Oracle RAC fixed issues in 6.2

Table 1-14 lists the issues fixed in 6.2.

Table 1-14 Issues fixed in 6.2

Incident	Description

3589293	Configuration of the CSSD resource for SF Oracle RAC in Solaris zone environments fails with the following error if Intelligent Monitoring Framework (IMF) is enabled for the CSSD resource:
	
	`VCS ERROR V-16-2-13710`
	`Resource(cssd) - imf_register entry point failed with exit code(1)`

3321004	Installation of Oracle Clusterware using Oracle response file fails.
3348812	During HAIP configuration, the SF Oracle RAC web installer fails to update the /etc/hosts file with the HAIP alias.

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 6.2.1

This section lists the incidents that are fixed in Symantec Storage Foundation Cluster File System High Availability in 6.2.1.

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

See "Veritas Volume Manager fixed issues" on page 33.

Table 1-15 Symantec Storage Foundation Cluster File System High Availability 6.2.1 fixed issues

Incident	Description
3604071	High CPU usage consumed by the vxfs thread process.
3233276	With a large file system, primary to secondary migration takes longer duration.
3656155	Import of VxVM diskgroup triggered from CVMVolDg agent fails with error message.

Symantec Storage Foundation Cluster File System High Availability fixed issues in 6.2

Table 1-16 lists the incidents that are fixed in Symantec Storage Foundation Cluster File System High Availability in 6.2.

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

See "Veritas Volume Manager fixed issues" on page 33.

Table 1-16 Symantec Storage Foundation Cluster File System High Availability 6.2 fixed issues

Incident	Description
3642894	VxFS agents for VCS should honor the customization of `VCS_LOG`.
3582470	Data change object (DCO) volume gets created using the same node's disks for both plexes.
3573908	Multiple `cbrbk.tmp$$` files in the /var/tmp folder on each node do not clean up properly.
3552008	The `vxconfigrestore` (vxvol resync) operation hangs on the master node while recovering a stripe-mirror volume.
3551050	The `vx*` commands are not able to connect to vxconfigd.
3547833	With Storage and I/Os from all three nodes, an I/O hang occurs.
3538683	After a master panic, the new master does not start plex resync when the older master comes online and joins the cluster.
3537519	The `vxdisk unexport` command hangs and then `vxconfigd` gets fault in an asymmetric array connection setup.
3534779	Internal stress testing on Cluster File System (CFS) hits a debug assert.

Table 1-16 Symantec Storage Foundation Cluster File System High Availability 6.2 fixed issues *(continued)*

Incident	Description
3528908	When the slave node contributing storage left the cluster, the `vxconfigd` command dumps core on the master node.
3523731	VxVM command check for device compliance prior to doing FSS operations on disks.
3517695	Storage failure in the FSS environment causes an I/O hang on all nodes.
3511444	Panics occur while checking memory pressure.
3508390	DCO object is unnecessarily marked BADLOG mode in cascade failure scenarios, and it results in requiring a full-recovery and can result in lost snapshots as well.
3505017	When the `da name` is the same as the existing `dm name`, the `vxdg addisk` operation from slave fails.
3496673	On a Flexible Storage Sharing (FSS) disk group, the read I/O performance is impacted.
3495811	When you create a disk group, the disk shows LMISSING state when the SCSI PGR operation fails.
3489167	Plex(es) on remote disks goes to DISABLED state because of a plex I/O error encountered after slave node reboot in cluster volume manger.
3484570	Accessing CVM message after decrementing reference count causes a panic.
3449152	The `vxtunefs`(1M) command fails to set the `thin_friendly_alloc` tunable in cluster file systems.
3444771	Internal noise test on cluster file system hits an debug assert when you are creating a file.
3430256	Space allocation for Volume: Single DAS disk in a disk group takes more preference than shared storage in that disk group.
3422704	Unique prefix generation algorithm redesign for forming cluster wide consistent name (CCN).
3411438	The `preferred` and `round-robin` read policy settings should be honoured irrespective of local connectivity to the plexes.
3394940	Anomaly numbers are displayed in the `vxstat` output.

Table 1-16 Symantec Storage Foundation Cluster File System High Availability 6.2 fixed issues *(continued)*

Incident	Description
3383625	When a cluster node that contributes the storage to the Flexible Storage Sharing (FSS) disk group rejoins the cluster, the local disks brought back by that node do not get reattached.
3373747	Adding new nodes to the 22-node cluster causes Cluster File System (CFS) failures after CVM deletes 2 nodes.
3370753	Internal tests with SmartIO writeback SSD cache hit debug asserts.
3370722	Operating system (OS) installation on virtual machines fails in two-node clusters with writeback caching enabled.
3329603	The `vxconfigd` related error messages are observed in system log files on every node for large cluster setups.
3301085	The `vxdg adddisk` operation fails when adding nodes containing disks with the same name.
3300418	VxVM volume operations on shared volumes cause unnecessary read I/Os.
3286030	Vxattachd debug messages get displayed on the console during a reboot.
3283518	Disk group deport operation reports messages in the syslog for remote disks.
3283418	Remote writes hang due to heavy sync workload on the target node in FSS environments.
3281160	When autoreminor is set off, no error is thrown when you import a disk group having the same minor number as that of the existing imported disk group.
3214542	Disk group creation or addition of a new disk to an existing disk group fails with "VxVM vxdg ERROR V-5-1-16087 Disk for disk group not found" error when the command is executed from the slave node.
3213411	A node join fails if a "disabled" Flexible Storage Sharing disk group exists in the cluster.
3198590	SmartIO VxVM caching is not supported for CFS.
3191807	CVM requires the T10 vendor provided ID to be unique.

Table 1-16 Symantec Storage Foundation Cluster File System High Availability 6.2 fixed issues *(continued)*

Incident	Description
3152304	When connectivity to some of the plexes of a volume is lost from all nodes, an I/O hang occurs.
3142109	The CFSMountAgent script does not parse the `MountOpt` properly if the `-O` overlay (native fs option) is used.
3117153	Disk group remains in deported state for remote disks after destroying the disk group from a node that is not exporting the disks.

3093407	When one host name in a cluster is a substring of other host name, with the superstring differentiated by a hyphen – for example, `abc` and `abc-01` – the `cfsmnatadm` utility may throw an error for some commands as follows:
	`# cfsmntadm modify /swapmnt/ add abc="rw"`
	`Nodes are being added...`
	`Error: V-35-117: abc already associated with cluster-mount`
	`Nodes added to cluster-mount /swapmnt`
	`#`

3079819	`vxconfigbackup` fails on Flexible Storage Sharing disk groups.
1203819	In some cases, inode and allocation maps inconsistencies occur in the event of a node crash in clusters.
640213	The node panics in case of overlapping reconfigurations due to race conditions.

Veritas File System fixed issues

This section describes the fixed issues of Veritas File System.

Veritas File System fixed issues in 6.2.1

Table 1-17 lists the incidents that are fixed in Veritas File System (VxFS) in 6.2.1.

Table 1-17 Veritas File System 6.2.1 fixed issues

Incident	Description
2059611	The system panics due to a NULL pointer dereference while flushing bitmaps to the disk.

Table 1-17 Veritas File System 6.2.1 fixed issues *(continued)*

Incident	Description
2439261	When the vx_fiostats_tunable value is changed from zero to non-zero, the system panics.
2560032	System panics after SFHA is upgraded from 5.1SP1 to 5.1SP1RP2 or from 6.0.1 to 6.0.5
3269553	VxFS returns inappropriate message for read of hole via Oracle Disk Manager (ODM).
3525858	The system panics when the Oracle Disk Manager (ODM) device (/dev/odm) is mounted in a Solaris Zone.
3567027	During the File System resize operation, the "fullfsck flag is set.
3601943	Truncating corrupted block map of a file may lead to an infinite loop.
3602322	Panic while flushing the dirty pages of the inode
3604750	The kernel loops during the extent re-org.
3615043	Data loss when writing to a file while dalloc is on.
3617191	Checkpoint creation takes a lot of time.
3621205	OpenSSL common vulnerability exposure (CVE): POODLE and Heartbleed.
3622323	Cluster Filesystem mounted as read-only panics when it gets sharing and/or compression statistics with the fsadm_vxfs(1M) command.
3767771	CVMVOLDG agent does not deport shared disk groups when these disks go into the disable state.
3647749	On Solaris, an obsolete v_path is created for the VxFS vnode.
3657482	Stress test on cluster file system fails due to data corruption
3689104	The module version of the vxcafs module is not displayed after the modinfo vxcafs command is run on Solaris.
3697964	The vxupgrade(1M) command fails to retain the fs_flags after upgrading a file system.
3715566	VxFS fails to report an error when the maxlink and nomaxlink options are set on file systems having disk layout version (DLV) lower than 10.
3719523	'vxupgrade' retains the superblock replica of old layout versions.

Table 1-17 Veritas File System 6.2.1 fixed issues *(continued)*

Incident	Description
3721466	After a file system is upgraded from version 6 to 7, the vxupgrade(1M) command fails to set the VX_SINGLEDEV flag on a superblock.
3725346	Trimming of underlying SSD volume was not supported for AIX and Solar is using "fsadm -R -o ssd" command.
3729104	Man pages changes missing for smartiomode option of mount_vxfs (1M)
3731678	During an internal test, a debug assert was observed while handling the error scenario.
3736772	The sfcache(1M) command does not automatically enable write-back caching on file system once the cache size is increased to enable write-back caching.
3739618	sfcache command with "-i" option maynot show filesystem cache statistic periodically.
3743912	Users could create sub-directories more than 64K for disk layouts having versions lower than 10.
3744424	OpenSSL common vulnerability exposure (CVE): POODLE and Heartbleed.
3756750	VxFS may leak memory when File Design Driver (FDD) module is unloaded before the cache file system is taken offline.

Veritas File System fixed issues in 6.2

Table 1-18 lists the incidents that are fixed in Veritas File System (VxFS) in 6.2.

Table 1-18 Veritas File System 6.2 fixed issues

Incident	Description
3641719	The `fallocate` may allocate a highly fragmented file when the size of the file is extremely large.
3597482	The `pwrite(2)` function fails with the `EOPNOTSUPP` error.
3589264	The `fsadm` command shows an incorrect option for the file system type in usage.
3563796	The file system `fullfsck` flag is set when the inode table overflows.

Table 1-18 Veritas File System 6.2 fixed issues *(continued)*

Incident	Description
3560968	The `delicache_enable` tunable is not persistent in the Cluster File System (CFS) environment.
3560187	The kernel may panic when the buffer is freed in the `vx_dexh_preadd_space()` function with the message `Data Key Miss Fault in kernel mode`.
3550103	After you upgrade or restart the system, mismatch in SSD cache usage may occur.
3520349	When there is a huge number of dirty pages in the memory, and a sparse write is performed at a large offset of 4 TB or above on an existing file that is not null, the file system hangs
3484336	The `fidtovp()` system call can panic in the `vx_itryhold_locked()` function.
3478017	Internal tests found assert in voprwunlock
3469644	The system panics in the `vx_logbuf_clean()` function.
3466020	The file system is corrupted with the error message `vx_direrr: vx_dexh_keycheck_1`.
3463464	Internal kernel functionality conformance tests hit a kernel panic due to null pointer dereference.
3457803	The file system gets disabled intermittently with metadata I/O errors.
3451284	While allocating extents during the write operation, if summary and bitmap data for file system allocation units get mismatched, a full `fsck` flag get set on the file system.
3449150	The `vxtunefs`(1M) command accepts garbage values for certain tunables.
3448702	Checkpoint creation of different file systems may get serialized.
3444775	Internal noise testing on cluster file system results in a kernel panic in `vx_fsadm_query()` function with an error message.
3434811	The `vxfsconvert`(1M) command in VxFS 6.1 hangs.
3424564	`fsppadm` fails with `ENODEV` and `file is encrypted or is not a database` errors.

Table 1-18 Veritas File System 6.2 fixed issues *(continued)*

Incident	Description
3417076	The `vxtunefs`(1M) command fails to set tunables when the file contains blank lines or white spaces.
3415639	The type of the `fsdedupadm`(1M) command always shows as MANUAL even when it is launched by the `fsdedupschd` daemon.
3413926	Internal testing hangs due to high memory consumption, resulting in fork failures.
3394803	A panic is observed in the VxFS routine `vx_upgrade7`() function while running the `vxupgrade` command (1M).
3340286	After a file system is resized, the tunable setting `dalloc_enable` is reset to a default value.
3335272	The `mkfs` (make file system) command dumps core when the log size provided is not aligned.
3332902	While shutting down, the system running the `fsclustadm`(1M) command panics.
3323866	Some Object Data Manager (ODM) operations fail with `ODM ERROR V-41-4-1-328-22 Invalid argument`.
3317368	File system operations needing a file system freeze may take longer in the presence of file level snapshots and when there is a heavy I/O load.
3301716	In a Veritas File System (VxFS), if a file system has compression enabled, the file system gets disabled due to the ENOSPC error. This occurs because of a defect in the delayed allocation feature.
3297840	VxFS corruption is detected during a dynamic LUN resize operation.
3264062	The allocated space may be leaked from the free space while you are unsharing shared extents.
3121933	There is a DB2 crash or corruption when `EOPNOTSUPP` is returned from VxFS.

