

Veritas InfoScale™ 7.4.3 Release Notes - Linux

Last updated: 2020-11-17

Legal Notice

Copyright © 2020 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third-party ("Third-Party Programs"). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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

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

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

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054
<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

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

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

infoscaledocs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Introduction and product requirements	13
	About this document	13
	VCS system requirements	14
	Supported Linux operating systems	14
	Required Linux RPMs for Veritas InfoScale	16
	Storage Foundation for Databases features supported in database environments	20
	Storage Foundation memory requirements	21
	Supported database software	22
	Supported hardware and software	22
	VMware Environment	22
	Number of nodes supported	22
Chapter 2	Changes introduced in InfoScale 7.4.3	24
	Support for Kubernetes on Linux	24
	Support for protection against ransomware	25
	Updates to volume encryption keys for FIPS compliance	26
	Support for configuring LLT over UDP multiport	26
	Support for SambaServer agent and SFAE service in systemd environments	27
	Limitations on supported upgrade paths	27
	Support for rolling upgrade of cluster nodes with different versions of VCS engine	27
Chapter 3	Limitations	29
	Virtualization software limitations	29
	Paths cannot be enabled inside a KVM guest if the devices have been previously removed and re-attached from the host	29
	Application component fails to come online [3489464]	30
	Storage Foundation software limitations	31
	Dynamic Multi-Pathing software limitations	31
	Veritas Volume Manager software limitations	32
	Veritas File System software limitations	33
	SmartIO software limitations	36

Replication software limitations	37
Softlink access and modification times are not replicated on RHEL5 for VFR jobs	38
VVR Replication in a shared environment	38
VVR IPv6 software limitations	38
VVR support for replicating across Storage Foundation versions	38
Cluster Server software limitations	38
Limitations related to bundled agents	39
Limitations related to VCS engine	41
Veritas cluster configuration wizard limitations	42
Limitations related to the VCS database agents	42
Security-Enhanced Linux is not supported on SLES distributions	43
Systems in a cluster must have same system locale setting	43
VxVM site for the disk group remains detached after node reboot in campus clusters with fire drill [1919317]	43
Limitations with DiskGroupSnap agent [1919329]	44
System reboot after panic	44
Host on RHEV-M and actual host must match [2827219]	45
Cluster Manager (Java console) limitations	45
Limitations related to LLT	45
Limitations related to I/O fencing	46
Limitations related to global clusters	47
Clusters must run on VCS 6.0.5 and later to be able to communicate after upgrading to 2048 bit key and SHA256 signature certificates [3812313]	48
Storage Foundation Cluster File System High Availability software limitations	48
cfsmntadm command does not verify the mount options (2078634)	48
Obtaining information about mounted file system states (1764098)	48
Stale SCSI-3 PR keys remain on disk after stopping the cluster and deporting the disk group	49
Unsupported FSS scenarios	49
Storage Foundation for Oracle RAC software limitations	49
Supportability constraints for normal or high redundancy ASM disk groups with CVM I/O shipping and FSS (3600155)	49
Limitations of CSSD agent	49
Oracle Clusterware/Grid Infrastructure installation fails if the cluster name exceeds 14 characters	50

Policy-managed databases not supported by CRSResource agent	50
Health checks may fail on clusters that have more than 10 nodes	50
Cached ODM not supported in Veritas InfoScale environments	50
Storage Foundation for Databases (SFDB) tools software limitations	51
Parallel execution of <code>vxsfadm</code> is not supported (2515442)	51
Creating point-in-time copies during database structural changes is not supported (2496178)	51
Oracle Data Guard in an Oracle RAC environment	51

Chapter 4	Known issues	52
	Documentation errata	52
	Issues related to installation and upgrade	53
	During upgrade from InfoScale 7.4.1 to 7.4.3, xprtld fails to start without node restart (4017516)	53
	Upgrade from SFRAC 6.2.1 to InfoScale 7.4.2 using yum may fail (4002582)	53
	Installer fails on SLES 11 and SLES 12 for older NTP versions (3912493)	54
	InfoScale deployments fail when secure boot is enabled (4005536)	54
	Switch fencing in enable or disable mode may not take effect if VCS is not reconfigured [3798127]	54
	During an upgrade process, the AMF_START or AMF_STOP variable values may be inconsistent [3763790]	54
	Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups (2574731)	55
	NetBackup 6.5 or older version is installed on a VxFS file system (2056282)	55
	Error messages in syslog (1630188)	56
	Ignore certain errors after an operating system upgrade—after a product upgrade with encapsulated boot disks (2030970)	56
	After a locale change restart the vxconfig daemon (2417547, 2116264)	57
	Dependency may get overruled when uninstalling multiple RPMs in a single command [3563254]	57

When using the response file, the installer must not proceed with the installation or upgrade, if you have not provided edge server details (3964335)	57
Unable to update edge server details by running the installer (3964611)	58
Storage Foundation known issues	58
Dynamic Multi-Pathing known issues	58
Veritas Volume Manager known issues	60
Veritas File System known issues	77
Virtualization known issues	81
Replication known issues	86
The secondary vradmind may appear hung and the vradmind commands may fail (3940842,3944301)	87
Data corruption may occur if you perform a rolling upgrade of InfoScale Storage or InfoScale Enterprise from 7.3.1 or earlier to 7.4 or later during replication (3951527)	87
vradmind may appear hung or may fail for the role migrate operation (3968642, 3968641)	88
After the product upgrade on secondary site, replication may fail to resume with "Secondary SRL missing" error [3931763]	88
<code>vradmind repstatus</code> command reports secondary host as "unreachable"(3896588)	89
RVGPrimary agent operation to start replication between the original Primary and the bunker fails during failback (2036605)	89
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]	90
In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon (1672410, 1672417, 1825031)	91
vradmind functionality may not work after a master switch operation [2158679]	91
Cannot relay layout data volumes in an RVG from concat to striped-mirror (2129601)	91
vradmind verifydata may report differences in a cross-endian environment (2834424)	92
vradmind verifydata operation fails if the RVG contains a volume set (2808902)	92
Plex reattach operation fails with unexpected kernel error in configuration update (2791241)	93

Bunker replay does not occur with volume sets (3329970)	93
SmartIO does not support write-back caching mode for volumes configured for replication by Volume Replicator (3313920)	93
During moderate to heavy I/O, the vradmin verifydata command may falsely report differences in data (3270067)	94
While vradmin commands are running, vradmind may temporarily lose heartbeats (3347656, 3724338)	94
Write I/Os on the primary logowner may take a long time to complete (2622536)	95
DCM logs on a disassociated layered data volume results in configuration changes or CVM node reconfiguration issues (3582509)	95
After performing a CVM master switch on the secondary node, both rlinks detach (3642855)	95
vradmin -g dg repstatus rvg displays the following configuration error: vradmind not reachable on cluster peer (3648854)	95
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)	96
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)	97
DCM plex becomes inaccessible and goes into DISABLED(SPARSE) state in case of node failure. (3931775)	98
Initial autosync operation takes a long time to complete for data volumes larger than 3TB (3966713)	98
Cluster Server known issues	98
Operational issues for VCS	99
Issues related to the VCS engine	102
Issues related to the bundled agents	108
Issues related to the VCS database agents	118
Issues related to the agent framework	123
Cluster Server agents for Volume Replicator known issues	126
Issues related to Intelligent Monitoring Framework (IMF)	129
Issues related to global clusters	132
Issues related to the Cluster Manager (Java Console)	132
VCS Cluster Configuration wizard issues	133

LLT known issues	135
I/O fencing known issues	138
Storage Foundation and High Availability known issues	143
Cache area is lost after a disk failure (3158482)	143
Installer exits upgrade to 5.1 RP1 with Rolling Upgrade error message (1951825, 1997914)	143
In an IPv6 environment, db2icrt and db2idrop commands return a segmentation fault error during instance creation and instance removal (1602444)	144
Process start-up may hang during configuration using the installer (1678116)	145
Not all the objects are visible in the VOM GUI (1821803)	145
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)	145
A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a dynamic storage tiering placement policy (1880081)	146
Storage Foundation Cluster File System High Availability known issues	146
During disk group creation with '-o same_enckey=yes', disk reattach fails if a disk comes online after losing connectivity (4003890)	146
Transaction hangs when multiple plex-attach or add-mirror operations are triggered on the same volume (3969500)	147
In an FSS environment, creation of mirrored volumes may fail for SSD media [3932494]	148
Mount command may fail to mount the file system (3913246)	148
After the local node restarts or panics, the FSS service group cannot be online successfully on the local node and the remote node when the local node is up again (3865289)	149
In the FSS environment, if DG goes to the dgdisable state and deep volume monitoring is disabled, successive node joins fail with error 'Slave failed to create remote disk: retry to add a node failed' (3874730)	149
DG creation fails with error "V-5-1-585 Disk group punedatadg: cannot create: SCSI-3 PR operation failed" on the VSCSI disks (3875044)	150
CVMVOLDg agent is not going into the FAULTED state. [3771283]	150

On CFS, SmartIO is caching writes although the cache appears as nocache on one node (3760253)	151
tail -f run on a cluster file system file only works correctly on the local node [3741020]	151
CFS commands might hang when run by non-root (3038283)	151
The fsappadm subfilemove command moves all extents of a file (3258678)	152
Certain I/O errors during clone deletion may lead to system panic. (3331273)	152
Panic due to null pointer de-reference in vx_bmap_lookup() (3038285)	152
In a CFS cluster, that has multi-volume file system of a small size, the fsadm operation may hang (3348520)	153
Storage Foundation for Oracle RAC known issues	153
Oracle RAC known issues	153
Storage Foundation Oracle RAC issues	154
Storage Foundation for Databases (SFDB) tools known issues	162
Clone operations fail for instant mode snapshot (3916053)	163
Sometimes SFDB may report the following error message: SFDB remote or privileged command error (2869262)	163
SFDB commands do not work in IPV6 environment (2619958)	163
When you attempt to move all the extents of a table, the dbdst_obj_move(1M) command fails with an error (3260289)	164
Attempt to use SmartTier commands fails (2332973)	164
Attempt to use certain names for tiers results in error (2581390)	164
Clone operation failure might leave clone database in unexpected state (2512664)	165
Clone command fails if PFILE entries have their values spread across multiple lines (2844247)	165
Clone command errors in a Data Guard environment using the MEMORY_TARGET feature for Oracle 11g (1824713)	165
Clone fails with error "ORA-01513: invalid current time returned by operating system" with Oracle 11.2.0.3 (2804452)	166
Data population fails after datafile corruption, rollback, and restore of offline checkpoint (2869259)	167
Flashsnap clone fails under some unusual archivelog configuration on RAC (2846399)	167
In the cloned database, the seed PDB remains in the mounted state (3599920)	168

Cloning of a container database may fail after a reverse resync commit operation is performed (3509778)	168
If one of the PDBs is in the read-write restricted state, then cloning of a CDB fails (3516634)	168
Cloning of a CDB fails for point-in-time copies when one of the PDBs is in the read-only mode (3513432)	169
If a CDB has a tablespace in the read-only mode, then the cloning fails (3512370)	169
SFDB commands fail when an SFDB installation with authentication configured is upgraded to InfoScale 7.4.3 (3644030)	169
Benign message displayed upon execution of <code>vxsfadm -a oracle -s filesnap -o destroyclone</code> (3901533)	170
Application isolation feature known Issues	170
Addition of an Oracle instance using Oracle GUI (dbca) does not work with Application Isolation feature enabled	171
Auto reattach of detached plexes may not happen for FSS disk groups when auto-mapping feature is used (3902004)	171
CPI is not supported for configuring the application isolation feature (3902023)	171
Thin reclamation does not happen for remote disks if the storage node or the disk owner does not have the file system mounted on it (3902009)	171
Cloud deployment known issues	172
Systems in GCP may get stuck in the LEAVING state when multiple nodes are restarted a cascaded manner	172
An error occurs during VVR or CVR configuration when alias IPs are assigned to GCP VM instances (3965275)	172
In an Azure environment, the systems under InfoScale control may panic due to CPU soft lockup [3929534]	172
In an Azure environment, an InfoScale cluster node may panic if any of the node is rebooted using Azure portal [3930926]	173
If you disable a public IP from the Azure portal, the corresponding AzureIP resource goes into UNKNOWN state [3928222]	173
After rolling upgrade phase 1, xprtld service fails to start on AWS instances (4004450)	174
Issues related to Veritas InfoScale Storage in Amazon Web Services cloud environments	174
Container deployment know issues	177
Creation of full sized snapshot of encrypted volume may lead to loss of data accessibility (4010594)	177

After restoring space-optimized snapshot to new PVC, mount on restored PVC may fail if the source snapshot volume is detached (4012858)	177
CSI controller pod remains in 'Terminating' state in case of graceful node shutdown or power-off (4011482)	178
CSI node pods does not get rescheduled on other worker nodes when its parent node is drained (4011384)	178
While restoring a snapshot, PVC goes into pending state after rebooting all nodes except the master node in the cluster (4014525)	179
InfoScale cluster on the worker node panics when network split occurs and the master node is disconnected from all the worker nodes. (4018090)	179
When a node in the Kubernetes cluster is shutdown gracefully, fencing is not invoked (4011482)	179
In case of storage failure, application IOs to the mountpoint inside container fails and pod goes into CreateContainerConfigError or Error state (4011219, 4014758, 4015259)	180
Pod gets stuck in 'Terminating' state and file system remains mounted (4017013)	180
When more than 20 containers are started simultaneously on a worker node, some of these containers may fail to start (4015069)	180
While stopping InfoScale on a Kubernetes worker node, the InfoScale cluster gets stuck in the LEAVING state (4018544)	181

Introduction and product requirements

This chapter includes the following topics:

- [About this document](#)
- [VCS system requirements](#)
- [Supported Linux operating systems](#)
- [Storage Foundation for Databases features supported in database environments](#)
- [Storage Foundation memory requirements](#)
- [Supported database software](#)
- [Supported hardware and software](#)
- [VMware Environment](#)
- [Number of nodes supported](#)

About this document

This document provides information that is specific to version 7.4.3 of the Veritas InfoScale products.

Review this entire document before using the following products:

- Veritas InfoScale Foundation
- Veritas InfoScale Storage
- Veritas InfoScale Availability
- Veritas InfoScale Enterprise

The information in this document supersedes the information provided in the product-specific documents.

You can download the latest version of this document from the Veritas Service and Operations Readiness Tools (SORT) website at:

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

The following documents provide further information that is common to all the InfoScale for Linux products:

- *Veritas InfoScale Getting Started Guide*
- *Veritas InfoScale Installation Guide*

For information about the InfoScale product components and their capabilities, refer to the corresponding configuration and upgrade guides and administrator's guides.