Veritas Volume Manager fixed issues

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

Veritas Volume Manager fixed issues in 6.2.1

Table 1-19 lists the incidents that are fixed in Veritas Volume Manager (VxVM) in 6.2.1.

Table 1-19 Veritas Volume Manager 6.2.1 fixed issues

Incident	Description
2965910	Volume creation with the vxassist(1M) command dumps core when the non-disk parameters are specified along with the '-o ordered' option.
3322760	"vxdisk list" command may list detached disks, if the options are fullshared, partialshared, or local.
3414023	If a volume has preexisting snapshots, reattaching a detached plex would resync the extra data.
3420347	Write errors occur with "vxdg flush" running.
3424704	Some VxVM commands fail on using the localized messages.
3442024	The vxdiskadm(1M) option #1(adding disk to disk group and initialization) does not support Extensible Firmware Interface(EFI) disks on Solaris operating system.
3506450	When migrating from PowerPath to Dynamic Multipathing (DMP), you may observe a system panic or VxVM configuration daemon (vxconfigd) core.
3518470	The 'vxdg adddisk' command fails with the following message: "*disk name* previous removed using -k option" Even though the disk was not removed using the -k option.
3564260	VVR commands are unresponsive when replication is paused and resumed in a loop.
3571434	DMP does not support NetApp cDOT cluster mode array with more than 2 nodes.
3588012	The vxconfigd(1M) daemon experiences a memory leak while printing I18N messages.
3594158	The spinlock and unspinlock are referenced to different objects when interleaving with a kernel transaction.
3599977	During a replica connection, referencing a port that is already deleted in another thread causes a system panic.
3616671	The number of transaction log files is reset to default value 10, even though it is explicitly set to the no_limit value.

Table 1-19 Veritas Volume Manager 6.2.1 fixed issues *(continued)*

Incident	Description
3621232	The vradmin ibc command cannot be started or executed on Veritas Volume Replicators (VVR) secondary node.
3625890	vxdisk resize operation on CDS disks fails with an error message of "Invalid attribute specification"
3628860	The vxdisk(1M) command shows disks with status as nolabel in the outputs.
3628933	Volumes remain in the DETACHED state, even after nodes with storage join back to the cluster.
3631025	The I/O statistics derived using the vxstat utility with "-ft" or "-fi" options along with -i option do not provide the system name and its origin (source or destination).
3631989	The vxconfigd(1M) daemon becomes unresponsive on the joiner node.
3644291	Some of the disk groups that are composed by 2TB LUNs fail to be auto-imported after a system restart.
3644687	In a Veritas Volume Replicator (VVR) environment, any node on the Primary cluster may panic during the I/O error handling.
3645370	vxevac command fails to evacuate disks with Dirty Region Log(DRL) plexes.
3646712	The Huawei ASL sometimes skips claiming a path to a disk as it identifies it incorrectly as a command device.
3648719	The server panics while adding or removing LUNs or HBAs.
3651034	Temporary loss of storage results in application I/O failure.
3651422	"sfcache app" command fails for Sybase application on Solaris.
3656539	I/O may hang when SmartIO VxVM block level is configured in only few nodes of cluster, while there are volumes with instant snapshots.
3658079	If the size of backup slice of volumes is larger than 64GB, Veritas Volume manager (VxVM) wraps it around it and exports it to LDOM (Logical Domains) guest domain.
3669863	"sfcache list" command doesn't show storage devices in the output if "cachearea" is created on stripe-mirror volume.

Table 1-19 Veritas Volume Manager 6.2.1 fixed issues *(continued)*

Incident	Description
3671975	The system panics when Veritas Volume manager (VxVM) cachearea is configured in Cluster Volume Manager (CVM).
3672759	The vxconfigd(1M) daemon may core dump when DMP database is corrupted.
3674614	Restarting the vxconfigd(1M) daemon on the slave (joiner) node during node-join operation may cause the vxconfigd(1M) daemon to become unresponsive on the master and the joiner node.
3677359	VxDMP (Veritas Dynamic MultiPathing) causes system panic after a shutdown or reboot.
3688156	The output of vxdisk -o local|partialshared list command lists disks in the error state.
3719478	ZFS zpool auto LUN expansion is not detected when dmp_native_support is enabled.
3727939	Data corruption may occur due to stale device entries in the /dev/vx/[r]dmp directories.
3736156	Solaris 11 update 1 system fails to come up after enabling dmp_native_support and reboot.
3747701	In Cluster Volume Manager (CVM) environment, node-join process fails and node remains in joining state.
3749557	System hangs because of high memory usage by vxvm.

Veritas Volume Manager fixed issues in 6.2

Table 1-20 lists the incidents that are fixed in Veritas Volume Manager (VxVM) in 6.2.

Table 1-20 Veritas Volume Manager fixed issues

Incident	Description
3622068	After mirroring an encapsulated root disk, the rootdg disk group fails to get imported if any disk in the disk group becomes unavailable.
3614182	The first system reboot after migration from Solaris Multi-Pathing (MPXIO) to Symantec Dynamic Multi-Pathing (DMP) native takes a long time.

Table 1-20 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
3603792	In Solaris 11 SRU 16, Postinstall temporarily stops for three hours when upgrading to SFHA 6.2.
3584311	The vxconfigd daemon hangs with "vol_rv_transaction_prepare+0005C8" on secondary site.
3580962	A panic occurs in VxDMP under high I/O load, and may cause complete storage disconnection.
3577957	Database instance is terminated while rebooting a node.
3573262	System is crashed by vxio when running snapshot operations on solaris (sparc) servers.
3566493	Orphan cpmap objects cannot be removed after you disassociate unfinished snap plexes.
3555230	The `vxconfigd` daemon hangs in Veritas Volume Replicator (VVR) when writing to SRL volume during replication.
3554608	Mirroring a volume creates a larger plex than the original on a CDS disk.
3553407	The SmartMove feature cannot work on layered volumes that do not have thin disks.
3544980	vxconfigd V-5-1-7920 di_init() failed message after SAN tape online event.
3544972	620:dmp:coredump while rebooting the OS after dmp installation.
3542272	The `vxconfigbackupd` daemon never exits after reboot. The daemon remains active for a disk group because configuration has been changed after the backup is initiated.
3539548	Duplicate disks and I/O error occurs after dynamic LUN allocation.
3521726	When using Symantec Replication Option, system panics happens due to double freeing IOHINT memory.
3513392	Secondary panics when rebooted while heavy IOs are going on primary.
3506336	Address deadlock between message processing on DR and Quiescing of IOs.
3503852	With multiple Replicated Volume Groups (RVGs), if you detach storage on a secondary master then reattach it back, rlinks are not able to connect. The rlink state is different on one of the three rlinks. .

Table 1-20 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
3498228	The `vxconfigd` core dump occurs after port disable or enable operation with migration from PP to DMP.
3495553	DV:6.1 The `vxconfigd` daemon hangs on secondary in vol_ru_transaction_prepare.
3495548	The `vxdisk rm` command fails with devices when Operating System Naming (OSN) scheme is used for devices controlled by the EMC Powerpath.
3490458	After managing class under PP, some of the devices are seen in error state.
3489572	Slave nodes panic when volume with DCO hits storage failure while volume is online.
3482026	The `vxattachd`(1M) daemon reattaches plexes of manually detached site.
3478019	When VxVM fails to assign a unique name to a new DCL volume, the `vxsnap prepare` command fails silently without giving an error.
3475521	During a system reboot, the following error message is displayed on the console: `es_rcm.pl:scripting protocol error`.
3456153	When Veritas Volume Replicator (VVR) replication is in progress, a Cluster Volume Manager (CVM) slave node reboot causes an I/O hang.
3455460	The `vxfmrshowmap` and the `verify_dco_header` utilities fail.
3450758	The slave node was not able to join CVM cluster and resulted in panic.
3446415	A pool may get added to the file system when the file system shrink operation is performed on FileStore.
3440790	The `vxassist` command with parameter mirror and the vxplex command(1M) with parameter att hang.
3428025	When heavy parallel I/O load is issued, the system that runs Symantec Replication Option (VVR) and is configured as VVR primary crashes.
3417044	System becomes unresponsive while creating a VVR TCP connection.
3415188	I/O hangs during replication in Veritas Volume Replicator (VVR).
3411668	Network and host endian difference is not handled in the nmcom_print_sock_storage() function.
3403390	After a crash, the linked-to volume goes into NEEDSYNC state.

Table 1-20 Veritas Volume Manager fixed issues *(continued)*

Incident	Description
3399131	For PowerPath (PP) enclosure, both DA_TPD and DA_COEXIST_TPD flags are set.
3385905	Data corruption occurs after VxVM makes cache area offline and online again without a reboot.
3385753	Replication to the Disaster Recovery (DR) site hangs even though Replication links (Rlinks) are in the connected state.
3374200	A system panic or exceptional IO delays are observed while executing snapshot operations, such as, refresh.
3368361	When siteconsistency is configured within a private disk group (with LUNs mapped only to local server) and CVM is up, then the reattach operation of a detached site fails.
3326964	VxVM hangs in Clustered Volume Manager (CVM) environments in the presence of FMR operations.
3317430	The `vxdiskunsetup` utility throws error after upgrade from 5.1SP1RP4.
3287940	LUNs from any EMC CLARiiON arrays that have Not Ready state are shown in the "online invalid" state by Veritas Volume Manager (VxVM).
3279932	The `vxdisksetup` and `vxdiskunsetup` utilities fail for disks that are part of a deported disk group, even if "-f" option is specified.
	Heavy I/O loads on primary sites result in transaction/session timeouts between the primary and secondary sites.
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.
2049952	The `vxrootadm` command shows incorrect messages in Japanese with localization for Solaris.
1390029	The `vxconfigrestore` command fails when there is a dot in the disk group name, i.g., test.2

Symantec ApplicationHA fixed issues

This section describes the fixed issues of Symantec ApplicationHA.

Symantec ApplicationHA 6.2.1 fixed issues

Table 1-21 lists the fixed issue for Symantec ApplicationHA in 6.2.1.

Table 1-21 Symantec ApplicationHA 6.2.1 fixed issues

Incident	Description
3640275	ApplicationHA fails to exit maintenance mode and resume application monitoring

Symantec ApplicationHA 6.2 fixed issues

Table 1-22 Symantec ApplicationHA 6.2 fixed issues

Incident number	Description
3335745	While upgrading ApplicationHA, the install program does not provide users with the option to specify keyless licensing.
3491170	ApplicationHA installer displays incorrect EULA path for non-English locales
3491158	If you have configured Oracle database instances for detail monitoring with ApplicationHA 6.0 or earlier, you cannot perform a live upgrade to ApplicationHA 6.1.

Symantec Dynamic Multi-Pathing fixed issues

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

Dynamic Multi-Pathing fixed issues in 6.2.1

There is no fixed issue in 6.2.1.

Dynamic Multi-Pathing fixed issues in 6.2

Table 1-23 lists the fixed issues for Dynamic Multi-Pathing 6.2.

Table 1-23 Dynamic Multi-Pathing 6.2 fixed issues

Incident	Description
3565212	IO failure during controller giveback operations on Netapp FAS31700 in ALUA mode.

Table 1-23 Dynamic Multi-Pathing 6.2 fixed issues *(continued)*

Incident	Description
3544980	vxconfigd V-5-1-7920 di_init() failed message after SAN tape online event.
3544972	620:dmp:coredump while rebooting the OS after dmp installation.
3543284	FIO device not visible.
3542713	vxdmpadm listenclosure all displays a different ENCL from array console/VOM.
3526500	DMP I/O getting timeout lot earlier than 300 seconds if I/O statistics daemon is not running.
3520991	vxconfigd core dumps during vxdisk scandisks.
3502923	ESX panic while running add/remove devices from smartpool with no license installed on server.
3399323	The reconfiguration of DMP DB failed.
3373208	DMP wrongly sends APTPL bit 0 to array.

Veritas Operations Manager fixed issues

This section describes the incidents that are fixed related to Veritas Operations Manager (VOM).

Veritas Operations Manager fixed issues in 6.2.1

Table 1-24 covers the incidents that are fixed related to Veritas Operations Manager (VOM) in 6.2.1.

Table 1-24 Veritas Operations Manager fixed issues in 6.2.1

Incident	Description
3517052	mismatch between VOM and Sort for SPVU calculation.
3526358	VOM 6.0 statistics 'Live' option shows incorrect 'free memory'.
3526792	VOM doesn't display RVG and host relationships properly in CVR environment.
3554764	Disk type resources from CVM were being shown twice in VOM availability perspective.
3558088	Switch configuration fails if password has single quotes (').