For information about installing and configuring and your databases with the InfoScale products, refer to the database-specific installation and configuration guides.

For the latest information on updates, patches, and known issues regarding this release, see the Late Breaking News (LBN) at:

https://www.veritas.com/support/en_US/article.100048832

VCS system requirements

This section describes system requirements for VCS.

The following information does not apply to SF Oracle RAC installations.

VCS requires that all nodes in the cluster use the same processor architecture and run the same operating system version. However, the nodes can have different update levels or service pack levels for a specific version of RHEL, OL, SLES or CentOS.

Note: The system from where you install VCS must run the same Linux distribution as the target systems.

Supported Linux operating systems

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

Table 1-1 Supported operating systems

Operating systems	Kernel version
Red Hat Enterprise Linux 7	Update 7 (3.10.0-1062.el7)
Red Hat Enterprise Linux 8	Update 1 (4.18.0-147.el8)
Oracle Linux 7 (RHEL Compatible mode)	Update 7 (3.10.0-1062.el7)
Oracle Linux 8 (RHEL compatible mode)	Update 1 (4.18.0-147.el8.x86_64)
Community Enterprise Operating System (CentOS) 7	Update 7 (3.10.0-1062.el7)
Community Enterprise Operating System (CentOS) 8	Update 1 (4.18.0-147.el8)
SUSE Linux Enterprise 12	SP4 (4.12.14-94.41-default) SP5(4.12.14-120.1-default)
SUSE Linux Enterprise 15	SP1 (4.12.14-195-default)

Note: All subsequent kernel versions and patch releases on the supported operating system levels are supported, but you should check the Veritas Services and Operations Readiness Tools (SORT) website for additional information that applies to the exact kernel version for which you plan to deploy.

Note: Only 64-bit operating systems are supported on the AMD Opteron or the Intel Xeon EM64T (x86_64) Processor line.

Note: SmartIO and FSS are not supported on platforms for which device drivers of Fusion-io SSD cards are not available.

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

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

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

Required Linux RPMs for Veritas InfoScale

Make sure you install the following operating system-specific RPMs on the systems where you want to install or upgrade Veritas InfoScale. Veritas InfoScale will support any updates made to the following RPMs, provided the RPMs maintain the ABI compatibility.

Note: The required RPM versions should be equal or later than the list in the following table.

Table 1-2 Required RPMs

Operating system	Required RPMs
RHEL 7	glibc-2.17-55.el7.i686
Note: Veritas recommends that you install RHEL 7 as the operating system of Server GUI.	glibc-2.17-55.el7.x86_64
	glibc-2.17-55.el7.i686
	net-tools.x86_64
	binutils.x86_64
	coreutils.x86_64
	ed.x86_64
	ethtool.x86_64
	findutils.x86_64
	glibc.x86_64
	kmod.x86_64
	ksh.x86_64
	libacl.x86_64
	libgcc.x86_64
	libhbaapi.x86_64
	libhbalinux.x86_64
	libstdc++.x86_64
	ncurses-libs.x86_64
	openssl-libs.x86_64
	perl-Exporter.noarch
	perl-Socket.x86_64
	perl.x86_64
	policycoreutils-python-utils.noarch
	policycoreutils.x86_64
	zlib.x86_64

Table 1-2 Required RPMs (continued)

Operating system	Required RPMs
RHEL 8	chrony.x86_64 glibc-2.28-42.el8.i686 glibc-2.28-42.el8.x86_64 glibc-2.28-42.el8.i686 binutils.x86_64 compat-openssl10.x86_64 coreutils.x86_64 ed.x86_64 ethtool.x86_64 findutils.x86_64 glibc.x86_64 kmod.x86_64 ksh.x86_64 libacl.x86_64 libgcc.x86_64 libnsl.x86_64 libstdc++.x86_64 ncurses-compat-libs.x86_64 perl-Exporter.noarch perl-Socket.x86_64 perl.x86_64 policycoreutils-python-utils.noarch policycoreutils.x86_64 zlib.x86_64

Table 1-2 Required RPMs *(continued)*

Operating system	Required RPMs
SLES 12	glibc-32bit-2.19-17.72.x86_64 bc.x86_64 binutils.x86_64 coreutils.x86_64 ed.x86_64 ethhtool.x86_64 findutils.x86_64 glibc.x86_64 kmod-compatible.x86_64 kmod-compatible.x86_64 libHBAAPI2.x86_64 libacl1.x86_64 libgcc_s1.x86_64 libhbalinux2.x86_64 libstdc++6.x86_64 libz1.x86_64 mksh.x86_64 policycoreutils-python-utils.noarch policycoreutils.x86_64

Table 1-2 Required RPMs (continued)

Operating system	Required RPMs
SLES 15	<div>chrony.x86_64</div> <div>glibc-32bit-2.19-17.72.x86_64</div> <div>bc.x86_64</div> <div>binutils.x86_64</div> <div>coreutils.x86_64</div> <div>ed.x86_64</div> <div>ethtool.x86_64</div> <div>glibc.x86_64</div> <div>kmod-compatible.x86_64</div> <div>libacl1.x86_64</div> <div>libgcc_s1.x86_64</div> <div>libstdc++6.x86_64</div> <div>libz1.x86_64</div> <div>mksh.x86_64</div> <div>parted.x86_64</div> <div>policycoreutils-python-utils.noarch</div> <div>policycoreutils.x86_64</div> <div>Note: Ensure that you install the iputils and insserv-compatible package manually. The iputils package is required for the ping command to work and insserv-compatible package is required to enable the vxdbd daemon (# /opt/VRTS/bin/sfae_config enable).</div>

See “Supported database software” on page 22.

Storage Foundation for Databases features supported in database environments

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

Table 1-3 SFDB features supported in database environments

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

See [“Supported database software”](#) on page 22.

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

Storage Foundation memory requirements

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

Supported database software

For the latest information on supported databases, see the database support matrices at:

- IBM DB2:
https://www.veritas.com/content/support/en_US/doc/112638608-112638611-1
- Oracle:
https://www.veritas.com/content/support/en_US/doc/112632971-112632974-1
Additionally, visit the following Oracle support site for information on patches that may be required by Oracle for each release.
- Sybase:
https://www.veritas.com/content/support/en_US/doc/112512557-113400602-1

Supported hardware and software

For the latest information on the supported hardware and software, see the appropriate compatibility list at:

https://www.veritas.com/content/support/en_US

Click **Documentation**, and on the Documentation tab, and select the appropriate **Product**, **Document Type**, and **Version** filters.

Before installing or upgrading the InfoScale products, review the current compatibility list to confirm the compatibility of your hardware and software. For information on specific setup requirements, see the corresponding Configuration and Upgrade Guide.

VMware Environment

[Table 1-4](#) lists the support VMWare ESX versions in 7.4.3.

Table 1-4 Supported VMWare ESX versions

Operating System	Update
VMware vSphere 6.0.0	update 3
VMware vSphere 6.5	update 1

Number of nodes supported

Veritas InfoScale supports cluster configurations up to 128 nodes.

SFHA, SFCFSHA, SF Oracle RAC: Flexible Storage Sharing (FSS) only supports cluster configurations with up to 64 nodes.

SFHA, SFCFSHA: SmartIO writeback caching only supports cluster configurations with up to 2 nodes.

Changes introduced in InfoScale 7.4.3

This chapter includes the following topics:

- [Support for Kubernetes on Linux](#)
- [Support for protection against ransomware](#)
- [Updates to volume encryption keys for FIPS compliance](#)
- [Support for configuring LLT over UDP multiport](#)
- [Support for SambaServer agent and SFAE service in systemd environments](#)
- [Limitations on supported upgrade paths](#)
- [Support for rolling upgrade of cluster nodes with different versions of VCS engine](#)

Support for Kubernetes on Linux

InfoScale Enterprise supports Kubernetes to orchestrate and schedule containers in InfoScale clusters. While containers virtualize the applications from the underlying virtual or physical infrastructure, Kubernetes orchestrates, and scales containerized applications in real time.

InfoScale provides the following advanced functionality for containers and Kubernetes that enterprise applications need:

- **Application availability**—High availability (HA) of data is one of the prime considerations when you plan for business continuity in case of disasters. InfoScale manages HA and recovery automation for your applications by monitoring critical application processes and resources. InfoScale provides persistent storage for stateful container applications. InfoScale fencing and

arbitration prevent data corruption and speed up recovery. The use of VCS agents and persistent volumes ensure that there is no loss of data when containers fail. Consequentially data is always available no matter where the containers are scheduled in the container ecosystem.

- **Advanced storage management**—The InfoScale Container Storage Interface (CSI) plugin provides high-performance shared storage for the Kubernetes clusters using the fast storage that is directly attached to the Kubernetes cluster nodes. InfoScale Storage provides highly available, persistent storage that conforms to CSI specifications for enterprise applications. It does so by using high-performance parallel storage access on shared storage (SAN) or in Flexible Storage Sharing (FSS) environments. InfoScale also supports Velero, which is a third-party application, to provide snapshot lifecycle management.
- **Application migration**—InfoScale supports moving non-containerized applications into a container environment by copying the application data storage volumes onto the Kubernetes cluster nodes. The InfoScale CSI plugin then presents the same data to the applications that are running in the container.

The key benefits to using InfoScale with Kubernetes are:

- High-performance parallel storage that provides better performance and reliability than NFS
- Optimized resource utilization with the ability to use either existing SAN storage or the InfoScale advanced FSS option that provides better performance than SAN at a reduced cost
- Migration of data from legacy to containerized MySQL applications
- Support for in-depth application and infrastructure monitoring with the VCS application HA agents that run inside the containers or inside the sidecar container
- Integrated I/O fencing and arbitration to protect against data corruption and to provide fast recovery in the event of a failure
- Snapshot copies of the production data for analytics and disaster recovery
- Volume cloning to address storage disk and node failures
- Support for statefulset applications like MySQL

Support for protection against ransomware

To ensure compliance with the security and governance requirements of critical data, support for the following feature is added to the Veritas File System (VxFS) module:

- Write Once Read Many (WORM) storage system
- Secure clock
- Audit logging

For details, refer to the following documents for this release:

- *Storage Foundation Administrator's Guide*
- *Storage Foundation Cluster File System High Availability Administrator's Guide*

Updates to volume encryption keys for FIPS compliance

For the encrypted volumes that are created with disk group version 300 or later, InfoScale uses the FIPS 140-2 standards to validate the wrapping key (KEK) that secures volume encryption. For the encrypted volumes that are created with disk group version 290 or earlier, you must first upgrade to disk group version 300 or later. Thereafter, you can use the re-key operation to ensure that InfoScale uses the FIPS standards to validate the KEK.

For details on volume encryption, refer to the *Storage Foundation Cluster File System High Availability Administrator's Guide*.

Support for configuring LLT over UDP multiport

LLT uses UDP sockets for communication among the cluster nodes and creates one UDP socket for each LLT link. In Flexible Storage Sharing (FSS) environments, read-write operations may be performed on remote disks, and one socket per LLT link may not be enough for large data volumes. More sockets are needed to achieve the parallelism and the throughput that is required to meet the needs of applications that generate large volumes of data.

To create additional sockets per link, you can configure LLT over UDP multiport. These sockets are reserved for I/O shipping only.

Note: For the multiport feature to work, LLT requires at least six consecutive network port numbers to be configured.

When you configure LLT over UDP multiport, you can use the following new tunables to enable ports:

- `set-udpports` to change the port range to be used for I/O shipping, if you do not want to use the port range from 50002 onwards.

- `set-udpthreads` to specify the number of threads to be created for each socket.
- `set-udpsockets` to specify the number of sockets to be created for each link.

For details, refer to the *Cluster Server Configuration and Upgrade Guide*.

Support for SambaServer agent and SFAE service in systemd environments

The SambaServer agent is updated to support the functioning of file shares in systemd environments, in single-instance as well as multi-instance setups.

For details, refer to the *Cluster Server Bundled Agents Reference Guide*.

The SFAE service is updated to support the functioning of Storage Foundation for Databases (SFDB) tools in systemd environments.

Limitations on supported upgrade paths

The following limitations apply in case of upgrades:

- You can upgrade to InfoScale 7.4.3 only if the base version of your currently installed product is 7.4.1 or 7.4.2.
- Upgrades from earlier InfoScale versions to 7.4.3 are allowed on Red Hat Enterprise Linux only and not on CentOS and Oracle Linux. However, fresh installations of InfoScale 7.4.3 are allowed on all the supported Linux distributions.
- Upgrades from Oracle RAC configurations with earlier InfoScale versions to 7.4.3 are not allowed. However, fresh installations of InfoScale 7.4.3 with Oracle RAC configurations are allowed.

Support for rolling upgrade of cluster nodes with different versions of VCS engine

The rolling upgrade process minimizes the downtime of a cluster during an upgrade to the amount of time that it takes to fail over a service group. InfoScale 7.4.2 and later versions let you configure clusters with nodes that run different versions of the VCS engine. To support such configurations, the installer provides a rolling upgrade option for InfoScale components that upgrades the kernel RPMs and the VCS agent-related RPMs in the same process. The rolling upgrade may involve some application downtime while a node is upgraded, but there is zero cluster downtime.

Note: You must perform a rolling upgrade on a completely configured cluster.

For details, refer to the following documents for this release:

- *Cluster Server Configuration and Upgrade Guide*
- *Storage Foundation Cluster File System High Availability Configuration and Upgrade Guide*

Limitations

This chapter includes the following topics:

- [Virtualization software limitations](#)
- [Storage Foundation software limitations](#)
- [Replication software limitations](#)
- [Cluster Server software limitations](#)
- [Storage Foundation Cluster File System High Availability software limitations](#)
- [Storage Foundation for Oracle RAC software limitations](#)
- [Storage Foundation for Databases \(SFDB\) tools software limitations](#)

Virtualization software limitations

This section describes the virtualization software limitations in this release of the following products:

- Veritas InfoScale Foundation
- Veritas InfoScale Storage
- Veritas InfoScale Availability
- Veritas InfoScale Enterprise

Paths cannot be enabled inside a KVM guest if the devices have been previously removed and re-attached from the host

LUNs are exported to the KVM guest via virtio-scsi interface. When some physical link between the host and the SAN array fails for a certain time (45-60 seconds by default), the HBA driver in the host will remove the timed-out devices. When the

link is restored, these devices will be re-attached to the host; however, the access from inside the KVM guest to these devices cannot be automatically restored too without rebooting the system or manually re-attaching the devices. For DMP, these subpaths will remain in DISABLED state.