Table 1-24 Veritas Operations Manager fixed issues in 6.2.1 *(continued)*

Incident	Description
3560777	Issues with disk discovery when LUNZ devices are present.
3563150	VVR RVGs displayed on the wrong host in CMS Console with host aliasing.
3576824	vxdclid is opening too many file descriptors.
3577772	Layout is not displayed for volumes created by LVM.
3585566	Keyless License not detected.
3593632	Base Release displays Incorrect for SFHA-RHEL6_x86_64-6.2.
3615769	Add host failing with MH having english-UK locale.
3629745	VxFS file system report usage is not getting updated, reports are incorrect.
3636983	VOM 6.1 shows mh status as "Not Installed"
3642699	Set credential operation for secure databases fails.
3651166	VOM console sees Solaris 11 SFHA 6.1.1 as 6.1
3658760	Failed to remove MHs from VOM GUI and vomadm cli with --force
3663528	xprtld daemon cannot be started after install on RHEL 7.
3671234	HBA vendor name discovery corrupts the SF family sfm.dump and this results in filling up DB logs with error messages.
3688166	Adding host fails in VOM 6.1.

Veritas Operations Manager fixed issues in 6.2

There is no fixed issue for Veritas Operations Manager fixed issues in 6.2.

Cluster Volume Manager fixed issues

Table 1-25 describes the incidents that are fixed in Cluster Volume Manager (CVM) in 6.2.

Table 1-25 Cluster Volume Manager 6.2 fixed issues

Incident	Description
640213	The node panics in case of overlapping reconfigurations due to race conditions.
3592783	For the FSS partially shared storage configuration, after you run `hastop -all` and then `hastart`, remote disks, which are not part of the disk group, are not visible in `vxdisk -o alldgs list`.
3528908	When the slave node contributing storage left the cluster, the `vxconfigd` command dumps core on the master node.
2705055	The `preferred` and `round-robin` read policy settings should be honoured irrespective of local connectivity to the plexes.

Vxexplorer tool fixed issue

This section describes the incidents that are fixed related to vxexplorer tool fixed issue.

Vxexplorer tool fixed issue in 6.2.1

Table 1-26 covers the incidents that are fixed related to vxexplorer tool fixed issue in 6.2.1.

Table 1-26 vxexplorer tool fixed issue in 6.2.1

Incident	Description
3739916	On VRTSspt 6.2.0, the VRTSexplorer fails on Solaris 10 for SPARC platform after the patch 119783-25 is applied.
3739942	On Veritas Support Package 6.2.0, for vxexplorer, an End of Service Life (EOSL) message is displayed in error..
3739946	Veritas Support Package6.2.0,installs the vxgetcore utility on a wrong path

Vxexplorer tool fixed issue in 6.2

There is no fixed issue for Vxexplorer tool fixed issue in 6.2.

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 6.2.1

Table 1-27 lists the fixed issues for LLT, GAB, and I/O fencing in 6.2.1.

Table 1-27 LLT, GAB, and I/O fencing 6.2.1 fixed issues

Incident	Description
3642031	In the postinstall phase, GAB may get stuck trying to run devlinks.
3663740	In a rare scenario, divide by zero error is seen when lltshow -p command is run.
3667755	Strict permission check for CPS configuration files.
3691202	Sometimes fencing in the customized mode fails to start when the number of files present in the current working directory is large.
3698093	The system may panic when a client of GAB calls the GAB API to get port parameters.

LLT, GAB, and I/O fencing fixed issues in 6.2

Table 1-28 lists the fixed issues for LLT, GAB, and I/O fencing in 6.2.

Table 1-28 LLT, GAB, and I/O fencing fixed issues

Incident	Description
3156922	The CP server process, vxcpserv, communicates with client nodes only on those VIPs that are available when CP server process starts.
3335137	Fencing configuration fails if SysDownPolicy is set to AutoDisableNoOffline in online service groups.
3473104	When virtual NICs are configured under LLT without specifying the MTU size 1500 in lltab, cluster does not function properly. For example, VCS engine commands may hang and print below message in the engine logs: `VCS CRITICAL V-16-1-51135 GlobalCounter not updated`
3548629	On Solaris 11, LLT, GAB and I/O fencing modules fails to configure when installed on an alternate boot environment.

Table 1-28 LLT, GAB, and I/O fencing fixed issues *(continued)*

Incident	Description
3331801	SMF services for VCS kernel components may go into maintenance state when installed in a new boot environment.
3031216	The dash (-) in a disk group name causes vxfentsthdw(1M) and Vxfenswap(1M) utilities to fail.
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.

Known issues

This section covers the known issues in this release.

- Issues related to installation and upgrade
- Symantec Storage Foundation known issues
- Symantec Storage Foundation and High Availability known issues
- Symantec Cluster Server known issues
- Symantec Dynamic Multi-Pathing known issues
- Symantec Storage Foundation Cluster File System High Availability known issues
- Symantec Storage Foundation for Oracle RAC known issues
- Symantec Storage Foundation for Databases (SFDB) tools known issues
- Symantec Storage Foundation for Sybase ASE CE known issues
- Symantec ApplicationHA known issues
- LLT known issues
- I/O fencing known issues
- GAB known issues

Issues related to installation and upgrade

This section describes the known issues during installation and upgrade.

On Oracle Solaris, drivers may not be loaded after stop and then reboot [3763550]

When the installer stops the processes, it uses the `rem_drv` command to prevent the drivers from being loaded back by the operating system. However, OS cannot load those drivers back after reboot, such as vxdmp and vxio.

Workaround: Start your product manually:

```
# /opt/VRTS/install/install<prod>62 -start
```

During an upgrade process, the AMF_START or AMF_STOP variable values may be inconsistent [3763790]

If the value of AMF_START or AMF_STOP variables in the driver configuration file is '0' before an upgrade, then after the upgrade is complete, the installer changes the value to 1. Simultaneously, the installer also starts the Asynchronous Monitoring Framework (AMF) process.

Workaround: To resolve the issue, stop the AMF process and change the AMF_START or AMF_STOP value to **0**.

On Solaris, the vxcafs driver may fail to load after upgrading the SF stack from version 6.2 to 6.2.1 with the installer response file [3747647]

When SF stack is upgraded from 6.2 to 6.2.1 using the response file, vxcafs driver fails to load due to which VxFS cache area fails to come online.

Workaround:

To resolve this issue

1. Execute the following on all nodes:

   # add_drv vxcafs

2. Bring the cache online by executing the following on all nodes:

   # sfcache online <cachearea_name>

3. Restart the system.

After a patch upgrade to version 6.2.1, the installer may fail to stop the Asynchronous Monitoring Framework (AMF) process [3737299]

After a patch upgrade to 6.2.1 version from 6.2, when the product is stopped using `/opt/VRTS/install/install<prod>62-stop`, the AMF process does not stop.

Workaround: There is no workaround for this issue.

Uninstallation fails on global zone if product packages are installed on both global zone and local zone [3762814]

If you have a product installed on both local zone and global zone, if you uninstall the product on global zone, the packages fails to be uninstalled.

Workaround: Log into the local zone and uninstall the packages of product from the local zone first.

On Solaris 10, incorrect module versions of VxFS and FDD are observed when VRTSvxfs is installed using the Install Bundles or upgraded to 6.2.1 from version older than 6.2.0 [3763728]

Due to some driver unload issue, the VxFS, and FDD module versions appears as 6.2.0 instead of 6.2.1.

Workaround: On Solaris 10, after installing or upgrading the VRTSvxfs using install bundle, restart the system.

The tuneable values cannot persist after upgrade from the releases prior to 6.0 [3736830]

If you upgrade releases that are prior to 6.0 to 6.0 or later release, the tunable values which were set by vxtune will be replaced by default values.

Workaround: There is no workaround.

On Solaris 11, when you install the operating system together with SFHA products using Automated Installer, the local installer scripts do not get generated. (3640805)

On Solaris 11, when you use Automated Installer (AI) to install the Solaris 11 operating system together with SFHA products, the local installer scripts fail to get generated.

Workaround:

On the target system(s), execute the following script:

```
/opt/VRTSsfcpi62/bin/run-once
```

Node panics after upgrade from Solaris 11 to Solaris 11.1 on systems running version 6.0.1 or earlier (3560268)

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.

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.

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
```

New agent types are not detected after SFHA stack upgrade to 6.2 [3654406]

On upgrading SFHA stack to 6.2 on Solaris 11, the new types may not be imported. As a result, newly introduced types like SFCache are not seen.

Perform the following steps to verify that you are not impacted by this issue and also for remediation steps.

1. Verify whether SFCache type is present on all running nodes with the command `hatype -list|grep SFCache` command.
2. If SFCache is present in the `hatpye -list` output, then you are not impacted by this issue.
3. If SFCache is not present in the `hatype -list` output, continue with the following steps.
4. Stop VCS on all nodes with `hastop -all -force`.
5. Identify the node which has types.cf or types.cf.previous containing SFCache (`grep -l SFCache /etc/VRTSvcs/conf/config/types.cf /etc/VRTSvcs/conf/config/types.cf.previous`).
6. If SFCache is present in types.cf, start VCS on this node first (`hastart`). Once VCS goes into RUNNING state on this node, you can proceed starting VCS on other nodes.
7. If SFCache is not present in types.cf on any node but in types.cf.previous, then copy types.cf.previous to types.cf on this node. Copy the `.previous` files for other type definition files like OracleTypes.cf, SybaseTypes.cf, Db2udbTypes.cf and so on, which are included in the main.cf. Start VCS on this node. Once VCS goes into RUNNING state on this node, check whether “Frozen =1” attribute got configured for available service groups. If present , unfreeze all service groups using `hagrp -unfeeze <group_name> -persistent` command. You can now proceed starting VCS on other nodes.

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.

Installing VRTSvlic package during live upgrade on Solaris system non-global zones displays error messages [3623525]

While installing VRTSvlic package during live upgrade on Solaris system with non-global zones following error messages are displayed:

```
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: This message can be ignored. The vxlicinst, vxlicrep, vxlictest utilities are present in `/opt/VRTSvlic/sbin/` inside a non-global zone.

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

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.

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.

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

If you install or configure DMP, SF, SFCFSHA, SFRAC or VCS 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.

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.

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.

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.

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.

After Live Upgrade to Solaris 10 Update 10/Update 11, boot from an alternate boot environment fails [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 10/11, boot from an alternate boot environment may fail.

Workaround:

1 Run the `vxlufinish` command. Enter:

```
# vxlufinish
```

2 Manually delete the entries of all the volumes of shared disks that are mounted as CFS in the `/altroot.5.10/etc/vfstab` directory. Enter:

```
rm -rf /altroot.5.10/etc/vfstab
```

3 Restart the system.

On Solaris 11, if a reboot is performed during upgrade from 6.0PR1 to 6.2, the pkg verify VRTSsfmh command results in an error (3624856)

On Solaris 11, if a reboot is performed during upgrade from 6.0PR1 to 6.2, the `pkg verify VRTSsfmh` command results in the following error:

```
pkg verify VRTSsfmh
    PACKAGE
STATUS
    pkg://Symantec/VRTSsfmh
ERROR
            dir: var/opt/VRTSsfmh
                    Group: 'root (0)' should be 'other (1)'
            dir: var/opt/VRTSsfmh/etc
                    Missing: directory does not exist
            dir: var/opt/VRTSsfmh/logs
                    Group: 'root (0)' should be 'other (1)'
            dir: var/opt/VRTSsfmh/tmp
                    Group: 'root (0)' should be 'other (1)'
            file: opt/VRTSsfmh/web/operator/cgi-bin/firedrill.pl
                    Missing: regular file does not exist
```

Workaround:

- Set the "Symantec" publisher repository pointing to `VRTSpkgs.p5p`.

  ```
  # pkg set-publisher -P -g /mnt/release_train/sol/6.2/
  SxRT-6.2-2014-10-01a/dvd1-sol_sparc/sol11_sparc/pkgs/VRTSpkgs.p5p
  Symantec
  ```

- Run the `pkg fix VRTSsfmh` command.

  ```
  # pkg fix VRTSsfmh
  ```

vxlustart failed due to lumount error when performing Live Upgrade to Solaris 10 Update 11 (3035982)

Live Upgrade (LU) to Solaris 10 Update 11 using `vxlustart` fails with following error:

```
# lumount -n dest.7667 /altroot.5.10
ERROR: mount point directory </altroot.5.10> is not empty
ERROR: failed to create mount point </altroot.5.10> for file system
</dev/dsk/c1t1d0s0>
ERROR: cannot mount boot environment by name <dest.7667>
ERROR: vxlustart: Failed: lumount -n dest.7667 /altroot.5.10
```

Workaround: To perform Live Upgrade to Solaris 10 Update 11, use one of the following procedures for your operating system version.

To perform Live Upgrade from Solaris 10 Update 10 to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Use `vxlustart` to upgrade to Solaris 10 Update 11.

To perform Live Upgrade from Solaris 10 Update 9 or below to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Use `vxlustart` to upgrade to Solaris 10 Update 11.

To perform Live Upgrade from Solaris 9 to Solaris 10 Update 11

1. Install the Solaris 10 Update 10 LU packages (SUNWlucfg, SUNWlur, SUNWluu) instead of the Solaris 10 Update 11 LU packages.
2. Install the patch 121430-72. (Do NOT patch to a higher version of 121430, such as 121430-92.)
3. Use `vxlustart` to upgrade to Solaris 10 Update 11.

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.

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

Symantec does not support SFCFSHA, SFHA, SF Oracle RAC, or SF Sybase CE 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.

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

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

This issue will be resolved in a later release.

The Installer fails to unload GAB module while installation of SF packages [3560458]

The Installer succeeds to upgrade SF package from 6.1.1 or 6.0.5 to 6.2.1, but GAB module (for 6.1.1 or 6.0.5) fails to unload and remains in loaded state. The issue is seen with the recent updates of Solaris 11U1 (SRU 8) or Solaris 11U2OS 11U1 (SRU 8). During un-installation of SFCFSHA, SFHA, SF Oracle RAC, SF Sybase CE or VCS packages, unloading of GAB fails.

Workaround: Restart the system. Restarting the system will unload the module successfully.

On Solaris 11 non-default ODM mount options will not be preserved across package upgrade (2745100)

On Solaris 11, before the package upgrade if Oracle Disk Manager (ODM) is mounted with non-default mount options such as nocluster, nosmartsync etc, these mount options will not get preserved after package upgrade.

There is no workaround at this time.

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.

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.

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 Symantec Storage Foundation and High Availability Solutions 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.

Upgrades from previous SF Oracle RAC versions may fail on Solaris systems (3256400)

The `vxio` and `vxdmp` modules may fail to stop on Solaris systems during upgrades from previous SF Oracle RAC versions. As a result, the upgrade fails to complete successfully.

Workaround: If `vxio` and `vxdmp` fail to stop and no other issues are seen during upgrade, continue with the upgrade and restart the system when the product installer prompts. After the reboot, use the installer to start the product again by entering:

```
# /opt/VRTS/install/installsfha62 -start
```

Note: Do not use the response file to upgrade in this situation.

After a locale change restart the vxconfig daemon (2417547, 2116264)

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 the first phase of a rolling upgrade, make sure the CVM is online on all nodes without errors (2595441)

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

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 SF Oracle RAC 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.

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.

Only AppHA can be upgrade when AppHA and SF/DMP are installed on the same machine [3693146]

When you upgrade products on your machine, with 6.2.0 Application HA (AppHA) and 6.2.0 Storage Foundation (SF)/Dynamic Multi-Pathing (DMP) installed on it. CPI can detect the two products correctly, but the installer can only upgrade AppHA and fails to upgrade SF.

Workaround: To solve this issue, perform the following steps:

- If you upgrade from versions earlier than 6.2.0 to 6.2.1, upgrade SF and DMP to 6.2 with product specific script `install<prod>` first.
- If you upgrade from 6.2.0 to 6.2.1, upgrade the remaining product manually.

1. Stop SF/DMP.
2. Get the patches for SF/DMP. To get them, type: `#./installmr -listpatches`.
3. Install the patches for SF/DMP with the OS command.
4. Start SF/DMP.

Symantec Storage Foundation known issues

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

- Veritas Volume Manager known issues

Veritas Volume Manager known issues

`vxassist` fails to create volume with maxsize option on FSS disk group that has disk local to one node and exported for FSS disk group [3736095]

On a disk with connectivity to one node is exported for FSS disk group, creating volume with maxsize option fails with the following message:

```
VxVM vxassist ERROR V-5-1-18301  No volume can be created within the given
constraints
```

Workaround: Provide the size instead of "maxsize" when you create volume on FSS disk group that has disk connectivity on single node:

```
# vxassist -g <diskgroup> make <volume> <size>
```

`vxdisksetup -if` fails on PowerPath disks of sizes 1T to 2T [3752250]

`vxdisksetup -if` fails on PowerPath disks of sizes 1T to 2T with the following message:

```
VxVM vxdisksetup ERROR V-5-2-4006 Disk emcpower48 contains auto:\
none DA record emcpower48s2
```

Workaround:

1. Format the disk to EFI label:

   # format

2. Remove the formatted disk from VxVM control:

   ```
   # vxdisk rm emcpower48s2
   ```

3. Scan the disk again:

   ```
   # vxdisk scandisks
   ```

   The disk should show up as emcpower48, without the s2 suffix.

4. Set up the disk:

   ```
   # vxdisksetup -if emcpower48
   ```

`vxdmpraw` creates raw devices for the whole disk, which causes problems on Oracle ASM 11.2.0.4 [3738639]

Oracle recommends working with partitions and using them for Automatic Storage Management (ASM).

The `vxdmpraw` utility allows and creates raw devices for the whole dmp device and changes permission and ownership of dmp devices to Oracle user or group. This results in Oracle ASM discovering the whole disk as a “CANDIDATE” disk.

This leads to the following issues while creating ASM diskgroup using the whole disk:

- The disk used for ASM diskgroup doesn't show ASM tags on the `vxdisk` list.
- ASM alert log shows below errors:

```
On discovery: Device or resource busy
            On creating diskgroup: failed to update diskgroup resource ora
```

Workaround:

To avoid issues, create partition prior to create raw devices for DMP and remove raw devices of the whole disk before ASM discovery. See the following steps and the example:

1 Create a primary partition on the disk with entire space:

```
# fdisk /dev/vx/dmp/ibm_shark0_10
```

2 Use the `vxdmpraw` command to enable DMP devices:

```
# /etc/vx/bin/vxdmpraw enable oracle dba 765 ibm_shark0_10
```

3 Remove the raw device for the whole disk:

```
# ls -l /dev/vx/dmp |grep ibm_shark0_10$

brwxrw-r-x 1 grid oinstall 201, 368 Mar 6 15:43 ibm_shark0_10

# raw -qa |grep 368

/dev/raw/raw17: bound to major 201, minor 368

# raw /dev/raw/raw17 0 0
```

Note: In this example, for making changes persistent in system reboot, remove the raw17 from `/etc/vx/.vxdmprawdev`.

4 Use ASM diskstring as '/dev/raw/*'(default), and discover all available disks:

```
SQL> select name,path,header_status from v$asm_disk;

NAME PATH HEADER_STATU
----- ------ --------------
/dev/raw/raw18 CANDIDATE
```

5 Create ASM diskgroup with raw device of partition:

```
SQL> create diskgroup DATA10 external redundancy disk '/dev/raw/raw18';
```

If you use ASM_DISKSTRING with /dev/vx/dmp/*, then change the permission and owner for the whole disk to prevent ASM from discovering it.

```
# chmod 600 /dev/vx/dmp/ibm_shark0_10

# chown root:root /dev/vx/dmp/ibm_shark0_10
```

VRAS `verifydata` command fails without cleaning up the snapshots created [3558199]

The `vradmin verifydata` and the `vradmin syncrvg` commands leave behind residues if terminated abnormally. These residues can be snapshot volumes or mount points.

Workaround: Remove the snapshot volumes and unmount the mount points manually.

Root disk encapsulation fails for root volume and swap volume configured on thin LUNs (3538594)

Root disk encapsulation fails if the root disk configuration on a thin LUN includes volumes such as `var`, `usr`, or `home`, in addition to the root volumes and the swap volumes. Root disk encapsulation is not supported in this configuration.

Workaround:

There is no workaround.

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 fails on Solaris 11 during expansion of the LUN from array side. The `vxdisk resize` command does not claim correct LUN size on Solaris 11 during expansion of the LUN from array side. This is because of Oracle bug -19603615. On Solaris 11, the `vxdisk resize` command may exit without any error, returning incorrect LUN size or failing 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, the following workaround works:

After expansion of LUN from array side, run `format -d` command and then run `vxdisk resize` command.

SmartIO VxVM cache invalidated after relayout operation (3492350)

If a relayout operation is done on a volume that has SmartIO VxVM caching enabled, the contents of the cache for the volume may be invalidated.

Workaround:

This behavior is expected. There is no workaround.

Disk greater than 1TB goes into error state [3761474, 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
```

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
```

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
```

VxVM fails to create volume by the vxassist(1M) command with maxsize parameter on Oracle Enterprise Linux 6 Update 5 (OEL6U5) [3736647]

The data change object (DCO) volume can't be created when volume size gets too long with the maxsize parameter, otherwise it succeeds.

When Veritas Volume Manager (VxVM) calculates the maxsize parameter, it also accounts pending reclaimation disks in the maxsize_trans function. If some disks are not yet reclaimed, space from those disks is not available to create volume.

Workaround: To resolve this issue, follow the two steps:

1 #`vxdisk -o thin reclaim <diskgroup>`

2 #`vxassist -g <diskgroup> make vol maxsize <parameters>`

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.

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
```

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.

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
```

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
```

Upgrading from Symantec Storage Foundation and High Availability Solutions 5.x to 6.2.1 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 6.2.1 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.

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.

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.

Workaround:

There is no workaround available.

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
```

Importing a clone disk group fails after splitting pairs (3134882)

When you import a clone disk group with the `-o updateid` option, the GUIDs of all the objects are assigned new values. However, these values are not updated on the maps in the data change object (DCO). When you initiate a volume recovery, it fails on the volumes having instant DCO (version >= 20) because it does not find the objects corresponding to the GUIDs. In this situation, the DCO is considered corrupt and the volume remains inaccessible.

Workaround: You mainly need the `-o updateid` option when you import the clone disk group on the same host as the primary disk group. You can avoid using the option by doing one of the following:

- Import the clone disk group on a different host.
- Deport the primary disk group before you import the clone disk group.

If the import of the clone disk group with `-o updateid` option or the recovery of volume thereafter fails with a message about the DCO being corrupted, this error occurs because the GUIDs are not being updated on the DCO implicitly. If the workaround is not acceptable and you need to access the volume, you can remove the DCO. You can dissociate or remove the snapshots and then remove the DCO manually to let the recovery proceed.

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 cause other commands to run slowly because the VxVM configuration daemon (`vxconfigd`) is busy servicing `vxdisk scandisks`.

Workaround: Stop the `vxattachd` script and set EMC CLARiiON values, as follows:

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
```

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

The Veritas Volume Manager (VxVM) 6.2.1 includes several array names that differ from the array names in releases 5.1SP1 or prior. Therefore, if you upgrade to VxVM 6.2.1 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 6.2.1.

Workaround:

Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-29 shows the Hitachi arrays that have new array names.

Table 1-29 Hitachi arrays with new array names

Previous name	New name
TagmaStore-USP	Hitachi_USP
TagmaStore-NSC	Hitachi_NSC
TagmaStoreUSPV	Hitachi_USP-V
TagmaStoreUSPVM	Hitachi_USP-VM
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 6.2.1. 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
```

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

Multi-pathing 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
```

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`.

Restarting the vxconfigd daemon on the slave node after a disk is removed from all nodes may cause the disk groups to be disabled on the slave node (3591019)

The issue occurs if the storage connectivity of a disk is removed from all the nodes of the cluster and the `vxconfigd` daemon is restarted on the slave node before the disk is detached from the slave. All the disk groups are in the `dgdisabled` state on the slave node, but show as `enabled` on the other nodes.

If the disk was detached before the `vxconfigd` daemon is restarted, the issue does not occur.

In a Flexible Shared Storage (FSS) environment, removing the storage connectivity on a node that contributes DAS storage to a shared disk group results in global connectivity loss because the storage is not connected elsewhere.

Workaround:

To prevent this issue:

Before restarting the `vxconfigd` daemon, if a disk in a shared disk group has lost connectivity to all nodes in the cluster, make sure that the disk is in the `detached` state. If a disk needs to be detached, use the following command:

```
# vxdisk check diskname
```

To resolve the issue after it has occurred:

If `vxconfigd` is restarted before the disks got detached, remove the node from the cluster and rejoin the node to the cluster.

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.

Workaround:

Do one of the following:

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

The 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.

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.

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.

CVM fails to start if the first node joining the cluster has no connectivity to the storage (2787713)

If the first node joining the cluster has no connectivity to disks, the import of shared disk groups fails. Other nodes that join the cluster later assume that the auto-import of disk groups is already done as part of the existing cluster processing.

Workaround:

Perform a master switch to the node that has connectivity to the disks. Then import the disk groups manually.

Dynamic LUN expansion is not supported for EFI disks in simple or sliced format and non-EFI disks greater than 1TB in simple or sliced format.(2836798)

Dynamic LUN expansion is not supported for EFI (Extensible Firmware Interface) disks in simple or sliced format and non-EFI disks greater than 1TB in simple or sliced format. The recommended format is the Cross-platform Data Sharing (CDS) disk format.

Workaround:

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

CVMVolDg agent may fail to deport CVM disk group when CVMDeportOnOffline is set to 1

When CVMDeportOnOffline is set to 1, 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 and the disk group is required to be deported during offline, set the value of the `CVMDeportOnOffline` attribute to 1 for all of the resources.

DMP uses OS device physical path to maintain persistence of path attributes from 6.0 [3761441]

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 any path-level attributes that were defined in the /etc/vx/dmppolicy.info file.