This is a known limitation of KVM.

Workaround:

From the KVM host, tune the `dev_loss_tmo` parameter of the Fibre Channel ports to a very large value, and set the `fast_io_fail_tmo` parameter to 15.

To restore access to the timed-out devices

- 1** Add the following lines into `/dev/udev/rules.d/40-kvm-device` file:

```
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add", \
    RUN+="/bin/sh -c 'grep -q off \
    /sys/class/fc_remote_ports/%k/fast_io_fail_tmo;if [ $? -eq 0 ]; \
    then echo 15 > /sys/class/fc_remote_ports/%k/fast_io_fail_tmo 2> \
    /dev/null;fi;'"
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add", \
    RUN+="/bin/sh -c 'echo 8000000 > \
    /sys/class/fc_remote_ports/%k/dev_loss_tmo 2> /dev/null'"
```

- 2** Create the `/etc/modprobe.d/qla2xxx.conf` file with the following content:

```
options qla2xxx qlport_down_retry=8000000
```

- 3** Create the `/etc/modprobe.d/scsi_transport_fc.conf` with the following content:

```
options scsi_transport_fc dev_loss_tmo=8000000
```

- 4** Rebuild the `initrd` file and reboot.

Application component fails to come online [3489464]

In the KVM virtualization environment, if you try to bring an application resource online, the online operation fails. This behavior is observed both from the command line interface as well as the High Availability view of the Veritas Operations Manager Management Server.

Workaround: Perform the following steps:

- 1 Set the locale of the operating system (OS) to default value, and then retry the operation. For detailed steps, see OS vendor documentation.
- 2 Restart High Availability Daemon (HAD).

Storage Foundation software limitations

These software limitations apply to the following products:

- Veritas InfoScale Foundation
- Veritas InfoScale Storage
- Veritas InfoScale Enterprise

Dynamic Multi-Pathing software limitations

These software limitations apply to the following products:

- Veritas InfoScale Foundation
- Veritas InfoScale Storage
- Veritas InfoScale Enterprise

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment,change the default values for the DMP tunable parameters.

[Table 3-1](#) describes the DMP tunable parameters and the new values.

Table 3-1 DMP settings for NetApp storage attached environment

Parameter name	Definition	New value	Default value
dmp_restore_interval	DMP restore daemon cycle	60 seconds.	300 seconds.
dmp_path_age	DMP path aging tunable	120 seconds.	300 seconds.

The change is persistent across reboots.

To change the tunable parameters

- 1 Issue the following commands:

```
# vxddmpadm settune dmp_restore_interval=60

# vxddmpadm settune dmp_path_age=120
```

- 2 To verify the new settings, use the following commands:

```
# vxddmpadm gettune dmp_restore_interval

# vxddmpadm gettune dmp_path_age
```

LVM volume group in unusable state if last path is excluded from DMP (1976620)

When a DMP device is used by a native LVM volume group, do not exclude the last path to the device. This can put the LVM volume group in an unusable state.

Veritas Volume Manager software limitations

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

Snapshot configuration with volumes in shared disk groups and private disk groups is not supported (2801037)

A snapshot configuration with volumes in the shared disk groups and private disk groups is not a recommended configuration. In this release, this configuration is not supported.

SmartSync is not supported for Oracle databases running on raw VxVM volumes

SmartSync is not supported for Oracle databases that are configured on raw volumes, because Oracle does not support the raw volume interface.

Veritas InfoScale does not support thin reclamation of space on a linked mirror volume (2729563)

The thin reclamation feature does not support thin reclamation for a linked mirror volume.

Cloned disks operations not supported for FSS disk groups

In this release, the VxVM cloned disks operations are not supported with FSS disk groups. If you clone a disk in the FSS disk groups, the cloned device cannot be imported. If you prefer to use hardware mirroring for disaster recovery purposes, you need to make sure that such devices should not be used to create FSS disk groups.

For more information, see the *Administrator's Guide*.

Thin reclamation requests are not redirected even when the ioship policy is enabled (2755982)

Reclamation requests fail from nodes that do not have local connectivity to the disks, even when the ioship policy is enabled. Reclamation I/Os are not redirected to another node.

Veritas Operations Manager does not support disk, disk group, and volume state information related to CVM I/O shipping feature (2781126)

The Veritas Operations Manager (VOM) does not support disk, disk group, and volume state information related to the I/O shipping feature introduced in this release of Cluster Volume Manager. New states such as Ifailed, lmissing or LDISABLED are introduced when I/O shipping is active because of storage disconnectivity.

Veritas File System software limitations

The following are software limitations in this release of Veritas File System.

Online migration from Ext4 to VxFS is not supported on RHEL 8.1

InfoScale does not support online migration from Ext4 to VxFS on RHEL 8.1 systems.

Limitations while managing Docker containers

- Administrative tasks: All VxFS and VxVM administrative commands, such as resize, add volumes, reorganize volume sets, so on are supported only on host nodes. These administrative commands cannot be executed inside Docker containers.

- Package installation only on host nodes: Installation and configuration of InfoScale solutions inside containers is not supported.
- Root volume: Veritas does not recommend exporting root volumes to Docker containers.
- Data loss because volume devices are not synchronized: If a volume is exported to a Docker container, some VxVM operations, such as removing volumes, deporting a disk group, renaming a volume, remirroring a disk group or volume, or restarting VxVM configuration daemon (vxconfigd) , can cause the volume device to go out of sync, which may cause data loss.

Linux I/O Scheduler for Database Workloads

Veritas recommends using the Linux deadline I/O scheduler for database workloads on both Red Hat and SUSE distributions.

To configure a system to use this scheduler, include the `elevator=deadline` parameter in the boot arguments of the GRUB or LILO configuration file.

The location of the appropriate configuration file depends on the system's architecture and Linux distribution:

Configuration File

`/boot/grub2/grub.cfg`

Architecture and Distribution

RHEL7 x86_64, RHEL8 x86_64, SLES12 x86_64 and SLES15 x86_64

For the GRUB configuration files, add the `elevator=deadline` parameter to the kernel command.

For RHEL 7, add `elevator=deadline` to `GRUB_CMDLINE_LINUX` in `/etc/default/grub`.

For example:

```
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="vconsole.font=latacyrheb-sun16 vconsole.keymap=us rd.lvm
                    elevator=deadline crashkernel=auto rhgb quiet"
GRUB_DISABLE_RECOVERY="true"
```

and then rebuild the `/boot/grub2/grub.cfg` file to reflect the changes:

- On BIOS-based machines:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```
- On UEFI-based machines:

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

A setting for the elevator parameter is always included by SUSE in its LILO and GRUB configuration files. In this case, change the parameter from `elevator=cfq` to `elevator=deadline`.

Reboot the system once the appropriate file has been modified.

See the Linux operating system documentation for more information on I/O schedulers.

Recommended limit of number of files in a directory

To maximize VxFS performance, do not exceed 100,000 files in the same directory. Use multiple directories instead.

The vxlist command cannot correctly display numbers greater than or equal to 1 EB

The `vxlist` command and all of the other commands that use the same library as the `vxlist` command cannot correctly display numbers greater than or equal to 1 EB.

Limitations with delayed allocation for extending writes feature

The following limitations apply to the delayed allocation for extending writes feature:

- In the cases where the file data must be written to disk immediately, delayed allocation is disabled on that file. Examples of such cases include Direct I/O, concurrent I/O, FDD/ODM access, and synchronous I/O.
- Delayed allocation is not supported on memory mapped files.
- Delayed allocation is not supported with BSD quotas. When BSD quotas are enabled on a file system, delayed allocation is turned off automatically for that file system.
- Delayed allocation is not supported for shared mounts in a cluster file system.

Compressed files that are backed up using NetBackup 7.1 or prior become uncompressed when you restore the files

The NetBackup 7.1 release and prior does not support the file compression feature. If you back up compressed files using NetBackup 7.1 or a prior release, the files become uncompressed when you restore the files.

On SUSE, creation of a SmartIO cache of VxFS type hangs on Fusion-io device (3200586)

On SUSE, creating a SmartIO cache of VxFS type hangs on Fusion-io devices. This issue is due to a limitation in the Fusion-io driver.

Workaround:

To workaround the issue

- ◆ Limit the maximum I/O size:

```
# vxtune vol_maxio 1024
```

A NetBackup restore operation on VxFS file systems does not work with SmartIO writeback caching

A NetBackup restore operation on VxFS file systems does not work with SmartIO writeback caching.

VxFS file system writeback operation is not supported with volume level replication or array level replication

The VxFS file system writeback operation is not supported with volume level replication or array level replication.

SmartIO software limitations

The following are the SmartIO software limitations in this release.

Cache is not online after a reboot

Generally, the SmartIO cache is automatically brought online after a reboot of the system.

If the SSD driver module is not loaded automatically after the reboot, you need to load the driver and bring the cache disk group online manually.

To bring a cache online after a reboot

- 1 Load the SSD driver module with the `insmod` command.

See the Linux documentation for details.

- 2 Perform a scan of the OS devices:

```
# vxdisk scandisks
```

- 3 Bring the cache online manually:

```
# vxdg import cachedg
```

Writeback caching limitations

In the case of CFS, writeback caching is supported with the cache area created on direct attached storage (DAS) and SAN via a Fibre Channel. The cache area should not be shared between cluster nodes.

Writeback caching is only supported on two-node CFS only.

The `sfcache` operations may display error messages in the caching log when the operation completed successfully (3611158)

The `sfcache` command calls other commands to perform the caching operations. If a command fails, additional commands may be called to complete the operation. For debugging purposes, the caching log includes all of the success messages and failure messages for the commands that are called.

If the `sfcache` command has completed successfully, you can safely ignore the error messages in the log file.

Replication software limitations

These software limitations apply to the following products:

- Veritas InfoScale Storage
- Veritas InfoScale Enterprise

Softlink access and modification times are not replicated on RHEL5 for VFR jobs

When running a file replication job on RHEL5, softlink access and modification times are not replicated.

VVR Replication in a shared environment

Currently, replication support is limited to 8-node cluster applications.

VVR IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

- A replication configuration from an IPv4-only node to an IPv6-only node and from an IPv6-only node to an IPv4-only node is not supported, because the IPv6-only node has no IPv4 address configured on it and therefore VVR cannot establish communication between the two nodes.
- A replication configuration in which an IPv4 address is specified for the `local_host` attribute of a primary RLINK and an IPv6 address is specified for the `remote_host` attribute of the same RLINK.
- A replication configuration in which an IPv6 address is specified for the `local_host` attribute of a primary RLINK and an IPv4 address is specified for the `remote_host` attribute of the same RLINK.
- IPv6 is not supported in a CVM and VVR cluster where some nodes in the cluster are IPv4-only and other nodes in the same cluster are IPv6-only, or all nodes of a cluster are IPv4-only and all nodes of a remote cluster are IPv6-only.
- VVR does not support Edge and NAT-PT routers that facilitate IPv4 and IPv6 address translation.

VVR support for replicating across Storage Foundation versions

VVR supports replication between InfoScale Storage 7.4.2 and the prior major release of InfoScale Storage 7.4.1. Replication between versions is supported for disk group versions 290, 280, and 270. Both the Primary and Secondary hosts must be using a supported disk group version.

Cluster Server software limitations

These software limitations apply to the following products:

- Veritas InfoScale Availability

- Veritas InfoScale Enterprise

Limitations related to bundled agents

GoogleIP service group comes online even though OverlayIP resource is already online outside the cluster

When an IP address that is associated with an OverlayIP resource is already online elsewhere in the VPC network, the GoogleIP service group with that OverlayIP resource should not come online. However, the service group does come online, and then the IP resource faults in both the clusters.

Workaround: Do not configure the IP address that is associated with the OverlayIP resource for any other device within the same VPC network.

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.

Veritas 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.

Share agent limitations

To ensure proper monitoring by the Share agent, verify that the `/var/lib/nfs/etab` file is clear upon system reboot. Clients in the Share agent must be specified as fully qualified host names to ensure seamless failover.

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.

Application agent limitations

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

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.

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 onlineing the resource.

Limitation of VMwareDisks agent to communicate with the vCenter Server [3528649]

If VMHA is not enabled and the host ESX faults then even after the disks are attached to the target virtual machine, they remain attached to the failed virtual machine. This issue occurs because the request to detach the disks fails since the host ESX itself has faulted. The agent then sends the disk attach request to the vCenter Server and attaches the disks to the target virtual machine. Even though the application availability is not impacted, the subsequent restart of the faulted virtual machine fails. This issue occurs because of the stale link between the virtual machine and the disks attached. Even though the disks are now attached to the target virtual machine the stale link with the failed virtual machine still exists.

Workaround: Detach the disks from the failed virtual machine and then restart the virtual machine.

NFSRestart agent: In NFSv3, lock recovery is not supported with multiple NFS share service groups

In NFSv3, lock recovery is not supported with multiple NFS share service groups.

Workaround: Configure a single NFS share service group.

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.

Veritas cluster configuration wizard limitations

Wizard fails to configure VCS resources if storage resources have the same name [3024460]

Naming storage resources like disk group and volumes with the same name is not supported as the High Availability wizard fails to configure the VCS resources correctly.

Workaround: No workaround.

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

By default, the Veritas 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.

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.

Security-Enhanced Linux is not supported on SLES distributions

VCS does not support Security-Enhanced Linux (SELinux) on SLES11. [1056433]

Systems in a cluster must have same system locale setting

VCS does not support clustering of systems with different system locales. All systems in a cluster must be set to the same locale.

VxVM site for the disk group remains detached after node reboot in campus clusters with fire drill [1919317]

When you bring the DiskGroupSnap resource online, the DiskGroupSnap agent detaches the site from the target disk group defined. The DiskGroupSnap agent invokes VCS action entry points to run VxVM commands to detach the site. These commands must be run on the node where the disk group is imported, which is at the primary site.

If you attempt to shut down the node where the fire drill service group or the disk group is online, the node goes to a LEAVING state. The VCS engine attempts to take all the service groups offline on that node and rejects all action entry point

requests. Therefore, the DiskGroupSnap agent cannot invoke the action to reattach the fire drill site to the target disk group. The agent logs a message that the node is in a leaving state and then removes the lock file. The agent's monitor function declares that the resource is offline. After the node restarts, the disk group site still remains detached. [1272012]

Workaround:

You must take the fire drill service group offline using the `hagrp -offline` command before you shut down the node or before you stop VCS locally.

If the node has restarted, you must manually reattach the fire drill site to the disk group that is imported at the primary site.