Workaround:

To configure path-level attributes

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

The vxsnap print command shows incorrect value for percentage dirty [2360780]

The `vxsnap print` command can display the percentage of regions that differ between snapshots, shown as the %dirty. In SF 6.0, if this command is run while the volumes are online and being actively used, the shown %dirty may lag from

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

1 TB luns goes in error state with Solaris x86 (3761486, 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, 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.

Limitation to DMP support for ZFS root in the Oracle VM Server for SPARC guest (3221944)

DMP support for ZFS root is not supported in the Oracle VM Server for SPARC guest. If DMP meta devices are exported to Oracle VM Server for SPARC 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 Oracle VM Server for SPARC guest:

DMP supports non-root ZFS in the Oracle VM Server for SPARC guest. You can use DMP devices for non-root ZFS.

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.

System hangs on a boot up after Boot Environment upgrades to Solaris 11 Update 2 and SF 6.2 from Solaris 11 GA.[3628743]

The issue results from some kind of OS race condition causing a deadlock during the system boot after upgrade. This hang sometimes gets resolved after many hours. This is still being investigated further with Oracle support engagement for solution.

Workaround:

The issue can be avoided if you perform the following steps to upgrade to Solaris 11 Update 2 in a specified order:

1 Upgrade system to Solaris 11 update 1.

2 Upgrade SF to 6.2

3 Upgrade system to Solaris 11 update 2.

vxdisk export operation fails if length of hostprefix and device name exceeds 30 characters (3543668)

If the combined length of the hostprefix and the device name exceeds 30 characters, the vxdisk export operation fails with the following error message:

```
VxVM vxdisk ERROR V-5-1-18318 Device c6t50060E8005655501d86s2: Name too
long for export. Length of Hostprefix + Disk accessname should not exceed
30 characters. Please see vxdctl(1M) man page for information on setting
user-specifed hostprefix.
```

Workaround:

Use the enclosure-based naming (EBN) scheme instead of the operating system naming (OSN) scheme. OSN naming typically contains more characters and is not as intuitive. If the EBN name combined with the hostprefix exceeds 30 characters,

you can manually set the hostprefix to a smaller size using the `vxdctl set hostprefix=value` command, where *value* is the new hostprefix.

Virtualization known issues

Locale message displayed on Solaris 11 system for solaris10 brand zones

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.

Veritas File System known issues

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

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.

`vxresize` may fail when you shrink a file system with the "blocks are currently in use" error (3762935)

The vxresize shrink operation may fail when active I/Os are in progress on the file system which 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: Re-run the shrink operation after stopping the I/Os.

On Solaris11U2, /dev/odm may show 'Device busy' status when the system mounts ODM [3661567]

If the system tries to mount Oracle Disk Manager (ODM) in a mode which is not supported by the installed license, the later ODM mount may not succeed and shows /dev/odm device busy error.

Workaround: There are two ways to resolve it.

Remove /dev/odm mount point and recreate it. Or reboot the system and then mount /dev/odm.

Delayed allocation may be turned off automatically when one of the volumes in a multi-volume file system nears 100%(2438368)

Delayed allocation may be turned off automatically when one of the volumes in a multi-volume file system is in almost full usage, even if other volumes in the file system have free space.

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

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.

Enabling delayed allocation on a small file system may disable the file system (2389318)

When you enable the delayed allocation on a small file system, such as around 100 MB, the file system may be disabled. In this case, the following error message is 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 for turning off the delayed allocation for the file system.

Oracle Disk Manager (ODM) may fail to start after upgrade from 6.2 to 6.2.1 on Solaris 11 [3739102]

ODM may fail to start after upgrade from 6.2 to 6.2.1 on Solaris 11.

Workaround: Manually start ODM service by entering:

```
# /lib/svc/method/odm  start
```

On the cluster file system, clone dispose may fail [3754906]

In case of clone dispose on cluster, if Veritas File System (VxFS) fails to unlink clones, the specific fset is marked as bad incore and the fullfsck flag is marked on the file system.

Workaround: Run full fsck on the file system, and it will complete the extop processing required for the clone removal.

VRTSvxfs verification reports error after upgrading to 6.2.1 [3463479]

Upgraded to 6.2.1, the VRTSvxfs package cannot pass the verification check with the pkg verify VRTSvxfs command. You can see error messages similar to the following:

```
# pkg verify VRTSvxfs
PACKAGE                                         STATUS
pkg://Symantec/VRTSvxfs                         ERROR
        driver: vxfs
                etc/name_to_major: 'vxfs' entry not present
```

Workaround: Use the following command to fix this issue:

```
# pkg fix VRTSvxfs
```

spfile created on VxFS and ODM may contain uninitialized blocks at the end (3760262)

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:
```

In a CFS cluster, that has multi-volume file system of a small size, the fsadm operation may hang (3348520)

In a CFS cluster, that has multi-volume file system of a small size, the fsadm operation may hang, when the free space in the file system is low.

Workaround: There is no workaround for this issue.

On the online cache device you should not perform the `mkfs` operation, because any subsequent `fscache` operation panics (3643800)

When the `mkfs` operation is performed on a volume already in use for SmartIO, caching can lead to unexpected results when the subsequent `sfcache` operations are performed.

Workaround: Workaround is not available.

When hard links are present in the file system, the sfcache list command shows incorrect cache usage statistics (3059125, 3760172)

If a hard link is present for a file that is loaded in the cache, the `sfcache list` command shows the cache usage for both files: the original file and the hard link. The resulting statistics are incorrect, because the cache usage is shown to be twice the actual usage.

For example:

```
# sfcache list -r /mnt1
/mnt1:
CACHE-USED(MB) MODE PINNED NAME
0 read no /mnt1/test_10
0 read no /mnt1/test_20
0 read no /mnt1/test_50
0 read no /mnt1/test_100
0 read no /mnt1/test_200
0 read no /mnt1/test_300
0 read no /mnt1/test_400
500 read yes /mnt1/test_500
0 read no /mnt1/test_1024
500 read yes /mnt1/dir/hardlink
500 read no /mnt1/dir
1000 read no /mnt1
```

```
# sfcache list fs1

Cachearea: fs1
Assoc Type: AUTO
Type: VxFS
Size: 1.00g
State: ONLINE
/dev/vx/dsk/sfcache_defaultdg/fs1:
FSUUID CACHE-USED(MB) MODE MOUNTPOINT
23642651-81a5-0d00-1a26-0000911ec26c 1000 read /mnt1
```

Workaround: There is no workaround for this issue.

Deduplication can fail with error 110 (3741016)

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.

A restored volume snapshot may be inconsistent with the data in the SmartIO VxFS cache (3760219)

The data in a volume snapshot may have data that is inconsistent with the VxFS level SmartIO cache. When the volume snapshot is restored and mounted, then before using that file system you should purge the corresponding cache data. Or, disable the caching for that file system.

Workaround:

Purge the file system data from the SmartIO cache after restoring the volume snapshot.

```
# sfcache purge {mount_point|fsuuid}
```

When in-place and relocate compression rules are in the same policy file, file relocation is unpredictable (3760242)

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.

Replication known issues

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

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.

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 [3761497]

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
```

Transactions on VVR secondary nodes may timeout waiting for I/O drain [3236772]

If the VVR secondary node receives updates out of order from the Primary, and a transaction starts on the secondary site, then the transaction may timeout waiting for I/O drain. This issue may occur in situations where the gaps created by out of order updates are not filled within the transaction timeout period.

Workaround:

Pause replication and make configuration changes.

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

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]

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
```

SRL resize followed by a CVM slave node join causes the RLINK to detach (3259732)

In a CVR environment, performing a CVM slave node join after an SRL resize may stop replication due to a detached RLINK.

Workaround:

There is no workaround for this issue.

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 6.2.1 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 [3343141]

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, 3724338)

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:

There is no workaround for this issue.

After performing a CVM master switch on the secondary node, both rlinks detach (3642855)

If the VVR logowner (master) node on the secondary site goes down during initial synchronization, then during the RVG recovery (initiated on any secondary side node as a result of node crash), the replication links detach with the following error:

```
WARNING: VxVM VVR vxio V-5-0-187 Incorrect magic number or unexpected
upid (1) rvg rvg1
WARNING: VxVM VVR vxio V-5-0-287 rvg rvg1, SRL srl1: Inconsistent log
 - detaching all rlinks.
```

Workaround:

Restart replication using the autosync operation.

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 (3761555, 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
```

Symantec Storage Foundation and High Availability known issues

Cache area is lost after a disk failure (3158482)

SmartIO supports one VxFS cache area and one VxVM cache area. If you create one cache area, and the disk fails, the cache area becomes disabled. If you attempt to create a second cache area of the other type before the cache disk group is enabled, then the first cache area is lost. It cannot be brought online.

For example, first you created a VxFS cache area. The disk failed and the cache area is disabled. Now create the VxVM cache area. While creating VxVM cache area, SmartIO looks for an existing default cache area. Due to the failed disk, the

existing cache area cannot be found. So SmartIO creates a VxVM cache area with the same name. Now even if disk containing VxFS cache area comes up, SmartIO cannot access the original cache area. In this scenario, the VxFS cache area is lost. Losing the cache area in this case does not result into any data loss or data inconsistency issues.

Workaround:

Create a new VxFS cache area.

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.

Some SmartTier for Oracle commands do not work correctly in non-POSIX locales (2138030)

Some SmartTier for Oracle 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 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 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 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.

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

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

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 (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.

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.

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

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

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.

System encounters multiple VCS resource timeouts and agent core dumps [3424429]

The system encounters multiple VCS resource timeouts and agent core dumps without any specific reason.

The issue pertains to a hardware errata with the Intel Xeon CPUs where a processor can go into a low power sleep mode, but takes a long time to wake up. This can cause erratic scheduling behavior, leading to unexpected delays, expired timers, or occasional freezes. For more information, see the Oracle document: https://support.oracle.com/epmos/faces/BugDisplay?id=15659645

Workaround: Add the following lines to the `/etc/system` file and reboot the system:

```
set idle_cpu_prefer_mwait = 0
set idle_cpu_no_deep_c = 1
```

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.

Issues related to the VCS engine

This section describes the knonw issues about the VCS engine.

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.

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

The `ha` commands fail to work for one of the following reasons:

- If you first use a non-root user without a home directory and then create a home directory for the same user.
- If you configure security on a cluster and then un-configure and reconfigure it.

Workaround

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

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

Running `-delete -keys` 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.

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

SFHA Solutions 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 SFHA Solutions 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 SFHA Solutions on the node and then on the rest of the nodes in the cluster.

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 [3380586]

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.

VCS crashes if feature tracking file is corrupt [3603291]

VCS keeps a track of some specific features used in the VCS cluster. For example, if a Global service group is brought online then the feature is logged in a specific feature tracking file. If the file however is corrupt, then VCS may dump core when attempting to write data to the file.

Workaround: Delete the corrupt feature tracking file (`/var/vx/vftrk/vcs`) and restart VCS.

RemoteGroup agent on versions lower than 6.2 reports service group status as UNKNOWN [3638347]

When the RemoteGroup agent running on a VCS version lower than 6.2 tries to monitor a service group on a 6.2 cluster, it reports the service group status as UNKNOWN.

Workaround: No workaround.

RemoteGroup agent and non-root users may fail to authenticate after a secure upgrade [3649457]

On upgrading a secure cluster to 6.2, the following issues may occur with unable to open a secure connection error:

- The RemoteGroup agent may fail to authenticate with remote cluster.
- Non-root users may fail to authenticate.

Workaround

1. Set `LC_ALL=C` on all nodes before upgrade or perform the following steps after the upgrade on all nodes of the cluster:
   - Stop HAD.
   - Set `LC_ALL=C`.
   - Start HAD using `hastart`.
2. Reset LC_ALL attribute to the previous value once the non-root users are validated.

Issues related to the bundled agents

This section describes the known issues of the bundled agents.

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

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.

The workaround for this is to 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

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.

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

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 SFHA Solutions 6.0 or later versions.

When you upgrade SFHA Solutions 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 SFHA Solutions 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 SFHA Solutions 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 SFHA Solutions 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>
```

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.

IP Agent fails to detect the online state for the resource in an exclusive-IP zone [3592683]

IP Agent does not detect the online state for the resource inside an exclusive-IP zone monitoring an IPv6 address if the link-local address is down.

Workaround: Bring the link-local address of the device up for the IP agent to detect the IPv6 address state properly.

SFCache Agent fails to enable caching if cache area is offline [3644424]

SFCache agent cannot enable caching if cache area associate with this particular object is in offline state. User need to manually online the cache area to make sure that caching can be enabled/disabled.

Workaround: Online the cache area using `sfcache` command

```
# sfcache online <cache_area_name>
```

RemoteGroup agent may stop working on upgrading the remote cluster in secure mode [3648886]

RemoteGroup agent may report the resource state as UNKNOWN if the remote cluster is upgraded to VCS 6.2 in secure mode.

Workaround: Restart the RemoteGroup agent.

Issues related to the VCS database agents

This section describes the known issues about VCS database agents.

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

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

Workaround:

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

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

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

NOFAILOVER action specified for certain Oracle errors

The 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.

Oracle agent fails to offline pluggable database (PDB) resource with PDB in backup mode [3592142]

If the PDB is in backup mode and if you attempt to offline the corresponding PDB resource, this will cause PDB resource to go into “Unable to Offline” state.

Workaround: Manually remove the PDB from the backup mode before attempting to take the PDB resource offline.

Clean succeeds for PDB even as PDB staus is UNABLE to OFFLINE [3609351]

Oracle does not allow any operation on a PDB when the PDB is in backup mode. This is an expected behavior of Oracle. Therefore, a shutdown fails when it is initiated on a PDB in backup mode and returns an UNABLE TO OFFLINE status for the PDB. If PDB is removed from the backup mode using the SQL script, the agent framework is unable to change the UNABLE TO OFFLINE status of the PDB as clean is called. Since Oracle does not differntiate between clean and offline for PDB, clean succeeds for the PDB in spite of being in UNABLE TO OFFLINE state.

Workaround: No workaround.

Second level monitoring fails if user and table names are identical [3594962]

If the table inside CDB has same name as the user name, second level monitoring fails and Oracle agent fails to update the table. For example, if user name is `c##pdbuser1` and table is created as `c##pdbuser1.vcs`, then Oracle agent is unable to update it.

Workaround: Avoid having identical user and CDB table names.

Monitor entry point times out for Oracle PDB resources when CDB is moved to suspended state in Oracle 12.1.0.2 [3643582]

In Oracle-12.1.0.2.0, when CDB is in SUSPENDED mode, then the SQL command for PDB view (`v$pdbs`) hangs. Due to this, the monitor entry point in PDB gets timed out and there is no issue found in oracle-12.1.0.1.0 .

Workaround: No workaround.

Oracle agent fails to online and monitor Oracle instance if threaded_execution parameter is set to true [3644425]

In Oracle 12c, the threaded execution feature is enabled. The multithreaded Oracle Database model enables Oracle processes to execute as operating system threads in separate address spaces. If Oracle Database 12c is installed, the database runs in the process mode. If you set a parameter to run the database in threaded mode, some background processes on UNIX and Linux run with each process containing one thread, whereas the remaining Oracle processes run as threads within the processes.

When you enable this parameter, Oracle agent is unable to check smon (mandatory process check) and lgwr (optional process check) processes which where tradtionally used for monitoring and which now run as threads.

Workaround: Disable the threaded execution feature as it is no supported on Oracle 12C.

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 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.

Logs from the script executed other than the agent entry point goes into the engine logs [3547329]

The agent logs of C-based and script-based entry points get logged in the agent log when the attribute value of LogViaHalog is set to 1 (one). To restore to the older logging behavior in which C-based entry point logs were logged in agent logs and script-based entry point logs were logged in engine logs, you can set the LogViaHalog value as 0 (zero). However, it is observed that some C-based entry point logs continue to appear in the engine logs even when LogViaHalog is set to 1 (one). This issue is observed on all the database agents.

Workaround: No workaround.

Issues related to Intelligent Monitoring Framework (IMF)

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

Registration error while creating a Firedrill setup [2564350]

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

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

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

Workaround: No workaround.

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 `linkamf` displays syntax error [2858163]

Bash cannot interpret Perl when executed directly.

Workaround: Run `linkamf` as follows:

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

Error messages displayed during reboot cycles [2847950]

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

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

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

Workaround: No workaround.

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

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

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

Workaround: No Workaround.

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 shows 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.

ProPCV fails to prevent a script from running if it is run with relative path [3617014]

If the absolute path is registered with AMF for prevention and the script is run with the relative path, AMF fails to prevent the script from running.

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 the Cluster Manager (Java Console)

This section describes the known issues about Symantec Cluster Server 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.

VCS Cluster Configuration wizard issues

IPv6 verification fails while configuring generic application using VCS Cluster Configuration wizard [3614680]

The VCS Cluster Configuration wizard fails to check whether IPv6 IP is already plumbed while configuring a generic application through the Virtual IP page. The wizard does neither displays a warning if IPv6 IP is already plumbed elsewhere nor indicates whether it is reachable through a ping.

Workaround: Manually ensure that IPv6 is not plumbed elsewhere on the network before configuring the generic application through the wizard.

Browser shows 404 error and wizard fails to launch when VCS is installed with Jumpstart or upgraded with Live upgrade [3626253]

On Solaris 10 systems, when ApplicationHA or VCS is installed through Jumpstart or Live upgrade mechanism, the wizards cannot be launched. The browser displays the 404 – page not found error because VCS namespace values are not set in the `xprtld` configuration.

Workaround:

1. Boot the system to the newly created boot environment.
2. Ensure xprtld service is in online state

   ```
   # svcs /system/xprtld
   ```

3. Run the following commands:

   For VCS:

   ```
   # /opt/VRTSvcs/portal/admin/conf/configGen.pl
   ```

   For ApplicationHA

   ```
   # /opt/VRTSvcs/portal/admin/plugins/unix/conf/configGen.pl
   ```

Symantec Dynamic Multi-Pathing known issues

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

VxDMP console installation fails if installed in a folder having Japanese character in it (2670506)

When you install VxDMP on a host with a Japanese OS, if you select for installation a directory having a Japanese character in its name, the installation fails.

Workaround:

When installing VxDMP on a host with a Japanese OS, use the default installation directory:`C:\program files(x86)\symantec\DMP\`

I/O statistics are not available if overlapping filters are applied in the VxDMP tab of the vSphere Web Client (3306319)

The `Manage > VxDMP` tab of the vSphere Web Client does not support mixing filters for I/O statistics. In the Connctivity Map, the I/O statistics pie chart does not display if you specify an unsupported filter combination, such as specifying both a controller name and a DMP node name.

The VxDMP tab supports the following filters for I/O statistics:

- [vm=*virtual-machine-name*] ctlr=*ctlr-name*
- [vm=*virtual-machine-name*] dmpnodename=*dmp-device-name*

- [vm=*virtual-machine-name*] enclosure=*enclr-name* [portid=*array-portid*] [ctlr=*ctlr-name*]
- [vm=*virtual-machine-name*] pathname=*path-name*
- [vm=*virtual-machine-name*] pwwn=*array-port-wwn* [ctlr=*ctlr-name*]

Workaround:

There is no workaround.

Drop-down list for filter on LUNs tab in the Web client fails to load (3189918)

When you click on the drop-down list for filter on the LUNs tab in the Web client, the drop-down list fails to load.

Configuration of the VxDMP Console server and registering the DMP plugin fails if the installation directory contains special characters (2671487)

If you install DMP Console in a directory that has a name with unsupported special characters, then the post-installation configuration fails. The special characters that are not supported include the following: percent sign (%); ampersand (&); tilde (~); apostrophe ('); exclamation point (!); at sign (@); number sign (#); dollar sign ($); carat (^); plus sign (+); equal sign (=}; brackets ([ ]) and semicolon (;).

If the directory contains unsupported characters, you cannot configure Symantec DMP Console server, or register the DMP plug-in.

If you try to uninstall VxDMP from the system, then a fatal error occurs.

Workaround:

The installation directory name must not have special characters other than dash (-); underscore (_); parenthesis ((()); or period (.).

Issues related to installation (2642164)

Repair of the DMP console installation fails if the *InstallDirectory*/DMP/bin directory is missing, or if files are missing from this directory.

Workaround:

If the repair fails, you must uninstall and reinstall the package.

On ESXi 5.0, DMP may fail to fetch any data (2733610)

Small Footprint CIM Broker (SFCB) on ESXi 5.0 is incompatible with JRE 1.6 u29 or later. As a result, the VxDMP tab in the vSphere Client UI fails to fetch any data. The remote CLI displays the Operation execution failed message as shown in the following example:

```
C:\Program Files (x86)\Symantec\DMP\bin> vxdmpadm vcenter=10.209.62.226
vc_user=zen\administrator vc_passwd=Merlin@123 host=10.217.56.157 gettune
VxVM vxdmpcmd ERROR V-5-1-18620 Operation execution failed.
```

To avoid this issue, upgrade to ESXi 5.0 Update 1 or higher.

On ESXi 5.0, run the following on every reboot of ESX server:

```
# /etc/init.d/sfcbd-watchdog restart
```

Additionally, perform the following workaround after installing DMP:

Workaround for vSphere Client UI:

1. Add **sblim.wbem.httpPoolSize=0** in `cimclinet.properties` file, which is located on the Console server:

   ```
   %AllUsersProfile%\Symantec\DMP\Conf\
   ```

2. Restart the VxDMP console service.

Workaround for CLI:

1. Create a new file and name it as **sblim-cim-client2.properties**.
2. In the file, add **sblim.wbem.httpPoolSize=0**, and copy the file to the following location:

   ```
   VxDMP Console install directory\DMP\bin\
   ```

Workaround for Web Client: There is no workaround for vSphere Web Client.

Downloading the ESXi bundle for an ESXi server requires that the Internet Explorer Enhanced Security is off (2843571)

To be able to download the ESXi offline bundle from the home view, make sure that for the logged in user, the Internet Explorer Enhanced Security setting is set to "Off".

VxDMP fails to detect the virtual machines with special characters (3249544)

In the Web Client, VxDMP fails to detect the virtual machines that contain special characters in their names. The special characters that are not supported include the following: bracket ([); slash (\); carat (^); dollar sign ($); period (.); pipe (|); question mark (?); asterisk (*); plus sign (+); parenthesis ((()); and braces ({ }).

Workaround: Rename the virtual machine, such that it does not contain a special character.

The Connectivity Map of the VxDMP Web Client is blank (3581542)

When you select **Manage > VxDMP** tab, the VxDMP Web Client does not display the Connectivity Map.

Workaround:

Select the refresh operation on your browser (F5) to display the Connectivity Map.

Devices for newly added claim rules do not show in the list of storage arrays (3615915)

When you add a new claim rule, the devices claimed by that rule do not show in the list of storage arrays in the VxDMP tab of the vSphere Web Client.

Workaround:

Log in again to the vSphere Web Client.

The connection lines in the Connectivity Map fail to appear as expected when you resize the Web Client view (3356063)

In the VxDMP Web Client view, with multiple tabs open if you unpin the panels to resize the view or resize the browser, then the connection lines in the Connectivity Map do not appear as expected. You may observe this issue in the Connectivity Map tab after resizing the view in a different tab.

Workaround: To fix this issue, in the **Connectivity Map** view, click **Reset**.

The VxDMP Web Client does not display LUNs with names longer than 51 characters (3306263)

The VxDMP Web Client does not display LUNs with names longer than 51 characters.

Workaround:

Rename the LUN so that the name has 51 characters or less.

Adding a device to a SmartPool displays the progress as Loading (3570852)

After you add a device to the SmartPool, the progress bar displays Loading when you view the performance data (**ESX** > **Monitor** > **Performance** > **SmartPool**).

The issue is fixed by VMware in ESX 6.0.

Workaround:

Use the `F5` key to refresh the view.

DMP shows errors if the ESXi server is disconnected (3405262)

DMP shows errors if the ESXi server is disconnected. This issue is seen with both the DMP Web Client and the DMP vSphere Client. If you click the Settings tab, DMP displays an Action script error.

There is no workaround.

The DMP CLI for remote management of an ESXi 5.0 host requires configuration (3389530)

By default, the DMP CLI is not configured for remote management of an ESXi 5.0 host. The configuration step is not required for hosts with ESXi 5.0 update 1 or later.

Workaround:

Upgrade to ESXi 5.0 update 1

OR

Use the following procedure to configure the DMP CLI.

To configure the DMP CLI for remote management of an ESXi 5.0 host

1. Stop the console service.
2. Add the following line to the file `C:\ProgramData\Symantec\DMP\conf\cimclient.properties`:

   ```
   sblim.wbem.httpPoolSize=0
   ```

3. Restart the console service.

The vxcache device assigned from SmartPool to a Virtual Machine shows with different names (3556478)

The vxcache device assigned from SmartPool to a Virtual Machine shows with different names like vxcache0_1, vxcache0_2, vxcache0_3.

The vxcache device LUN serial number is created using the virtual machine UUID. Therefore, if the virtual machine UUID changes with operations like cloning, then the LUN serial number of vxcache device also changes. This causes the vxcache device name to change.

Workaround:

This behavior does not have any functional impact.

To display the original default device name, clear the vxcache device name with the following command:

```
# vxddladm -c assign names
```

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).

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

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

Workaround

To resolve this issue

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

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

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
  ```

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.

Panic due to null pointer de-reference in vx_bmap_lookup() (3038285)

If you use the `fsadm -b` command on a CFS secondary node to resize the file system, it might fail with the following error message printed in the syslog:

```
Reorg of inode with shared extent larger than 32768 blocks
can be done only on the CFS Primary node
```

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

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 known issues

This section lists the known issues in Oracle RAC.

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

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 SF Oracle RAC 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*.