If the secondary node has crashed or restarted, you must manually reattach the fire drill site to the target disk group that is imported at the primary site using the following command: `/opt/VRTSvcs/bin/hares -action $targetres joindg -actionargs $fdsitename $is_fenced -sys $targetsys.`

Limitations with DiskGroupSnap agent [1919329]

The DiskGroupSnap agent has the following limitations:

- The DiskGroupSnap agent does not support layered volumes.
- If you use the Bronze configuration for the DiskGroupSnap resource, you could end up with inconsistent data at the secondary site in the following cases:
 - After the fire drill service group is brought online, a disaster occurs at the primary site during the fire drill.
 - After the fire drill service group is taken offline, a disaster occurs at the primary while the disks at the secondary are resynchronizing.

Veritas recommends that you use the Gold configuration for the DiskGroupSnap resource.

System reboot after panic

If the VCS kernel module issues a system panic, a system reboot is required [293447]. The supported Linux kernels do not automatically halt (CPU) processing. Set the Linux "panic" kernel parameter to a value other than zero to forcibly reboot the system. Append the following two lines at the end of the `/etc/sysctl.conf` file:

```
# force a reboot after 60 seconds
kernel.panic = 60
```

Host on RHEV-M and actual host must match [2827219]

You must configure the host in RHEV-M with the same name as in the `hostname` command on a particular host. This is mandatory for RHEV Manager to be able to search the host by hostname.

Cluster Manager (Java console) limitations

This section covers the software limitations for Cluster Manager (Java Console).

VCS Simulator does not support I/O fencing

When running the Simulator, be sure the `UseFence` attribute is set to the default, “None”.

Using the KDE desktop

Some menus and dialog boxes on Cluster Manager (Java Console) may appear misaligned or incorrectly sized on a KDE desktop. To ensure the proper appearance and functionality of the console on a KDE desktop, use the Sawfish window manager. You must explicitly select the Sawfish window manager even if it is supposed to appear as the default window manager on a KDE desktop.

Limitations related to LLT

This section covers LLT-related software limitations.

Limitation related to LLT on RHEL [3983163]

Limitation for manual configuration of LLT on RHEL 8

On RHEL 8.0 and later systems, network services does not come up after a system restart. This issue occurs because RHEL 8 mandates that the NetworkManager service is used to control the networking interfaces. If the **NM_CONTROLLED** parameter is set to **no** for a network device, the device is no longer controlled by the NetworkManager. Therefore, the device does not come up automatically after a system restart.

Workaround:

On RHEL 8 and later systems, use the `nmcli` or the `nmtui` tools to configure network interfaces. If you configure LLT manually, ensure that you set the `NM_CONTROLLED` parameter to **yes** in the interface files (`ifcfg`) in the `/etc/sysconfig/network-scripts/` directory.

Limitation of LLT support over UDP or RDMA using alias IP [3622175]

When configuring the VCS cluster, if alias IP addresses are configured on the LLT links as the IP addresses for LLT over UDP or RDMA, LLT may not work properly.

Workaround: Do not use alias IP addresses over UDP or RDMA.

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 *Cluster Server Administrator's Guide* for a description of the problems a failed interconnect can create and the protection I/O fencing provides.

In a cluster using SCSI-3 based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on both the data disks and coordinator disks. In a cluster using CP server-based fencing, I/O fencing implements data protection by placing the SCSI-3 PR keys on data disks and similar registrations on CP server. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordination points and data disks to prevent possible difficulties with subsequent cluster startup.

Using the reboot command rather than the shutdown command bypasses shutdown scripts and can leave keys on the coordination points and data disks. Depending on the order of reboot and subsequent startup events, the cluster may warn of a possible split brain condition and fail to start up.

Workaround: Use the shutdown -r command on one node at a time and wait for each node to complete shutdown.

Uninstalling VRTSvxvm causes issues when VxFEN is configured in SCSI3 mode with dmp disk policy (2522069)

When VxFEN is configured in SCSI3 mode with dmp disk policy, the DMP nodes for the coordinator disks can be accessed during system shutdown or fencing arbitration. After uninstalling VRTSvxvm RPM, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm RPM 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 state, VCS triggers a fencing race to avoid data corruption. However, the restarted node loses the fencing race and panics.

Workaround: No workaround.

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.

Clusters must run on VCS 6.0.5 and later to be able to communicate after upgrading to 2048 bit key and SHA256 signature certificates [3812313]

In global clusters, when you install or upgrade VCS to 7.4.3 and you upgrade to 2048 bit key and SHA256 signature certificates on one site and the other site is on VCS version lower than 6.0.5, the clusters fail to communicate. The cluster communication will not be restored even if you restore the trust between the clusters. This includes GCO, Steward and CP server communication.

Workaround: You must upgrade VCS to version 6.0.5 or later to enable the global clusters to communicate.

Storage Foundation Cluster File System High Availability software limitations

These software limitations apply to the following products:

- Veritas InfoScale Storage
- Veritas InfoScale Enterprise

cfsmntadm command does not verify the mount options (2078634)

You must confirm the mount options are correct which are then passed to the `cfsmntadm` command. If the mount options are not correct, the mount fails and the CFSMount resource will not come online. You can check the VCS engine log file for any mount failure messages.

Obtaining information about mounted file system states (1764098)

For accurate information about the state of mounted file systems on Linux, refer to the contents of `/proc/mounts`. The `mount` command may or may not reference this source of information depending on whether the regular `/etc/mtab` file has been replaced with a symbolic link to `/proc/mounts`. This change is made at the discretion of the system administrator and the benefits are discussed in the mount online

manual page. A benefit of using `/proc/mounts` is that changes to SFCFSHA mount options are accurately displayed for all nodes.

Stale SCSI-3 PR keys remain on disk after stopping the cluster and deporting the disk group

When all nodes present in the Veritas InfoScale cluster are removed from the cluster, the SCSI-3 Persistent Reservation (PR) keys on the data disks may not get preempted. As a result, the keys may be seen on the disks after stopping the cluster or after the nodes have booted up. The residual keys do not impact data disk fencing as they will be reused or replaced when the nodes rejoin the cluster. Alternatively, the keys can be cleared manually by running the `vxfcntlclearpre` utility.

For more information on the `vxfcntlclearpre` utility, see the *Veritas InfoScale Administrator's Guide*.

Unsupported FSS scenarios

The following scenario is not supported with Flexible Storage Sharing (FSS):
Veritas NetBackup backup with FSS disk groups

Storage Foundation for Oracle RAC software limitations

These software limitations apply to Veritas InfoScale Enterprise.

Supportability constraints for normal or high redundancy ASM disk groups with CVM I/O shipping and FSS (3600155)

Normal or high redundancy ASM disk groups are not supported in FSS environments or if CVM I/O shipping is enabled.

Configure ASM disk groups with external redundancy in these scenarios.

Limitations of CSSD agent

The limitations of the CSSD agent are as follows:

- The CSSD agent restarts Oracle Grid Infrastructure processes that you may manually or selectively take offline outside of VCS.

Workaround: First stop the CSSD agent if operations require you to manually take the processes offline outside of VCS.

For more information, see the topic "Disabling monitoring of Oracle Grid Infrastructure processes temporarily" in the *Storage Foundation for Oracle RAC Configuration and Upgrade Guide*.

- The CSSD agent detects intentional offline only when you stop Oracle Clusterware/Grid Infrastructure outside of VCS using the following command: `crsctl stop crs [-f]`. The agent fails to detect intentional offline if you stop Oracle Clusterware/Grid Infrastructure using any other command.

Workaround: Use the `crsctl stop crs [-f]` command to stop Oracle Clusterware/Grid Infrastructure outside of VCS.

Oracle Clusterware/Grid Infrastructure installation fails if the cluster name exceeds 14 characters

Setting the cluster name to a value that exceeds 14 characters during the installation of Oracle Clusterware/Grid Infrastructure causes unexpected cluster membership issues. As a result, the installation may fail.

Workaround: Restart the Oracle Clusterware/Grid Infrastructure installation and set the cluster name to a value of maximum 14 characters.

Policy-managed databases not supported by CRSResource agent

The CRSResource agent supports only admin-managed database environments in this release. Policy-managed databases are not supported.

Health checks may fail on clusters that have more than 10 nodes

If there are more than 10 nodes in a cluster, the health check may fail with the following error:

```
vxgettext ERROR V-33-1000-10038
Arguments exceed the maximum limit of 10
```

The health check script uses the `vxgettext` command, which does not support more than 10 arguments.[2142234]

Cached ODM not supported in Veritas InfoScale environments

Cached ODM is not supported for files on Veritas local file systems and on Cluster File System.

Storage Foundation for Databases (SFDB) tools software limitations

The following are the SFDB tools software limitations in this release.

Parallel execution of `vxsfadm` is not supported (2515442)

Only one instance of the `vxsfadm` command can be run at a time. Running multiple instances of `vxsfadm` at a time is not supported.

Creating point-in-time copies during database structural changes is not supported (2496178)

SFDB tools do not support creating point-in-time copies while structural changes to the database are in progress, such as adding or dropping tablespaces and adding or dropping data files.

However, once a point-in-time copy is taken, you can create a clone at any time, regardless of the status of the database.

Oracle Data Guard in an Oracle RAC environment

SFDB tools cannot be used with RAC standby databases. SFDB tools can still be used with the primary database, even in a Data Guard Oracle RAC environment.

Known issues

This chapter includes the following topics:

- [Documentation errata](#)
- [Issues related to installation and upgrade](#)
- [Storage Foundation known issues](#)
- [Replication known issues](#)
- [Cluster Server known issues](#)
- [Storage Foundation and High Availability known issues](#)
- [Storage Foundation Cluster File System High Availability known issues](#)
- [Storage Foundation for Oracle RAC known issues](#)
- [Storage Foundation for Databases \(SFDB\) tools known issues](#)
- [Application isolation feature known Issues](#)
- [Cloud deployment known issues](#)
- [Container deployment know issues](#)

Documentation errata

The following issues exist with the documentation for this release:

- The *Storage Foundation for Oracle RAC 7.4.3 Configuration and Upgrade Guide - Linux* retains legacy content about upgrade scenarios and processes. However, only fresh installations of InfoScale 7.4.3 with Oracle RAC configurations are allowed. Upgrades from Oracle RAC configurations with earlier InfoScale versions to 7.4.3 are not allowed. Ignore the upgrade-related content from that document.

- You may find some content that is applicable to platforms other than Linux. Ignore such content for InfoScale 7.4.3, because this release only supports the Linux platform.

Issues related to installation and upgrade

This section describes the known issues during installation and upgrade. These known issues apply to the following products:

- Veritas InfoScale Foundation
- Veritas InfoScale Storage
- Veritas InfoScale Availability
- Veritas InfoScale Enterprise

During upgrade from InfoScale 7.4.1 to 7.4.3, xprtld fails to start without node restart (4017516)

When you upgrade an InfoScale product by using the product installer, the xprtld service starts automatically along with the product stack. However, when you upgrade an InfoScale 7.4.1 product to the corresponding 7.4.3 version, the xprtld service may fail to start. The product installer may prompt you to start the service on specific cluster nodes by using the `/sbin/shutdown -r` command.

After you follow the instructions in the prompt, the xprtld service starts without the need for any further action.

Workaround: No workaround

Upgrade from SFRAC 6.2.1 to InfoScale 7.4.2 using yum may fail (4002582)

When you perform an upgrade from SFRAC 6.2.1 to InfoScale 7.4.2 using the yum utility, the upgrade may fail and display the following error message:

```
%preun(VRTSveki-6.2.0.000-RHEL7.x86_64) scriptlet failed, exit status 1
Error in PREUN scriptlet in rpm package VRTSveki-6.2.0.000-RHEL7.x86_64
```

This issue occurs because the old VRTSveki package does not support the yum upgrade utility. It tries to unload the new vek kernel module while executing the pre-uninstall script, which fails because the other dependent InfoScale kernel modules are already loaded. Eventually, the old VRTSveki package fails to be uninstalled.

Workaround:

When upgrading from SFRAC 6.2.1 to InfoScale 7.4.2, perform the following steps using the yum utility:

- 1 Upgrade VRTSveki package using the yum utility.

```
# yum upgrade VRTSveki
```

- 2 Upgrade other InfoScale packages.

```
# yum upgrade VRTS*
```

Installer fails on SLES 11 and SLES 12 for older NTP versions (3912493)

The Veritas product installer fails to install the packages on SLES 11 and SLES 12 if the NTP RPM version is `ntp-4.2.8p4-1.3` or earlier.

Workaround: Install NTP version `ntp-4.2.8p6-8.2` or later.

InfoScale deployments fail when secure boot is enabled (4005536)

InfoScale deployment is not supported when secure boot is enabled. Kernel components failed to load when the operating system is booted in the secure mode.

Switch fencing in enable or disable mode may not take effect if VCS is not reconfigured [3798127]

When you choose not to reconfigure Veritas Cluster Server (VCS), and set the fencing in enable or disable mode, it may not take effect. This is because the fencing mode switch relies on VCS reconfiguration.

Workaround: If you want to switch the fencing mode, when the installer shows "Do you want to re-configure VCS?", enter `y` to reconfigure VCS .

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**.

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

```
# hagrps -list Frozen=1
```

- 2 Unfreeze all the frozen service groups:

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

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

If you have NetBackup 6.5 or older version installed on a VxFS file system and before upgrading to InfoScale Foundation 7.4.3, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (`/usr/opensv`), then while upgrading to SF 7.4.3, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure RPMs `VRTSspbx`, `VRTSat`, and `VRTSiscsi`. This causes NetBackup to stop working.

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

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

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

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

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

Error messages in syslog (1630188)

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

```
Jul  6 10:58:50 swlx62 setroubleshoot: SELinux is preventing the  
semanage from using potentially mislabeled files  
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For  
complete SELinux messages. run sealert -l ed8978dl-0b1b-4c5b-a086-  
67da2a651fb3  
Jul  6 10:58:54 swlx62 setroubleshoot: SELinux is preventing the  
semanage from using potentially mislabeled files  
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For  
complete SELinux messages. run sealert -l ed8978dl-0b1b-4c5b-a086-  
67da2a651fb3  
Jul  6 10:58:59 swlx62 setroubleshoot: SELinux is preventing the  
restorecon from using potentially mislabeled files
```

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

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

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

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


Or

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

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

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

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

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

Dependency may get overruled when uninstalling multiple RPMs in a single command [3563254]

When performing uninstallation of multiple RPMs through a single comment, the system identifies and follows the specified dependency among the RPMs as the uninstallation progresses. However, if the pre-uninstallation script fails for any of the RPMs, the system does not abort the task but instead uninstalls the remaining RPMs.

For example, if you run `rpm -e VRTS11lt VRTSgab VRTSvxfen` where the RPMs have a dependency between each other, the system bypasses the dependency if the pre-uninstallation script fails for any RPM.

Workaround: Uninstall the RPMs independently.

When using the response file, the installer must not proceed with the installation or upgrade, if you have not provided edge server details (3964335)

When installing or upgrading using the response file, configuration takes place even if you have not provided details for `edgeserver_host` and `edgeserver_port`. This

causes the configuration to fail, without giving the user a chance to provide the required details. The following error message is displayed:

```
CPI ERROR Answer for 'Enter the edge server's hostname/ip:' is
required with responsefile or silent mode
```

Workaround:

You must manually configure the edge server using the following command:

```
/opt/VRTSvlic/tele/bin/TelemetryCollector -update
--hostname=telemetry.veritas.com --port=443
```

Note that **telemetry.veritas.com** and **443** are default values for the host name and port number of the Veritas Cloud Receiver.

Unable to update edge server details by running the installer (3964611)

If you provide incorrect details for an edge server during an installation or upgrade, you are unable to update the details of the edge server by running the installer again.

Workaround:

You can update the edge server details by entering the following command:

```
/opt/VRTSvlic/tele/bin/TelemetryCollector -update
--hostname=telemetry.veritas.com --port=443
```

Note that **telemetry.veritas.com** and **443** are default values for the host name and port number of the Veritas Cloud Receiver.

Storage Foundation known issues

This section describes the known issues in this release of Storage Foundation (SF). These known issues apply to the following products:

- Veritas InfoScale Foundation
- Veritas InfoScale Storage
- Veritas InfoScale Enterprise

Dynamic Multi-Pathing known issues

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

kdump functionality does not work when DMP Native Support is enabled on Linux platform [3754715]

The issue occurs because of filters which are required for Dynamic Multi-pathing (DMP) Native support to work. For working of DMP Native Support, we reject all the devices in LVM filters except /dev/vx/dmp. This means that the kdump device is also excluded. The DMP devices are not present as part of initramfs at boot and hence kdump is not able to capture the crash dump of the system.

Workaround: There are two ways to solve this issues.

- Workaround 1:

1. Copy vxvm lvm.conf.

```
# cp /etc/lvm/lvm.conf /etc/lvm/lvm.conf.vxvm
```

2. Copy original lvm.conf back.

```
# cp /etc/lvm/lvm.conf.orig /etc/lvm/lvm.conf
```

3. Remove kdump initrd.

```
# rm -rf /boot/initrd-2.6.32-504.el6.x86_64kdump.img
```

4. Restart kdump.

```
# service kdump restart
```

5. Copy VxVM lvm.conf back.

```
# cp /etc/lvm/lvm.conf.vxvm /etc/lvm/lvm.conf
```

The drawback for this workaround is that you have to perform these steps every time after reboot and whenever the kdump initrd is re-generated.

- Workaround 2:

Add the filter for the dump device in accept section in lvm.conf file. But here you need to make sure that DUMP device is *not configured* on root device i.e “/”. On the system, if the dump device is configured on top of root and if we accept the root device, then Root LVM will not come under DMP and Root LVM will be monitored by Native Multipathing and not DMP.

On SLES machine, after you enable the switch ports, some paths may not get enabled automatically [3782724]

If disabled host-side switch ports are enabled without running Veritas Volume Manager (VxVM) device discovery by using `vxdisk scandisks` OR `vxctl enable` command in between, some paths may not get enabled automatically.

Workaround: Run the # `vxddmpadm enable path=<path_name>` command to enable the paths which were not automatically enabled.

Veritas Volume Manager known issues

This section lists the known issues that are related to the Veritas Volume Manager (VxVM) component of InfoScale.

Issues with host prefix values in case of NVME disks (4017022)

As NVME devices are local disks, their names are prepended with host prefix value by default.

NVME devices appear in the error state in the output of the `vxdisk list` command if the following conditions are satisfied. A host prefix is set to an empty value, and then the `vxddladm -c assign names` command is run to clear the previous names and autogenerate the names again.

Note: This issue does not occur if a host prefix is unset by using the `vxctl unset hostprefix` command.

Workaround: Either unset host prefix or set it to a non-empty value, so that the hostname string can be set as the host prefix.

If the hostname or the host prefix starts with "nvme", all the disks that are claimed by the JBOD and ASLs of the local connection type may also appear in the error state.

Workaround: No workaround

vxsnap prepare may fail in case of a volume set (3993242)

The `vxsnap prepare` command may fail when run for a volume set (vset), and it may display the following error: "VxVM vxsnap ERROR V-5-1-19171 Cannot perform prepare operation on cloud volume."

Workaround

Run the `vxsnap prepare` command separately for each individual volume in the volume set, instead of running it for the entire volume set.

vradmin delsec fails to remove a secondary RVG from its RDS (3983296)

This error occurs if the `vradmin addsec` command does not reset the value of the `mConfigStatus` attribute to **0** (zero) after adding a secondary RVG to an RDS.

Workaround:

Restart the `vradmind` daemon.

FSS disk group creation fails for clusters with eight or more nodes that have several directly attached disks (3986110)

The creation of an FSS disk group or the addition of a disk to an existing FSS disk group fails and logs the following error:

```
VxVM vxdg ERROR V-5-1-10127 associating disk-media smicro101_exosx100_0 with
Slave failed to create remote disk
```

This issue occurs when a cluster has eight or more nodes and several directly attached disks. In case of such a configuration, a race condition occurs during operations like disk group creation or disk addition to an existing group. The race condition deletes the disk records from the kernel and consequently fails to add disks to a disk group.

Workaround:

In a cluster with eight or more nodes and several directly attached disks, create the disk group with a few disks at a time instead of with all the disks in one go.

Kernel-level warnings may appear in system logs during direct I/O to underlying device (3990118, 3998171)

During direct I/O operations on an underlying device, kernel-level warnings like the following may appear in the system logs:

```
"no blkcg associated for bio on block-device"
"(blkcg_release) <= 0 after switching to atomic at lib/percpu-refcount.c"
```

Such one-time warnings may be logged in rare cases, and they occur due to a limitation in the way InfoScale interacts with the kernel. You can safely ignore these warnings.

Workaround:

None

Multiple Issues with Root Disk Encapsulation on RHEL

Multiple issues have been observed with root disk encapsulation in Veritas Infoscale Storage 7.0 and later versions. The issues have been listed as known issues in the Release Notes with the following tracking IDs:

- Servers with UEFI firmware (1842096)
See [“Machine fails to boot after root disk encapsulation on servers with UEFI firmware \(1842096\)”](#) on page 66.
- `vxunroot` cannot encapsulate a root disk when the root partition has XFS mounted on it (3614362)
See [“vxunroot cannot encapsulate a root disk when the root partition has XFS mounted on it \(3614362\)”](#) on page 71.
- `Device.map` must be up to date before doing root disk encapsulation (3761585, 2202047)
See [“device.map must be up to date before doing root disk encapsulation \(2202047\)”](#) on page 67.
- XFS file system is not supported for root disk encapsulation.
See [“XFS file system is not supported for RDE”](#) on page 79.

For details, refer to the TechNote at:

https://www.veritas.com/content/support/en_US/article.100033121

Some of the above issues are related to boot interface (grub2, UEFI) updates while other issues may be related to the operating system (RHEL Bug ID: 1399517). Veritas is working with the OS vendors to resolve the issues and recommends that root disk encapsulation not be used until the issues are resolved.

Workaround : None.

Core dump issue after restoration of disk group backup (3909046)

After you restore a disk group backup using the `vxconfigrestore` command, it is possible that some configuration copies remain in a disabled state. As a result, VxVM generates a core dump when you view the list of disk groups after the restore operation.

Stack Trace:

```
#0  0x00000033a3432625 in raise () from /lib64/libc.so.6
#1  0x00000033a3433e05 in abort () from /lib64/libc.so.6
#2  0x00000033a342b74e in __assert_fail_base () from /lib64/libc.so.6
```

```
#3 0x00000033a342b810 in __assert_fail () from /lib64/libc.so.6
#4 0x00000000005060f1 in req_dg_get_info_common (clnt=0x1af1750,
dg=0x7fc330004bb0) at
dg.c:3261
#5 0x00000000005059f5 in req_dg_get_info_name (clnt=0x1af1750,
req=0x1b03f78)
at dg.c:3057
#6 0x00000000005b2e2d in vold_process_request (arg=0x18f5f20) at
request.c:1997
#7 0x00000033a3807a51 in start_thread () from /lib64/libpthread.so.0
#8 0x00000033a34e896d in clone () from /lib64/libc.so.6
```

Workaround: Restart the VxVM configuration daemon.

```
# vxconfigd -kr reset
```

VxVM tunables not updated on SLES 12 SP2 systems with 4.4 kernel (3916902)

The VxVM tunables are not updated on SLES 12 SP2 systems (with 4.4 kernel) even after rebooting the systems. This is because changes in the 4.4 kernel API cause a failure in reading the VxVM tunable file.

Workaround: There is no workaround for the issue.

Failed verifydata operation leaves residual cache objects that cannot be removed (3370667)

When you use the verify data command, and type

```
# vradmin -g dgname verifydata rvgname IPaddress cachesize=size
```

the command may fail and leave residual cache objects that cannot be removed.

Workaround:

To solve this problem, choose different ways based on different residual cache objects.

To explicitly clean up the cache object that is associated to SO snapshots:

1. List the SO snapshots that are created on a cache object by typing:

```
# vxcache -g dgname listvol volumename
```

2. Unmount the listed snapshots.
3. Remove the snapshot volume. Type:

```
# vxedit -g dname -fr rm volumename
```

It also removes the cache object.

To clean up the cache object that is not associated to the snapshot volume but associated to the cache volume:

1. Stop the cache object by typing:

```
# vxcache -g dname stop cacheobject_name
```

2. Remove the cache object. Type:

```
# vxedit -g dname -rf rm cacheobject_name
```

It also removes the cache volume.

LUNs claimed but not in use by VxVM may report “Device Busy” when it is accessed outside VxVM (3667574)

When a LUN claimed by Veritas Volume Manager (VxVM) is accessed, the open on the device gets cached for performance improvement. Due to this, some OS utilities which require exclusive access reports `Device Busy`.

Workaround:

To solve this issue, either exclude these LUNs from the VxVM view or disable them by typing `vxddmpadm disable dmpnodename=<> CLI`.

For details, refer to the TechNote at:

https://www.veritas.com/content/support/en_US/article.100014895

If the disk with CDS EFI label is used as remote disk on the cluster node, restarting the vxconfigd daemon on that particular node causes vxconfigd to go into disabled state (3873123)

When you restart the `vxconfigd` daemon, or run the `vxctl enable` command, you may encounter this error:

```
VxVM vxctl ERROR V-5-1-1589 enable failed: Error in disk group
configuration copies
```

This is because one of the cases for EFI remote disk is not properly handled in the disk recovery part when you enable the `vxconfigd` daemon.

Workaround:

To solve this issue, follow the steps:

- 1 Take the node on which issue is seen out of cluster by running proper VCS command to stop the node.
- 2 Enable the `vxconfigd` daemon by running:

```
# vxctl enable
```

- 3 Restart the node by running proper VCS command.

Unable to set master on the secondary site in VVR environment if any pending I/O's are on the secondary site (3874873)

There is deadlock situation with the cluster reconfiguration and the network disconnection (serialization) on RVG object. Wherein, the reconfiguration quiesces the disk level I/O's and it expects the replica object to be disconnected. The Rlink cannot be disconnected unless the underlying I/O's are completed and the reconfig thread quiesces these I/Os at disk level.

Workaround:

Pause the Rlink on the primary site and then set master on the secondary slave node.

After installing DMP 6.0.1 on a host with the root disk under LVM on a cciss controller, the system is unable to boot using the `vxdump_kernel` command [3599030]

The Dynamic Multi-Pathing (DMP) Native support feature is not supported for the COMPAQ SMART controllers which use device names of the form `/dev/cciss/cXdxpX`. When the `dmp_native_support` feature is enabled, it creates a new `initrd` image with a Logical Volume Manager (LVM) filter in `lvm.conf.filter=["a|/dev/vx/dmp/.*/", "r|.*/"]`. The filter only allows access to devices under `/dev/vx/dmp`. But `/dev/vx/dmp/cciss`, where the root disks DMP nodes are located, are not allowed.

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.

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.

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

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

Certain new servers in the market such as IBM x3650 M2, Dell PowerEdge T610, come with support for the UEFI firmware. UEFI supports booting from legacy MBR type disks with certain restrictions on the disk partitions. One of the restrictions is that each partition must not overlap with other partitions. During root disk encapsulation, it creates an overlapping partition that spans the public region of

the root disk. If the check for overlapping partitions is not disabled from the UEFI firmware, then the machine fails to come up following the reboot initiated after running the commands to encapsulate the root disk.

Workaround:

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

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

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

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

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

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

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

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

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
```

VxVM starts before OS device scan is done (1635274)

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

Workaround:

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

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

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

Workaround:

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

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

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

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

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

Workaround:

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

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

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

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

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:

```
# vxddmpadm setattr enclosure enc11 recoveryoption=throttle \
  iotimeout=600
```

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

```
# vxddctl enable
```

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

If a CLARiiON LUN is mapped to a different host than its snapshot, a trespass on one of them could cause a trespass on the other. This behavior could result in a loop for these LUNs, as DMP tries to fail back the LUNs if the primary paths are available.

Workaround

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

```
# vxddm padm settune dmp_monitor_ownership=off
```

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 environment which includes devices that are managed by EMC PowerPath, a storage loss causes Veritas Volume Manager commands to be delayed. In the event of storage loss, VxVM sends SCSI inquiry 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 dgroup
```

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.