SF Oracle RAC issues

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

ASM disk groups configured with normal or high redundancy are dismounted if the CVM master panics due to network failure in FSS environment or if CVM I/O shipping is enabled (3600155)

Disk-level remote write operations are paused during reconfiguration for longer than the default ASM heartbeat I/O wait time in the following scenarios:

- CVM master node panics
- Private network failure

As a result, the ASM disk groups get dismounted.

Workaround: See to the Oracle metalink document: 1581684.1

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|FaultOnMonitorTimeouts"
CSSD  AlertOnMonitorTimeouts 4
CSSD  FaultOnMonitorTimeouts 0
```

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

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

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.

The vxconfigd daemon fails to start after machine reboot (3566713)

The `shutdown -r` command makes sure that the file contents on the OS file system are written properly to the disk before a reboot. The `volboot` file is created in the OS file system, and is used to bring up the `vxconfigd` daemon after the system reboot. If the machine reboots for any reason without proper shutdown, and the `volboot` file contents are not flushed to the disk, `vxconfigd` will not start after the system reboots.

Workaround:

You must rerun the `vxinstall` script to re-create the `volboot` file and to start the `vxconfigd` daemon and other daemons.

Health check monitoring fails with policy-managed databases (3609349)

The health check option of the Symantec Cluster Server agent for Oracle fails to determine the status of the Oracle resource in policy-managed database environments. This is because the database SID is dynamically created during the time of the health check as a result of which the correct SID is not available to retrieve the resource status.

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

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

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

- On galaxy: 192.168.12.1
- On nebula: 192.168.12.2

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

CVMVolDg agent may fail to deport CVM disk group

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

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

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.

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.

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.

CVM requires the T10 vendor provided ID to be unique (3191807)

For CVM to work, each physical disk should generate a unique identifier (UDID). The generation is based on the T10 vendor provided ID on SCSI-3 vendor product descriptor (VPD) page 0x83. In some cases, the T10 vendor provided ID on SCSI-3 VPD page 0x83 is the same for multiple devices, which violates the SCSI standards. CVM configurations should avoid using such disks.

You can identify the T10 vendor provided ID using the following command:

```
# sq_inq --page=0x83 /dev/diskname
```

On VxVM you can identify the T10 vendor provided ID using the following command:

```
# /etc/vx/diag.d/vxscsiinq -e 1 -p 0x83 /dev/vx/rdmp/diskname
```

You can verify the VxVM generated UDID on the disk using the following command:

```
# vxdisk list diskname | grep udid
```

Preserving Flexible Storage Sharing attributes with vxassist grow and vxresize commands is not supported (3225318)

Preservation of FSS attributes using `vxasssist grow` and `vxresize` is not supported. FSS attributes include the attributes that are specified on the command line as well as the attributes implicitly assumed for FSS disk groups. These attributes are not reused with further `vxassist` orpations on the volume such as the `growby` and the `vxresize` commands.

Workaround:

There is no workaround for this issue.

FSS Disk group creation with 510 exported disks from master fails with Transaction locks timed out error (3311250)

Flexible Storage Sharing (FSS) Disk group creation for local disks that are exported may fail if the number of disks used for disk group creation is greater than 150, with the following error message:

```
VxVM vxdg ERROR V-5-1-585 Disk group test_dg: cannot create: Transaction
    locks timed out
```

A similar error can be seen while adding more that 150 locally exported disks (with `vxdg adddisk`) to the FSS disk group, with the following error message:

```
VxVM vxdg ERROR V-5-1-10127 associating disk-media emc0_0839 with emc0_0839:
        Transaction locks timed out
```

Workaround:

Create an FSS disk group using 150 or less locally exported disks and then do an incremental disk addition to the disk group with 150 or less locally exported disks at a time.

vxdisk export operation fails if length of hostprefix and device name exceeds 30 characters (3543668)

If the combined length of the hostprefix and the device name exceeds 30 characters, the vxdisk export operation fails with the following error message:

```
VxVM vxdisk ERROR V-5-1-18318 Device c6t50060E8005655501d86s2: Name too
long for export. Length of Hostprefix + Disk accessname should not exceed
30 characters. Please see vxdctl(1M) man page for information on setting
user-specifed hostprefix.
```

Workaround:

Use the enclosure-based naming (EBN) scheme instead of the operating system naming (OSN) scheme. OSN naming typically contains more characters and is not as intuitive. If the EBN name combined with the hostprefix exceeds 30 characters, you can manually set the hostprefix to a smaller size using the `vxdctl set hostprefix=`*`value`* command, where *value* is the new hostprefix.

Change in naming scheme is not reflected on nodes in an FSS environment (3589272)

In a Flexible Storage Sharing (FSS) environment, if you change the naming scheme on a node that has local disks, the remote disk names are not reflected with the corresponding name change. If you change the naming scheme on a node where exported disks are present, to reflect the updated remote disk names, you must either export the disks again or restart the node where the remote disks are present

Workaround:

There is no workaround for this issue.

vxassist does not create data change logs on all mirrored disks, if an FSS volume is created using DM lists (3559362)

When a Flexible Storage Sharing (FSS) volume is created using DM lists, the `vxassist` command does not create data change logs on all the mirrored disks; the number of DCO mirrors is not equal to the number of data mirrors. The `vxassist` command creates a two-way DCO volume.

Workaround:

Manually add a DCO mirror using the `vxassist -g` *`diskgroup`* `mirror` *`dco_volume`* command.

Symantec Storage Foundation for Databases (SFDB) tools known issues

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

Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.2.1 (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 6.2.1.

When upgrading from SF, SFCFSHA, SFHA or SFRAC version 5.0 to SF, SFCFSHA, SFHA or SFRAC 6.2.1, 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.

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.

SFDB commands do not work in IPV6 environment (2619958)

In IPV6 environment, SFDB commands do not work for SF, SFCFSHA, SFHA or SFRAC.

Workaround:

There is no workaround at this point of time.

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.

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 when performed on multiple snapshot configurations (3066532)

When you perform a Reverse Resync operation on multiple snapshot configurations, SFDB reports the following error message:

```
$ 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 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.

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.

Workaround: 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 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.

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.

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.

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

Workaround: To use FlashSnap, modify the above configuration to *.log_archive_dest_1='location=/tpcc_arch'. For example,

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

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.

On Solaris 11.1 SPARC, setting up the user-authentication process using the sfae_auth_op command fails with an error message (3556996)

The debug logs display the missing `ps` utility as the 'ucb' package was absent in the default operating system installation. Due to which, the user-authentication process fails and the following error message is reported:

```
#/opt/VRTS/bin/sfae_auth_op -o setup
Setting up AT
Starting SFAE AT broker

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

Workaround: Install the `pkg:/compatibility/ucb` package such that the `ps` utility is available in `/usr/ucb/ps`.

In the cloned database, the seed PDB remains in the mounted state (3599920)

In Oracle database version 12.1.0.2, when a container database (CDB) is cloned, the **PDB$SEED** pluggable database (PDB) remains in the mounted state. This behavior is observed because of the missing datafiles in the cloned database for all point-in-time copies.

When you attempt to open the cloned seed database, the following error is reported:

```
"ORA-01173" oracle error.
...
SFDB vxsfadm ERROR V-81-0564 Oracle returned error.

Reason: ORA-01122: database file 15 failed verification check
ORA-01110: data file 15: '/tmp/test1/data/sfaedb/newtbs1.dbf'
ORA-01202: wrong incarnation of this file - wrong creation time
...
```

Workaround: There is no workaround for this issue.

Cloning of a container database may fail after a reverse resync commit operation is performed (3509778)

After a reverse resync operation is performed, the cloning of a container database may fail with the following error message:

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

```
Reason: ORA-01503: CREATE CONTROLFILE failed
ORA-01189: file is from a different RESETLOGS than previous files
ORA-01110: data file 6: '/tmp/testRecoverdb/data/sfaedb/users01.dbf'
```

Workaround: There is no workaround for this issue.

If one of the PDBs is in the read-write restricted state, then cloning of a CDB fails (3516634)

Cloning a container database (CDB) for point-in-time copies fails if some of the pluggable databases (PDBs) are open in the restricted mode. The failure occurs with the following error message:

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

Reason: ORA-65106: Pluggable database #3 (PDB1) is in an invalid state.
```

Workaround: There is no workaround for this issue.

Cloning of a CDB fails for point-in-time copies when one of the PDBs is in the read-only mode (3513432)

For Oracle version 12.1.0.1, cloning a container database (CDB) fails if one of the pluggable databases (PDBs) is in the read-only mode. The failure occurs with the following error message:

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

Reason: ORA-00376: file 9 cannot be read at this time
ORA-01111: name for data file 9 is unknown - rename to correct file
ORA-01110: data file 9: '/ora_base/db_home/dbs/MISSING00009'...
```

Workaround: There is no workaround for this issue.

If a CDB has a tablespace in the read-only mode, then the cloning fails (3512370)

For Oracle version 12.1.0.1, when a container database (CDB) has a tablespace in the read-only mode for all point-in-time copies, cloning of that CDB fails with the following error message:

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

Reason: ORA-01122: database file 15 failed verification check
ORA-01110: data file 15: '/tmp/test1/data/sfaedb/newtbs1.dbf'
```

```
ORA-01202: wrong incarnation of this file - wrong creation time
...
```

Workaround: There is no workaround for this issue.

If any SFDB installation with authentication setup is upgraded to 6.2, the commands fail with an error (3644030)

The commands fail with the error message similar to the following:

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

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

Action: Verify that the prodhost is reachable. If it is, verify
that
the vxdbd daemon is enabled and running using the [
/opt/VRTS/bin/sfae_config
status ] command, and enable/start vxdbd using the [
/opt/VRTS/bin/sfae_config
enable ] command if it is not enabled/running. Also make sure you are
authorized to run SFAE commands if running in secure mode.
```

Workaround: Set up the authentication for SFDB again. See *Storage and Availability Management for Oracle Databases* or *Storage and Availability Management for DB2 Databases*.

Symantec Storage Foundation for Sybase ASE CE known issues

This section lists the known issues in SF Sybase CE for this release.

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

SFHA Solutions 6.2.1 doesn't support VCS 6.2.1 on VCSVM and AppHA 6.0 on AppVM running at the same time

SFHA Solutions 6.2.1 doesn't support VCS 6.2.1 on VCSVM and AppHA 6.0 on AppVM running at the same time(3746362)

Workaround: If you plan to upgrade VCSVM to VCS621, perform a full upgrade for both VCSVM and AppVM.

App.RestartAttempts setting does not take effect if value is set to 2 or more

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. (2508392)

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 (for example, a graceful reboot of the guest domain).

Compatibility with other clustering products

Symantec Storage Foundation and High Availability Solutions runs on Symantec Cluster Server (VCS). The version of VCS used by SFHA Solutions is a customized version of VCS. Many components have been removed in order to provide a lighter

footprint inside the virtual machine. You cannot run both SFHA Solutions and VCS together inside the same virtual machine. There is no method to upgrade from SFHA Solutions to VCS.

Additionally SFHA Solutions does not co-exist with other clustering solutions offered by Symantec. These include, Symantec Storage Foundation High Availability, Clustered File System, Clustered File System High Availability and Clustered Volume Manager.

Application monitoring configuration freezes

This issue occurs if you configure application monitoring on systems where host names start with a hyphen. (2038685)

The application monitoring configuration may freeze and the SFHA Solutions view in the vSphere Client may not display the status of the application. If the configured application fails, SFHA Solutions takes no action.

Symantec recommends that you rename systems whose host names start with a hyphen before installing SFHA Solutions and configuring application monitoring on those systems.

Symantec Storage Foundation and High Availability Solutions commands do not display the time as per the locale settings

This issue occurs with all the SFHA Solutions commands that display the date and time stamp in the output. The date and time stamp do not display as per the locale settings on the system. They are displayed only in English. (2142740)

ApplicationHA fails to work if Veritas Operations Manager is uninstalled

The Managed Host components of Veritas Operations Manager (VOM) are installed on the control domain and the guest domain, during the ApplicationHA installation. (2361128, 2323516)

Uninstallation of VOM removes the VRTSsfmh package which breaks the ApplicationHA functionality. The VRTSsfmh package contains the 'Veritas Storage Foundation Messaging Service' (xprtld) that is used by both, ApplicationHA and VOM.

Note: This issue also occurs when you uninstall the Veritas Operations Manager Central Server.

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:

```
# rpm -ivh VRTSsfmh-*.rpm
```

```
# pkgadd -d 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

This issue is typically observed in case of IE7 browser.

Symantec High Availability view of Veritas Operations Manager 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. (2379946, 2379707)

If you click **Ok** on the error message and then click another virtual machine on the VOM Management Server console, 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 or tab, 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

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. (2607134)

Workaround

Revert the VCS configuration to the read-only mode by using the following command:

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

Configuration option of SFHA Solutions installer malfunctions

When you run the Symantec Storage Foundation and High Availability Solutions installer, it displays the following option to configure SFHA Solutions: **Configure an Installed Product**.