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
```

vxunroot cannot encapsulate a root disk when the root partition has XFS mounted on it (3614362)

If the root partition has the XFS file system mounted on it, you cannot change the root partition's Universally Unique Identifier (UUID). However, changing the UUID of the partitions of the root disk is necessary in root disk encapsulation. Given the limitation above, Veritas does not support root disk encapsulation where the root partition has an XFS file system.

Workaround:

None.

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 Storage Sharing (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.

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

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

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

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.

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

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

Workaround: There is no workaround for this issue.

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.

Systems may panic after GPT disk resize operation (3930664)

After you resize the GPT disks using the following command, you may experience a system panic issue, # `vxdisk resize <disk_name> length=<new_size>`. This issue occurs if your deployment setup includes GPT disks partition.

No workaround available to resolve this issue. So you must not resize GPT disks, and to recover the system, wait for the system to restart.

If LVM volume group has mirror volume, the conversion operation to VxVM fails (3930536)

The `vxvmconvert` utility enables you to convert the Logical Volume Manager (LVM) volume group to Veritas Volume Manager (VxVM) disk groups. This conversion may fail, if the LVM volume group has mirrored volumes.

Currently , there is no workaround available to resolve the issue.

If recovery of columns on EC volumes fails, recovery of other columns on the other volumes also fails (3930435)

If some faults are identified during resize operation on erasure coded volume, some columns of erasure coded volumes are detached. Once these faults are resolved, the columns are reattached to erasure coded volume and recovery operation is triggered. The recovery operation may fail with the `DCO experienced IO errors during the operation` error. The error occurs when the DCO has stale maps and cannot allocate new maps for the recovery operation. This failure may also result in failure of recovery on columns on other non-erasure coded volumes.

Workaround

Add maps in DCO using the following command and trigger the recovery manually:

```
# vxsnap -g <dg> addmap <volume> <no_of_maps>
```

Restarting vxconfigd during relay layout operation causes the volume to go in an intermediate state.(3959429)

Relayout operation is performed in multiple iterations and each iteration involves transaction. If `vxconfigd` is down or restarted during the relay layout operation, the transaction fails. This results into relay layout failure and the volumes undergoing relay layout goes into an intermediate state.

Workaround

When vxconfigd is up and running on all cluster nodes, restart the relay layout operation:

```
# vxrelayout -g <dgname> start <volname/vsetname>
```

Veritas File System known issues

This section lists the known issues that are related to the Veritas File System (VxFS) component of InfoScale.

On an SELinux-enabled RHEL 7.7 or later system with DLV 10 or earlier, the mount operation on the filesystem fails after the upgrade (3992626)

This issue occurs because DLV 10 or earlier cannot store the SELinux extended attributes, and therefore, VxFS cannot perform the local mount after the upgrade. All the file systems that are created on VxFS 6.x or VxFS 7.0 use DLV 10 or earlier. The SELinux policy for RHEL 7.7 and later supports VxFS as a persistent storage for the SELinux extended attributes. DLV 11 and later can store the `security.selinux` extended attributes, however, they are only supported with VxFS 7.1 and later.

Workaround:

Before upgrading VxFS on RHEL 7.7 and later systems, disable SELinux on VxFS 6.x or 7.0. After upgrading the filesystem on a node to VxFS 7.4.2, use the `vxupgrade` command to upgrade an earlier VxFS disk layout to DLV 12 or later. Unmount the disk and re-enable SELinux. The node can then proceed for cluster mount normally.

On CFS, if delayed allocation and delayed extending write are both enabled on an Inode, data behaviour becomes unpredictable on that Inode (3982121)

This issue occurs if the delayed allocation and the delayed extending writes features are enabled at the same time. In such a scenario, a race condition occurs between multiple threads working on the same Inode. The race condition might cause the data on that specific Inode to behave unpredictably.

Workaround:

Ensure that only one among the delayed allocation and the delayed extending writes features is enabled at a time.

Cluster may hang if CFS is FCL-enabled and its DLV is greater than or equal to 14 (4002222)

VxFS worker threads manage the operations that are related to the File Change Log (FCL) feature. The worker threads can get stuck in a deadlock if the disk layout version of an FCL-enabled cluster file system is greater than or equal to 14. If such a deadlock occurs, the cluster may become unresponsive.

Workaround: None.

The VxFS file system with local scope enabled may hang if two or more nodes are restarted simultaneously (3944891)

In a scenario where the VxFS file system has local scope enabled, restarting two or more nodes simultaneously causes a deadlock which leads the file system to hang.

Workaround:

Reboot all the nodes.

Docker does not recognize VxFS backend file system

When VxFS is used as backing filesystem to run the docker daemon, the following error is displayed:

```
Backing Filesystem: unknown
```

The link for this issues in Github is: <https://github.com/docker/docker/issues/14847>

Workaround:

VxFS is recognized as backing filesystem in the Docker upstream.

On RHEL7 onwards, Pluggable Authentication Modules(PAM) related error messages for Samba daemon might occur in system logs [3765921]

After adding Common Internet File System(CIFS) share and the CIFS share is might not be accessible from windows client, the PAM related error messages for Samba daemon might occur.

This issue occurred because the `/etc/pam.d/samba` file is not available by default on RHEL 7 onwards and the `obey_pam_restrictions` attribute from `smb.conf` file, which is Samba configuration file, is set to `yes`, where default is `no`. This parameter controls whether or not Samba should obey PAM's account and session management directives. The default behavior is to use PAM for clear text

authentication only and to ignore any account or session management. Samba always ignores PAM for authentication in the case of `encrypt passwords = yes`.

Workaround: Set `obey pam restrictions = no` in the `/opt/VRTSvcs/bin/ApplicationNone/smb.conf` file before configuring `cfsshare` and adding share.

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.

XFS file system is not supported for RDE

The Root Disk Encapsulation (RDE) feature is not supported if the root partition is mounted with XFS file system.

Workaround: There is no workaround available.

The command tab auto-complete fails for the `/dev/vx/` file tree; specifically for RHEL 7 (3602082)

The command tab auto-complete operation fails because the following RPM is installed on the machine:

```
"bash-completion-2.1-6.el7.noarch"
```

This somehow overwrites the default auto-complete rules. As a result, some issues are observed with the VxFS commands. However, the issue is not observed with all the VxFS commands. The issue is observed with the `mkfs(1M)` command, but is not observed with the `mount(1M)` command.

Workaround: Please remove the "bash-completion-2.1-6.el7.noarch" RPM, so that the command tab auto-complete does not fail for the `/dev/vx/` file tree.

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.

During a deduplication operation, the spoold script fails to start (3196423)

This issue occurs because a port is not available during the operation; therefore the spoold script fails to start with the the following error:

```
DEDUP_ERROR INIT: exec spoold failed (1024)
```

Workaround:

Check the `spoold.log` file for specific error messages, and if the log indicates a port is not available, you can check if the port is in use with the `netstat/lsof` command. If the port is not open, you can retry the deduplication operation; if the port is open, you can wait for the port to close, and then try the deduplication operation again.

For example, the following error message in the `spoold.log` file indicates that port 51003 is not available:

```
ERR [140399091685152]: -1: NetSetup: NetBindAndListen returned error,
unable to bind to localhost:51003
```

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.

Virtualization known issues

This section describes the virtualization known issues in this release.

Configuring application for high availability with storage using VCS wizard may fail on a VMware virtual machine which is configured with more than two storage controllers [3640956]

Application configuration from VCS wizard may fail on VMware virtual machine which is configured with multiple SCSI controllers.

Workaround: There is no workaround available.

Host fails to reboot when the resource gets stuck in ONLINE|STATE UNKNOWN state [2738864]

In a Red Hat Enterprise Virtualization environment, if a host reboot is performed on which the KVMGuest resource monitoring the virtual machine is ONLINE, then the host reboot fails. This is because the VDSM is getting stopped before VCS could shutdown the virtual machine. In this case, the virtual machine state remains ONLINE|STATE UNKNOWN, and hence VCS stop fails eventually failing the host reboot as well.

Workaround: Switch the service group to other node before initiating a host reboot.

VM state is in PAUSED state when storage domain is inactive [2747163]

If the storage domain associated with the running virtual machine becomes inactive, the virtual machine may go to **paused** state.

Workaround: Make sure that the storage domain is always active when running virtual machine.

Switching KVMGuest resource fails due to inadequate swap space on the other host [2753936]

Virtual machine fails to start on a host if the host does not have the sufficient swap space available.

Workaround: Please make sure that each host has sufficient swap space available for starting the virtual machine.

Policies introduced in SLES 11SP2 may block graceful shutdown if a VM in SUSE KVM environment [2792889]

In a SUSE KVM environment, virtual machine running SLES11 SP2 inside may block the virtual machine graceful shutdown request due to some policies introduced in SLES 11SP2. SUSE recommends turning off the policy with `polkit-gnome-authorization` for a virtual machine.

Workaround: Make sure that all the policies blocking any such request are turned off.

Load on libvirtd may terminate it in SUSE KVM environment [2824952]

In a SUSE KVM environment, occasionally `libvirtd` process may get terminated and `/etc/init.d/libvirtd status` command displays:

```
#/etc/init.d/libvirtd status
Checking status of libvirtd                                dead
```

This may be due to heavy load on `libvirtd` process.

Workaround: Restart the `libvirtd` process and run:

```
# service libvirtd stop
# service libvirtd start
```

Offline or switch of KVMGuest resource fails if the VM it is monitoring is undefined [2796817]

In a SUSE KVM environment, if a running virtual machine is undefined using `virsh undefine` command, an attempt to offline or switch the KVM guest resource monitoring that VM fails because the agent is not able to get the information from the KVM hypervisor.

Workaround: To undefine the VM on a particular node, first switch the service group containing the KVMGuest resource to another node and then undefine the VM on the first node.

Increased memory usage observed even with no VM running [2734970]

Increased memory usage was observed on the hosts even when VMs were either not running or had stopped. This is due to the RHEV behavior.

Workaround: No workaround.

Resource faults when it fails to ONLINE VM because of insufficient swap percentage [2827214]

In virtualization environment, if VCS fails to start the virtual machine due to unavailability of required virtualization resources such as CPU, memory, or disks, the resource goes into FAULTED state.

Workaround: Make sure that the required virtualization resources are always available in a virtualization environment.

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

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

Workaround: Start the `libvirtd` process manually.

Virtual machine may return the not-responding state when the storage domain is inactive and the data center is down (2848003)

In a Red Hat Enterprise Virtualization Environment, if the storage domain is in an inactive state and the data center is in down state, the virtual machine may return a not-responding state and the KVMGuest resource in OFFLINE state.

Workaround: To resolve this issue:

- 1 Activate the storage domain in RHEV-M.
- 2 Check that the data center is in the up state.

Guest virtual machine may fail on RHEL 6.1 if KVM guest image resides on CVM-CFS [2659944]

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

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

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

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

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

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

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

CD ROM with empty file vmPayload found inside the guest when resource comes online [3060910]

When you unset the DROpts attribute on a KVMGuest resource and online the resource on the host, a CD ROM with an empty file vmPayload is available inside the guest.

The KVMGuest agent adds a CD ROM to the virtual machine configuration when you online a KVMGuest resource with the DROpts attribute set. The CD ROM carries some site-specific parameters to be used inside the guest. When you offline the same resource, the agent removes the CD ROM, but for some reason, the CD ROM does not get removed completely. If you unset the DROpts attribute and online the resource later, a CD ROM with an empty file vmPayload continues to be available inside the guest.

Workaround: This does not impact the functionality of the virtual machine in any way and can be ignored.

VCS fails to start virtual machine on another node if the first node panics [3042806]

In the KVM environment, if a node on which a virtual machine is running panics, then VCS fails to start that virtual machine on another node. This issue occurs because KVM Hypervisor is not able to acquire lock on the virtual machine. This issue is due to KVM Hypervisor behavior and is very rarely observed.

Workaround: Restart libvirtd process to resolve this issue. Command to restart libvirtd:

```
# service libvirtd restart
```

VM fails to start on the target node if the source node panics or restarts during migration [3042786]

If a virtual machine (VM) migration is initiated and the source node (node on which VM was running) panics or is restarted forcefully, VM fails to start on any other node in a KVM environment. This issue is due to the KVM locking mechanism. The VM start fails with the following error:

```
error: Failed to start domain VM1
error: Timed out during operation: cannot acquire state change lock
```

Workaround: Restart (kill and start) the libvirtd daemon on the second node using the following command:

```
# service libvirtd restart
```

High Availability tab does not report LVMVolumeGroup resources as online [2909417]

The High Availability tab does not automatically report the online status of activated LVMVolumeGroup resources in the following case:

- If you created the VCS cluster as part of the High Availability Configuration Wizard workflow.

Workaround: Start the LVMVolumeGroup resources from the High Availability tab. For more information, see the High Availability Solutions Guide for VMware.

Cluster communication breaks when you revert a snapshot in VMware environment [3409586]

If VCS is running on the guest operating system when a VMware virtual machine snapshot is taken, the virtual machine snapshot contains the run-time state of the cluster. When you restore the snapshot, the state of the cluster which is restored can be inconsistent with other nodes of the cluster. Due to the inconsistent state, VCS is unable to communicate with other nodes of the cluster.

Workaround: Before you take a snapshot of the virtual machine, Veritas recommends that you stop VCS services running inside the virtual machine.

VCS may detect the migration event during the regular monitor cycle due to the timing issue [2827227]

In a virtualization environment, VCS detects the virtual machine migration initiated outside VCS and changes the state accordingly. However, occasionally, VCS may miss the migration event due to timing issue and detect the migration during the regular monitor cycle. For example, if you set OfflineMonitorInterval as 300sec, it takes up to 5 minutes for VCS to report ONLINE on the node where the virtual machine got migrated.

Workaround: No workaround available.

Replication known issues

This section describes the replication known issues in this release of Veritas InfoScale Storage and Veritas InfoScale Enterprise.

The secondary vradmind may appear hung and the vradmind commands may fail (3940842,3944301)

The secondary vradmind daemon may appear to be hung and the vradmind commands may fail, even if the replication is ON.

This issue is observed in cases where heavy I/Os are in progress.

To address the heavy I/Os, the secondary vradmind continues to perform certain IOCTLs and hence fails to respond to the heartbeat from the primary vradmind. As a result, the connection between the two vradmind is lost. This behavior continues during the next connection attempts too, due to which the vradmind commands fail and the secondary vradmind appears hung.

Workaround:

On both the sites, modify the `/etc/vx/vras/vras_env` file to resolve the issue. Make the following edits:

- Uncomment the following commands:

```
#export VRAS_ENABLE_STATS
#VRAS_ENABLE_STATS=on
```

and

```
#export IPM_HEARTBEAT_TIMEOUT
#IPM_HEARTBEAT_TIMEOUT=30
```

- Set the `VRAS_ENABLE_STATS` to OFF
- Increase the `IPM_HEARTBEAT_TIMEOUT` value

After you modify the files, restart the vradmind daemon on both the sites:

```
/usr/sbin/vxstart_vvr stop
```

```
/usr/sbin/vxstart_vvr start
```

Data corruption may occur if you perform a rolling upgrade of InfoScale Storage or InfoScale Enterprise from 7.3.1 or earlier to 7.4 or later during replication (3951527)

A rolling upgrade to InfoScale 7.4 or later is not supported in a Volume Replicator (VVR) environment. With InfoScale 7.3.1 or earlier, you could pause replication, perform a rolling upgrade, and then safely resume replication. However, to upgrade to 7.4 or later, you must first stop any ongoing replication, perform a full upgrade of the product and the disk groups on both the sites, and then start replication.

Workaround:

To upgrade to InfoScale 7.4 or later in a VVR environment

1. Stop replication.
2. Perform a full upgrade on the primary and the secondary sites to the same InfoScale version.
3. Upgrade the disk groups on both the sites.
4. Start replication.

vradmin may appear hung or may fail for the role migrate operation (3968642, 3968641)

While performing the role migration operation, the new primary vradmin daemon may appear to be hung even if the VVR role migration is complete.

However, in certain situations, the vradmin commands may fail with following error message without completing the operation:

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

This issue is observed intermittently.

Workaround:

1. Restart vradmin on all cluster nodes:

```
# /etc/init.d/vras-vradmin.sh restart
```
2. Re-enter the command that failed.

After the product upgrade on secondary site, replication may fail to resume with "Secondary SRL missing" error [3931763]

When you attempt to resume the replication after the product upgrade on secondary site is complete, the replication may fail to resume with a configuration error "Secondary SRL missing".

This issue occurs because even after the product upgrade is complete, the Storage Replicator Log (SRL) volume remains disassociated from the Replicated Volume Group (RVG).

During a product upgrade, the installer pauses the replication and performs several tasks that include dissociation and association of SRL volume. Due to some internal error, the installer fails to reassociate the SRL volume to the RVG. As a result, when

you attempt to resume the replication from the primary site to the upgraded secondary site, it fails to start with a "Secondary SRL missing" error.

Workaround: Perform the following steps to restart the replication

1. On the upgraded site, associate the SRL to RVG

```
#vxvol -g DiskGroup_name aslog RVG_name SRL_name
```

2. Start RVG

```
# vxrvlg -g DiskGroup_name -f start RVG_name
```

3. Stop replication at primary site

```
# vradmin -g DiskGroup_name -f stoprep RVG_name
```

4. Start replication at primary site

```
# vradmin -g DiskGroup_name -a startrep RVG_name  
Secondary_hostname
```

vradmin repstatus command reports secondary host as "unreachable"(3896588)

The vradmin repstatus command output incorrectly reports all secondary hosts as unreachable if even one of the secondary hosts is unreachable in CVR/VVR multi-secondary environments.

Workaround: Run the following command to obtain the correct status:

```
# vradmin -g dg_name printrvlg rvlg_name
```

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 vxlink 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.

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

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

For example:

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

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

`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 restart
```

- 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:

```
# vxrvrg -g diskgroup stop rvrg
```
- 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 rvrg vol
```
- 7 Start the RVG. Enter the following:

```
# vxrvrg -g diskgroup start rvrg
```
- 8 Resume or start the applications.

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
```

Plex reattach operation fails with unexpected kernel error in configuration update (2791241)

In a VVR environment with layered volumes, if a DCM plex becomes detached because of a storage failure, reattaching the plex after fixing the storage issue fails with the following error:

```
VxVM vxplex ERROR V-5-1-10128 Unexpected kernel error in configuration
update
```

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 Storage Foundation HA 7.4.3 if you have configured or plan to configure bunker replication using VVR with volume sets.

Workaround:

Contact Veritas Technical Support for a patch that enables you to use this configuration.

SmartIO does not support write-back caching mode for volumes configured for replication by Volume Replicator (3313920)

SmartIO does not support write-back caching mode for volumes that are configured for replication by Volume Replicator (VVR).

Workaround:

If you have configured volumes for replication by VVR, do not enable write-back caching

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 *Storage Foundation Administrator's Guide*.

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.

DCM logs on a disassociated layered data volume results in configuration changes or CVM node reconfiguration issues (3582509)

If you have configured layered data volumes under an RVG that has DCM protection enabled and at a later point disassociate the data volume from the RVG, you must manually remove the DCM logs from the volume. Leaving DCM logs on a layered data volume after it has been disassociated from the RVG, may result configuration changes, or the CVM node reconfiguration to not work properly.

Workaround:

If the disk group has a layered volume, remove DCM logs after disassociating the volumes from the RVG.

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.

`vradmin -g dg repstatus rvg` displays the following configuration error: `vradmin not reachable on cluster peer` (3648854)

`vradmin -g dg rep status rvg` displays the following configuration error:

`vradmin` is not reachable on the cluster peer

However, replication is an ongoing process. The reason is that an unclean disconnect left the `vradmind` port open and in the **TIME_WAIT** state. An instance is as following:

```
# netstat -n | grep 8199
tcp        0          0 1:44781      1:8199
TIME_WAIT
tcp        0          0 1:44780      1:8199
TIME_WAIT
```

The following error message appear in `/var/vx/vras/log/vradmind_log_A`:

```
VxVM VVR Notice V-5-20-0 TAG_D IpmHandle:recv peer closed errno=0
VxVM VVR Debug V-5-20-8690 VRASCache TAG_E Cache_RLink
repstatus UPDATE message created for rlink rlk_192.168.111.127_rvg1
VxVM VVR Warning V-5-20-0 TAG_C IpmHandle::handleTo
vvr_sock_host_serv failed for 1111031
VxVM VVR Warning V-5-20-0 TAG_C IpmHandle::open: getaddrinfo
error(could not resolve srchost 1111032, error: Connection refused)
```

Workaround: Restart the `vradmind` daemon.

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

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

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

For example:

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

DCM plex becomes inaccessible and goes into DISABLED(SPARSE) state in case of node failure. (3931775)

In FSS environment, when a CVR is configured on primary and secondary site, DCM log plex is created by default on each volume. This log gets created with mirrors across local disks of single node, instead of mirrors across multiple nodes due to `vxassist` command limitation. This limitation restricts mirror, stripe separation, or confinement to allocate the log plexes to associate with the volume. Hence, DCM plex becomes inaccessible and goes into DISABLED (SPARSE) in case of node failure.

Workaround

1. While creating volumes, create and associate DCM logs manually from different nodes using the following command:

```
# vxassist -g <disk_group_name> addlog <volume_name> logtype=dcm  
  <local_disks_across_different_nodes>
```

2. Create RVG on the data volume.

Initial autosync operation takes a long time to complete for data volumes larger than 3TB (3966713)

If SmartMove is enabled and autosync is in progress, the SmartSync operation performs a difference-based synchronization, which is faster than full synchronization. However, for data volumes larger than 3TB, the SmartMove feature gets disabled if the allocated DCM plexes are not sufficiently sized. Therefore, the autosync operation performs full synchronization and synchronizes the entire volume.

Workaround

1. To enable the smartmove feature for volumes larger than 3TB, use the following command:

```
# vxassist -g <disk_group_name> addlog <volume_name> logtype=dcm  
  loglen=<size>
```

2. where, the *size* is a minimum of 1024 blocks.

Cluster Server known issues

This section describes the known issues in this release of Cluster Server (VCS). These known issues apply to the following products:

- Veritas InfoScale Availability

- Veritas InfoScale Enterprise

Operational issues for VCS

This section describes the Operational known issues for VCS.

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

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

Workaround: You need to restore at least one path.

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

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

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

Switching service group with DiskGroup resource causes reservation conflict with UseFence set to SCSI3 and powerpath environment set [2749136]

If `UseFence` is set to `SCSI3` and `powerpath` environment is set, then switching the service group with `DiskGroup` resource may cause following messages to appear in `syslog`:

```
reservation conflict
```

This is not a Veritas InfoScale issue. In case `UseFence` is set to `SCSI3`, the diskgroups are imported with the reservation. This message gets logged while releasing and reserving the disk.

Workaround: Refer to the TechNote at:

https://www.veritas.com/content/support/en_US/article.100006747

Stale NFS file handle on the client across failover of a VCS service group containing LVMLogicalVolume resource (2016627)

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

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

NFS cluster I/O fails when storage is disabled [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.

VVR configuration may go in a primary-primary configuration when the primary node crashes and restarts [3314749]

The AutoResync attribute of the RVGPrimary and RVGSharedPri agent control whether the agent must attempt to automatically perform a fast-failback resynchronization of the original primary after a takeover and after the original primary returns. The default value of this attribute is 0, which instructs the agent not to perform a fast-failback resynchronization of the original primary after a takeover and after the original primary returns. The takeover is performed automatically since the default value of the AutoTakeover attribute of the RVGPrimary and RVGShared agents is 1. Thus, the default settings of AutoTakeover and AutoResync set to 1 and 0 respectively cause the first failover to succeed when the original primary goes down, and on return of the original primary, the Replicated Data Set (RDS) ends up with a primary-primary configuration error.

Workaround: Set the default value of the AutoResync attribute of the RVGPrimary agent to 1 (one) when you want the agent to attempt to automatically perform a fast-failback resynchronization of the original primary after a takeover and after the original primary returns. This prevents the primary-primary configuration error. Do not set AutoResync to 1 (one) if you intend to use the Primary-Elect feature.

Moreover, if you want to prevent VCS from performing an automatic takeover and fast-failback resynchronization, set AutoTakeover and AutoResync attributes to 0 for all the RVGPrimary and RVGSharedPri resources in your VCS configuration. For more information, refer to the RVGPrimary and RVGSharedPri agent sections

of the *Replication Agents chapter* in the *Cluster Server Bundled Agents Reference Guide*.

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.

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 *Cluster Server Configuration and Upgrade Guide*.

VCS fails to stop volume due to a transaction ID mismatch error [3292840]

If VCS imports a disk group *A* on node *sys1*, which implies that the DiskGroup resource is online on *sys1*. If you run `vxdg -C import <dg_name>` outside VCS on node *sys2*, then the disk group gets imported on node *sys2* and `-c` clears the import locks and host tag. However on node *sys1*, disk group *A* continues to appear as imported and enabled, and hence, VCS continues to report the resource state as ONLINE on node *sys1*. Subsequently, when VCS detects the imported disk group on *sys2*, it deports the disk group from *sys2*, and imports it on *sys1* to resolve concurrency violation. At this point, the disk group deported from node *sys2* is shown as imported and enabled on node *sys1*. If you stop any volume from within or outside VCS, it fails with the `Transaction ID mismatch` error, but the read and write operations continue to function so the data continues to be accessible. This situation may lead to data corruption if the disk group appears enabled on multiple nodes. This issue is due to the Volume Manager behavior.

Workaround: Do not import a disk group using `-c` option if that diskgroup is under VCS control.

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 *Cluster Server Configuration and Upgrade 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 known issues about the VCS engine.

Invalid argument message in the message log due to Red Hat Linux bug (3872083)

An error message regarding the rtkit-daemon occurs due to the Red Hat Linux (RHEL) bug https://bugzilla.redhat.com/show_bug.cgi?id=999986

We have bypassed the system functionality for RHEL7, but the dependency check is performed before bypassing the systemctl. This is why the warning messages are logged.

Workaround:

There is no functionality effect. You can ignore the message.

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

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

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

The `hacf -cmdtoctf` 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 -cmdtoctf` command.

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

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

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

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

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

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

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

If the top level resource which does not have any parent dependency is disabled then the other resources do not come online and the following message is displayed:

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

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

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

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

Workaround: Terminate the agent process.

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

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

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

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

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

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

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 `hargp -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 *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).

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.

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

Veritas InfoScale 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 Veritas InfoScale 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 Veritas InfoScale 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 emptied 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.

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

On upgrading a secure cluster to 6.2 or later release, 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.

Global Cluster Option (GCO) require NIC names in specific format [3641586]

The `gcoconfig` script requires the NIC names in the letters followed by numbers format. For example, NIC names can be `eth0`, `eth123`, `xyz111` and so on. The script fails to configure GCO between NICs which do not comply with this naming format.

Workaround: Rename the NIC name and use the letters followed by numbers format to configure GCO.

If you disable security before upgrading VCS to version 7.0.1 or later on secured clusters, the security certificates will not be upgraded to 2048 bit SHA2 [3812313]

The default security certificates installed with VCS 7.0 and the earlier versions are 1024 bit SHA1. If you disable security before upgrading VCS to version 7.0.1 or later on secured clusters, the installer will upgrade VCS but will not upgrade the security certificates. Therefore, merely enabling security after the VCS upgrade to 7.0.1 or later does not upgrade the security to 2048 bit SHA2 certificates.

Workaround:

When you upgrade VCS to version 7.0.1 or later releases, run the `installer -security` command and select the `reconfigure` option to upgrade the security certificates to 2048 bit SHA2.

Java console and CLI do not allow adding VCS user names starting with ‘_’ character (3870470)

When a user adds a new user name, VCS checks if first character of the user name is part of the set of allowed characters. The ‘_’ character is not part of the permitted set. So the user name starting with ‘_’ is considered invalid.

Workaround: Use another user name which starts with a character permitted by VCS.

Issues related to the bundled agents

This section describes the known issues of the bundled agents.

Mounting an NFSv4 volume on the NFS client side fails

Mounting an NFSv4 volume on the NFS client side fails if `fsid` of the share is set to 0

This issue is a Red Hat issue. For details, visit:

<https://access.redhat.com/solutions/44693>

Workaround:

Make sure that `fsid` is not set to 0 (zero) for the NFS share

If multiple Mount resources uses the same block or volume, one of the resources may go into the OFFLINE or the UNKNOWN state (4001585)

Multiple NFS Mount resources can be associated with the same NFS server. While some resources may use the hostname of the NFS server in the `BlockDevice` attribute, others may use the IP address of the server. In such a scenario, if a Mount resource that has a hostname specified in the `BlockDevice` attribute comes online first, other resources that have an IP address in the `BlockDevice` attribute fail to come online or go into the UNKNOWN state. Similarly, if a Mount resource that has an IP address specified in the `BlockDevice` attribute comes online first, other resources that have a hostname in the `BlockDevice` attribute fail to come online or go into the UNKNOWN state.

Workaround:

When multiple NFS Mount resources that correspond to the same NFS server are configured, ensure that the `BlockDevice` attribute value is specified as either an IP address or a hostname for all of the resources

KVMGuest resource fails to work on VCS agent for RHEV3.5 (3873800)

When you configure RHEV3.5 guest as a resource in VCS (RHEV agent) of physical host, the KVMGuest resource does not probe.

Workaround:

To solve this issue, follow the steps:

- 1 The `havirtverify` utility fails since the `xpath` utility is not found on setup. Install the `perl-XML-XPath` package to fix it.
- 2 The monitor fails to match the cluster ID since you get FQDN host name, and on RHEV-M configuration you have plain host name.

Change to FQDN in the RHEVManager `CLUSTER > HOSTS`.

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

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

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

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

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

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

Workaround: No workaround for this known issue.