If you specify this option, the installer fails to configure SFHA Solutions. Instead, the installer starts stopping certain SFHA Solutions processes. (2621468)

Workaround

Do not use the installer option to configure an application. Instead, to configure Symantec Storage Foundation and High Availability Solutions for monitoring an application, use one of the following methods:

- If you have already installed SFHA Solutions, navigate to the following URL, and use the **Configure Application Monitoring** link to launch the Symantec Storage Foundation and High Availability Solutions Application Monitoring Configuration Wizard:

  ```
  https://<logicalDomainNameorIPaddress>:5634/vcs/admin/
  application_health.html?priv=ADMIN
  ```

- 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 SFHA Solutions, see the *Symantec ApplicationHA User's Guide*.

Heartbeat service group may fail to come online

If the high availability daemon (HAD) on the guest domain is restarted, the configured heartbeat service group (VCSAppMonHBSG) does not automatically come online. (2605506)

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 guest domain may retain stale values

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. (2611726)

Application monitoring may not resume on exiting maintenance mode

If you suspend application monitoring from the Symantec High Availability view, and later resume it by clicking **Exit Maintenance Mode**, application monitoring may not resume as expected.

If you then refresh the tab/view, the Configure Application Monitoring link may appear. This issue may occasionally occur in any setup because of variations in the time taken by internal commands to complete, after you click **Exit Maintenance Mode**. In some cases, before the internal commands are complete, the High Availability Daemon (HAD) module may shut down and stop application monitoring. (3640282)

Workaround

Perform the following steps:

1. From the command line, execute the following command:

   ```
   # /opt/VRTSvcs/bin/hastart -onenode
   ```

2. From the Symantec High Availability view, try to resume application monitoring by clicking **Exit Maintenance Mode**.

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

I/O fencing known issues

This section describes the known issues in this release of I/O fencing.

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 SFCFSHA cluster (application cluster), the installer also fails.

Workaround: Perform the following procedure on all of the nodes of the CP server.

To resolve this issue

1 Rename `cpsadm` to `cpsadmbin`:

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

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

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

3 Change the permissions of the new file to 775:

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

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
```

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.

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.

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.

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.

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.

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

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

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

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

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

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

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

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

Workaround: Start VxFEN again after some time.

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.

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.

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.

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 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.

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.

The `vxfenconfig -l` command output does not list Coordinator disks that are removed using the `vxdmpadm exclude dmpnodename=<dmp_disk/node>` command [3644431]

After you remove a Coordinator disk used by fencing or fencing disk group by running the `vxdmpadm exclude dmpnodename=<dmp_disk/node>` command, the removed disk is not listed in the `vxfenconfig -l` command output.

In case of a split brain, the `vxfen` program cannot use the removed disk as a coordination point in the subsequent fencing race.

Workaround: Run the `vxdmpadm include dmpnodename=<dmp_disk/node>` command to again enable the dmp disk. This disk will show up in subsequent `vxfencondig -l` output.

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.

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
    ```

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.

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*.

The upper bound value of FaultTolerance attribute of CoordPoint agent should be less than the majority of the coordination points. (2846389)

The upper bound value of `FaultTolerance` attribute of `CoordPoint` agent should be less than the majority of the coordination points. Currently this value is less than the number of coordination points.

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.

Software limitations

This section covers the software limitations of this release.

Note: For software limitations inherited from previous releases, see the 6.2 component product release notes. The latest product documentation is available on the Symantec website at:

https://sort.symantec.com/documents

Veritas Volume Manager software limitations

The following are software limitations in this release of Veritas Volume Manager.

It is not recommended to create a disk group with large number of objects, otherwise "Out of kernel memory" error may be reported while creating disk group or while splitting, joining or moving such a disk group (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.

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.

Disks on the Oracle VM Server for SPARC guest are claimed under other_disks category (2354005)

The disks on the Oracle VM Server for SPARC 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.

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.

Replication software limitations

The following are replication software limitations in this release of Symantec Storage Foundation and High Availability Solutions.

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.

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
```

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.

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.

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.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.

Node may panic if HAD process is stopped by force and then node is shut down or restarted [3640007]

A node may panic if the HAD process running on it is stopped by force and then it is shut down or restarted.This limitation is observed when you perform the following steps on a cluster node:

1 Stop the HAD process with the `force` flag.

```
# hastop -local -force
```

or

```
# hastop -all -force
```

2 Restart or shut down the node.

The node panics because forcefully stopping VCS on the node leaves all the applications, file systems, CVM, and other process online on that node. If the same node is restarted in this starte, VCS triggers a fencing race to avoid data curruption. However, the restarted node loses the fencing race and panics.

Workaround: No workaround.

Limitations related to bundled agents

Programs using networked services may stop responding if the host is disconnected

Programs using networked services (for example, NIS, NFS, RPC, or a TCP socket connection to a remote host) can stop responding if the host is disconnected from the network. If such a program is used as an agent entry point, a network disconnect can cause the entry point to stop responding and possibly time out.

For example, if the host is configured to use NIS maps as a client, basic commands such as `ps -ef` can hang if there is network disconnect.

Symantec recommends creating users locally. To reflect local users, configure:

/etc/nsswitch.conf

Volume agent clean may forcibly stop volume resources

When the attribute FaultOnMonitorTimeouts calls the Volume agent clean entry point after a monitor time-out, the `vxvol -f stop` command is also issued. This command forcibly stops all volumes, even if they are still mounted.

False concurrency violation when using PidFiles to monitor application resources

The PID files created by an application contain the PIDs for the processes that are monitored by Application agent. These files may continue to exist even after a node running the application crashes. On restarting the node, the operating system may assign the PIDs listed in the PID files to other processes running on the node.

Thus, if the Application agent monitors the resource using the PidFiles attribute only, the agent may discover the processes running and report a false concurrency violation. This could result in some processes being stopped that are not under VCS control.

Volumes in a disk group start automatically irrespective of the value of the StartVolumes attribute in VCS [2162929]

Volumes in a disk group are started automatically when the disk group is imported, irrespective of the value of the StartVolumes attribute in VCS. This behavior is observed if the value of the system-level attribute `autostartvolumes` in Veritas Volume Manager is set to On.

Workaround: If you do not want the volumes in a disk group to start automatically after the import of a disk group, set the autostartvolumes attribute to Off at the system level.

Online for LDom resource fails [2517350]

Online of LDom resource fails when the boot disk configured in the guest domain that is a part of the virtual disk multi-pathing group (mpgroup) and also the primary path to the virtual disk is not available.

This is due to the limitations in Oracle VM Server that do not allow retrying of other device paths that exist for the virtual disks, which are part of a virtual disk multi-pathing group, when booting a guest domain.

Workaround: None.

Zone agent registered to IMF for Directory Online event

The Directory Online event monitors the Zone root directory. If the parent directory of the Zone root directory is deleted or moved to another location, AMF does not provide notification to the Zone agent. In the next cycle of the zone monitor, it detects the change and reports the state of the resource as offline.

LDom resource calls clean entry point when primary domain is gracefully shut down

LDom agent sets failure policy of the guest domain to stop when primary domain stops. Thus when primary domain is shut down, guest domain is stopped. Moreover, when primary domain is shutdown, ldmd daemon is stopped abruptly and LDom configuration cannot be read. These operations are not under VCS control and VCS may call clean entry point.

Workaround: No workaround.

Application agent limitations

- ProPCV fails to prevent execution of script-based processes configured under MonitorProcesses.

Interface object name must match net<*x*>/v4static for VCS network reconfiguration script in Solaris 11 guest domain [2840193]

If the Solaris 11 guest domain is configured for DR and its interface object name does not match the `net<x>/v4static` pattern then the VCS guest network reconfiguration script (VRTSvcsnr) running inside the guest domain adds a new interface object and the existing entry will remain as is.

Share agent limitation (2717636)

If the Share resource is configured with VCS to share a system directory (for example, /usr) or Oracle Solaris 11 which gets mounted at boot time, the VCS share resource detects it online once VCS starts on the node after a panic or halt. This can lead to a concurrency violation if the share resource is a part of a failover service group, and the group has failed over to another node in the cluster. VCS brings down the Share resource subsequently. This is due to the share command behavior or Oracle Solaris 11, where a directory shared with share command remains persistently on the system across reboots.

Campus cluster fire drill does not work when DSM sites are used to mark site boundaries [3073907]

The campus cluster FireDrill agent currently uses the SystemZones attribute to identify site boundaries. Hence, campus cluster FireDrill is not supported in DSM enabled environment.

Workaround: Disable DSM and configure the SystemZones attribute on the application service group to perform the fire drill.

On Solaris 10, the online operation of IP resource may fail if `ifconfig -a` returns an error [3609861]

The IP agent uses the output of `ifconfig -a` to determine the next alias of free NIC to plumb IP. In rare and specific scenarios, the `ifconfig -a` command may return an error if it does not find an interface at the time of listing the interface. The IP resource online operation is affected by this and the resource may fault.

Workaround: Increase OnlineRetryLimit to a value higher than the default value.

Mount agent reports resource state as OFFLINE if the configured mount point does not exist [3435266]

If a configured mount point does not exist on a node, then the Mount agent reports the resource state as OFFLINE instead of UNKNOWN on that particular node. If an attempt is made for onlining the resource, it fails on that node as the mount point does not exist.

Workaround: Make sure that configured mount point exists on all nodes of the cluster or alternatively set the CreateMntPt attribute value of Mount agent to 1. This will ensure that if a mount point does not exist then it will create while onlining the resource.

Limitations related to VCS engine

Loads fail to consolidate and optimize when multiple groups fault [3074299]

When multiple groups fault and fail over at the same time, the loads are not consolidated and optimized to choose the target systems.

Workaround: No workaround.

Preferred fencing ignores the forecasted available capacity [3077242]

Preferred fencing in VCS does not consider the forecasted available capacity for fencing decision. The fencing decision is based on the system weight configured.

Workaround: No workaround.

Failover occurs within the SystemZone or site when BiggestAvailable policy is set [3083757]

Failover always occurs within the SytemZone or site when the BiggestAvailable failover policy is configured. The target system for failover is always selected based on the biggest available system within the SystemZone.

Workaround: No workaround.

Load for Priority groups is ignored in groups with BiggestAvailable and Priority in the same group[3074314]

When there are groups with both BiggestAvailable and Priority as the failover policy in the same cluster, the load for Priority groups are not considered.

Workaround: No workaround.

Symantec cluster configuration wizard limitations

Environment variable used to change log directory cannot redefine the log path of the wizard [3609791]

By default, the Symantec cluster configuration wizard writes the logs in `/var/VRTSvcs/log` directory. VCS provides a way to change the log directory through environment variable VCS_LOG, but this does not apply to the logs of VCS wizards.

Workaround: No workaround.

Cluster configuration wizard takes long time to configure a cluster on Solaris systems [3582495]

Some times the VCS cluster configuration wizard takes a long time (10 to 15 minutes) to configure a VCS cluster on Solaris systems. The wizard may appear stuck but it completes the configuration in some time.

Workaround: No workaround.

Limitations related to the VCS database agents

DB2 RestartLimit value [1234959]

When multiple DB2 resources all start at the same time with no dependencies, they tend to interfere or race with each other. This is a known DB2 issue.

The default value for the DB2 agent RestartLimit is 3. This higher value spreads out the re-start of the DB2 resources (after a resource online failure), which lowers the chances of DB2 resources all starting simultaneously.

Sybase agent does not perform qrmutil based checks if Quorum_dev is not set (2724848)

If you do not set the Quorum_dev attribute for Sybase Cluster Edition, the Sybase agent does not perform the qrmutil-based checks. This error in configuration may lead to undesirable results. For example, if qrmutil returns failure pending, the agent does not panic the system. Thus, the Sybase agent does not perform qrmutil-based checks because the Quorum_dev attribute is not set.

Therefore, setting Quorum_Dev attribute is mandatory for Sybase cluster edition.

Pluggable database (PDB) online may timeout when started after container database (CDB) [3549506]

PDB may take long time to start when it is started for the first time after starting CDB. As a result, the PDB online initiated using VCS may cause ONLINE timeout and the PDB online process may get cancelled.

Workaround: Increase the OnlineTimeout attribute value of the Oracle type resource.

Limitations related to global clusters

- Cluster address for global cluster requires resolved virtual IP.
  The virtual IP address must have a DNS entry if virtual IP is used for heartbeat agents.
- Total number of clusters in a global cluster configuration can not exceed four.
- Cluster may not be declared as faulted when Symm heartbeat agent is configured even when all hosts are down.
  The Symm agent is used to monitor the link between two Symmetrix arrays. When all the hosts are down in a cluster but the Symm agent is able to see the replication link between the local and remote storage, it would report the heartbeat as ALIVE. Due to this, DR site does not declare the primary site as faulted.
- Configuring Veritas Volume Replicator for Zone Disaster Recovery is not supported for zone root replication. Oracle Solaris 11 supports zone root only on ZFS file system.
- Configuring a cluster of mixed nodes such as a cluster between systems running on Solaris 10 and Solaris 11 versions is not supported in SFHA Solutions 6.2.1. The configuration is not supported through manual as well as CPI configuration.