Concurrency violation observed during migration of monitored virtual machine [2755936]

If a VCS service group has more than one KVMGuest resource monitoring virtual machine and one of the virtual machines is migrated to another host, a service group level concurrency violation occurs as the service group state goes into PARTIAL state on multiple nodes.

Workaround: Configure only one KVMGuest resource in a Service group.

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

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

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

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

configuration). VCS detects this concurrency violation and brings down the guest on the second node.

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

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

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

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

Workaround:

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

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

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

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

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

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

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

Workaround:

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

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

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

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

CoordPoint agent remains in faulted state [2852872]

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

Workaround: After HAD has stopped, reconfigure fencing.

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

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

Workaround: No workaround.

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

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

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

KVMGuest agent fails to recognize paused state of the VM causing KVMGuest resource to fault [2796538]

In a SUSE KVM environment, when a virtual machine is saved, its state is changed to paused and then shut-off. The paused state remains for a very short period of time, due to timing in case that the KVMGuest agent misses this state. Then the resource state will be returned as OFFLINE instead of INTENTIONAL OFFLINE, which causes the KVMGuest resource to fault and failover.

This is due to the limitation of SUSE KVM as it does not provide a separate state for such events.

Workaround: No workaround.

Concurrency violation observed when host is moved to maintenance mode [2735283]

When a Red Hat Enterprise Virtualization host running a virtual machine is moved to maintenance state, the virtual machine migration is initiated by RHEV. Veritas InfoScale detects the migration according to virtual machine state, such as "migrating". Due to timing issue RHEV Manager occasionally sends the virtual machine state as "up" even if the migration is in progress. Due to this state, the resource is marked ONLINE on the node to which it migrates and may cause concurrency violation.

Workaround: No workaround.

Logical volume resources fail to detect connectivity loss with storage when all paths are disabled in KVM guest [2871891]

In a KVM environment if all storage paths are disabled, then LVMLogicalVolume and LVMVolumeGroup resources fails to detect the loss of connectivity with the storage. This occurs because of LVM2 commands return success even if all the paths to storage are disabled. Moreover, the LVMVolumeGroup and LVMLogicalVolume agents report the resource state as ONLINE.

Workaround: Verify the multi-pathing environment and make sure that all the read and write operations to the disk are blocked when all paths to the storage are disabled.

Resource does not appear ONLINE immediately after VM appears online after a restart [2735917]

During a VM restart the resource does not come ONLINE immediately after the VM starts running. As the VM state is 'Reboot in Progress' it reports INTENTIONAL OFFLINE and after VM is UP the resource cannot immediately detect it as the next monitor is scheduled after 300 seconds.

Workaround: Reduce the OfflineMonitorInterval and set it to suitable value.

Unexpected behavior in VCS observed while taking the disk online [3123872]

If the VMwareDisks resource is configured for a disk connected to another virtual machine outside of an ESX cluster and if you bring the disk online on the configured node, you may observe unexpected behavior of VCS (like LLT connection break). The behavior is due to a known issue in VMware.

Workaround: Remove the disk from the other virtual machine and try again.

LVMLogicalVolume agent clean entry point fails to stop logical volume if storage connectivity is lost [3118820]

If storage connectivity is lost on a system on which the LVM resources are in ONLINE state and a volume is mounted using the Mount resource, LVMVolumeGroup agent monitor entry point detects the loss of connectivity and returns the resource state as offline. This causes agent framework to call clean entry point of LVMVolumeGroup agent; however, the state of the resource stays online. Agent framework waits for the clean entry point to return success so that the resource can be moved to the offline|faulted state. At this stage, the clean entry point fails as it is not able deactivate and export the volume group because the logical volume is mounted. There is no option available to forcefully deactivate and export the volume group. Hence, the service groups get stuck in this state. Even if the storage connectivity is restored, the problem does not resolve because the logical volume remains mounted. If the logical volume is unmounted, then the LVMVolumeGroup resource goes into FAULTED state and service group fails over.

Workaround: Manually unmount the logical volume.

VM goes into paused state if the source node loses storage connectivity during migration [3085214]

During virtual machine migrations in a RHEV environment, the VM may freeze in paused state if the source host loses storage connectivity. This issue is specific to RHEV environment.

Workaround: No workaround.

Virtual machine goes to paused state during migration if the public network cable is pulled on the destination node [3080930]

The virtual machine goes into paused state during migration if the public network cable is pulled on the destination node. This behavior depends on the stage at which the migration is disrupted. The virtual machine rolls back to the source node if the network cable is pulled during migration. Resource on the source node reports this as an online virtual machine that is in running state. On the destination node, the virtual machine goes into shut-off state.

If the virtual machine migration gets disrupted during the transfer from source to destination, it may happen that the virtual machine remains in paused state on the source node. In such a case, you must manually clear the state of the virtual machine and bring it online on any one node.

This operational issue is a behavior of the technology and has no dependency on Veritas InfoScale. This behavior is observed even if the migration is invoked outside VCS control. Due to the disruption in virtual machine migration, it may happen that the locking mechanism does not allow the virtual machine to run on any host, but again, this is a virtualization technology issue.

Workaround: No workaround. Refer to the virtualization documentation.

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.

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

Manual configuration of RHEVMInfo attribute of KVMGuest agent requires all its keys to be configured [3277994]

The RHEVMInfo attribute of KVMGuest agent has 6 keys associated with it. When you edit main.cf to configure RHEVMInfo attribute manually, you must make sure that all the keys of this attribute are configured in main.cf. If any of its keys is left unconfigured, the key gets deleted from the attribute and agent does not receive the complete attribute. Hence, it logs a Perl error `Use of uninitialized value` in the engine log. This is due to the VCS engine behavior of handling the attribute with key-value pair.

Workaround: Use `ha` commands to add or modify RHEVMInfo attribute of KVMGuest resource.

SambaServer agent may generate core on Linux if LockDir attribute is changed to empty value while agent is running [3339231]

If LockDir attribute is changed to an empty value while agent is running and debugging is enabled, the logging function may access invalid memory address resulting in SambaServer agent to generate core dump.

Workaround: When LockDir attribute is changed while agent is running, ensure that its new value is set to a non-empty valid value.

Independent Persistent disk setting is not preserved during failover of virtual disks in VMware environment [3338702]

VMwareDisks agent supports Persistent disks only. Hence, Independent disk settings are not preserved during failover of virtual disk.

Workaround: No workaround.

LVMLLogicalVolume resource goes in UNABLE TO OFFLINE state if native LVM volume group is exported outside VCS control [3606516]

If you export the LVM volume group without stopping LVM logical volumes, the LVMLLogicalVolume resource falsely reports online. If offline is initiated for LVMLLogicalVolume resource, it fails as the volume group was not exported cleanly and LVMLLogicalVolume Agent fails to deactivate the logical volume causing LVMLLogicalVolume to go in UNABLE TO OFFLINE state.

Workaround: Make sure volume group is deactivated and exported using VCS or manually deactivate the LVM logical volumes.

DiskGroup resource online may take time if it is configured along with VMwareDisks resource [3638242]

If a service group is configured with VMwareDisks and DiskGroup resource, the DiskGroup resource may take time to come online during the service group online. This is because VxVM takes time to recognize a new disk that is attached by VMwareDisks resource. A VMwareDisks resource attaches a disk to the virtual machine when the resource comes online and a DiskGroup resource, which depends on VMwareDisks resource, tries to import the disk group. If `vxconfigd` does not detect the new disk attached to the virtual machine, the DiskGroup resource online fails with the following error message because the resource is not up even after the resource online is complete.

```
VCS ERROR V-16-2-13066 ... Agent is calling clean for resource(...)
```

Workaround: Configure OnlineRetryLimit to appropriate value.

For example, if the DiskGroup resource name is `res_rawdg`:

```
# hares -override res_rawdg OnlineRetryLimit
# hares -modify res_rawdg OnlineRetryLimit 2
```

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

SFCache agent cannot enable caching if cache area associated with this particular object is in offline state. User needs 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 or later in secure mode.

Workaround: Restart the RemoteGroup agent.

VMwareDisks agent may fail to start or storage discovery may fail if SELinux is running in enforcing mode [3106376]

The VMwareDisks agent and discFinder binaries refer to the `libvmwarevcs.so` shared library. SELinux security checks prevent discFinder from loading the `libvmwarevcs.so` library, which requires text relocation. If SELinux is running in enforcing mode, these two executables may report the "Permission denied" error and may fail to execute.

Workaround:

Enter the following command and relax the security check enforcement on the Veritas `libvmwarevcs.so` library:

```
# chcon -t textrel_shlib_t '/opt/VRTSvcs/lib/libvmwarevcs.so'
```

Issues related to the VCS database agents

This section describes the known issues about VCS database agents.

Unsupported startup options with systemd enabled [3901204]

This is applicable when systemd is enabled on RHEL 7 and SLES 12 Linux distributions.

With systemd enabled, an Oracle single instance or Oracle RAC application does not support `SRVCTLSTART` and `SRVCTLSTART_RO` startup options.

With systemD enabled, an Oracle ASMInst application does not support SRVCTLSTART, SRVCTLSTART_OPEN, and SRVCTLSTART_MOUNT start up options.

VCS ASMDG resource status does not match the Oracle ASMDG resource status (3962416)

In an Oracle environment, the status of an ASMDG resource as gathered through the srvctl command may not match the status that is gathered through a SQL query. A mismatch in these values causes a conflict between the status of the VCS ASMDG resource (online) and the status of the Oracle ASMDG resource (online/starting). As a result, the resources that are dependent on ASMDG do not come online and go into the faulted state.

Workaround: No workaround.

ASMDG agent does not go offline if the management DB is running on the same (3856460)

If an offline is fired on the node on which Flex ASM is running and the same node has Management DB running on it, then the same would not go offline.

Workaround: Use commands to migrate the Management DB to another node before getting the Flex ASM offline. The following is a sample command that you can use to check if the Management DB is running on a node:

```
# /oracle/12102/app/gridhome/bin/srvctl status mgmtldb -verbose
Database is enabled
Instance -MGMTDB is running on node vcslx017. Instance status: Open.
```

The following is a sample command that you can use to migrate the Management DB to another node:

```
# /oracle/12102/app/gridhome/bin/srvctl relocate mgmtldb -node vcslx018
```

ASMDG on a particular does not go offline if its instances is being used by other database instances (3856450)

If you initiate an offline of the ASMDG group on a node which has its ASMInstance being used by one of more DB z resources from the cluster, then the offline would fail and a fault would get reported on both the ASM and DB level.

Workaround: Run the following SQL command to check the ASM DG running on the node:

```
SQL> select INST_ID, GROUP_NUMBER, INSTANCE_NAME,  
DB_NAME, INSTANCE_NAME||':'||DB_NAME client_id from gv$asm_client;
```

INST_ID	GROUP_NUMBER	INSTANCE_NAME	DB_NAME	CLIENT_ID
3	2	oradb2	oradb	oradb2:oradb
3	2	oradb3	oradb	oradb3:oradb
3	2	+ASM3	+ASM	+ASM3:+ASM
3	1	+ASM3	+ASM	+ASM3:+ASM
1	2	oradb1	oradb	oradb1:oradb
1	1	-MGMTDB	_mgmtdb	-MGMTDB:_mgmtdb
1	1	+ASM1	+ASM	+ASM1:+ASM
4	2	oradb4	oradb	oradb4:oradb

8 rows selected.

In the above table:

- oradb1 is using the ASMInstance 1
- oradb2 and oradb3 are using ASMInstance 3
- oradb4 is using ASMInstance 4

Use the following SQL to relocate the ASMPool to another node:

```
SQL> alter system relocate client 'oradb4:oradb';  
System altered.
```

If the command does not work, please refer Oracle documentation for further information on relocating the client.

Sometimes ASMDG reports as offline instead of faulted (3856454)

Sometimes, you may observe that the agent reports the ASMDG state for the node where the ASM instance is down as offline instead of as faulted, even when the

cardinality is violated. This occurs in scenarios in which the ASM instance is abruptly shut down.

Workaround: No workaround.

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 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 *Cluster Server Configuration and Upgrade 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.

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 status 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 differentiate 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$pdb`s) 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 come online and monitor Oracle instance if `threaded_execution` parameter is set to true (3644425)

In Oracle Database 12c or later, the threaded execution feature is enabled. The multithreaded Oracle Database model lets Oracle processes execute as operating system threads in separate address spaces. If Oracle Database 12c or later is installed, the database runs in the process mode. If you set a parameter to run the database in the threaded mode, only some background processes on UNIX and

Linux run with each process containing one thread. The remaining Oracle processes run as threads within those processes.

When you enable the `threaded_execution` parameter, the Oracle agent cannot check the `smon` (mandatory process check) and the `lgwr` (optional process check) processes. These processes were traditionally used for monitoring, and they now run as threads.

Workaround: Disable the threaded execution feature, because it is not supported on Oracle Database 12c or later.

Issues related to the agent framework

This section describes the known issues about the agent framework.

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

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

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

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

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

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

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

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

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

- 4 Wait for several minutes to ensure that VCS has executed all pending commands, and then execute any new VCS command.
- 5 If the delay persists, repeat step 2 or 3 as appropriate.
- 6 If the CPU usage returns to normal limits, revert the attribute changes to the backed up values to avoid the delay in detecting the resource fault.

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 Veritas 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.

VCS fails to process the `hares -add` command resource if the resource is deleted and subsequently added just after the VCS process or the agent's process starts (3813979)

When VCS or the agent processes start, the agent processes the initial snapshots from the engine before probing the resource. During the processing of the snapshots, VCS fails to process the `hares -add` command, thereby skipping the resource addition operation and subsequently failing to probe the resource.

Workaround: This behavior is by the current design of the agent framework.

Cluster Server agents for Volume Replicator known issues

The following are new additional Cluster Server agents for Volume Replicator known issues in 7.4.3 release.

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

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

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

Stale entries are found in sample `main.cf` file for RVGLogowner agent and RVGPrimary agent.

The stale entries are present in the `main.cf.seattle` and `main.cf.london` files on the RVGLogowner agent which includes CFSQlogckd resource. However, CFSQlogckd is not supported since VCS 5.0.

On RVGPrimary agent, the stale entries are present in file `main.cf.seattle` and `main.cf.london` and the stale entry includes the DetailMonitor attribute.

Workaround

1 For `main.cf.seattle` for RVGLogowner agent in the `cvm` group:

- Remove the following lines.

```
CFSQlogckd qlogckd (
    Critical = 0
)

cvm_clus requires cvm_vxconfigd
qlogckd requires cvm_clus
vxfsckd requires qlogckd
```

```
// resource dependency tree
//
//     group cvm
//     {
//         CFSfsckd vxfsckd
//         {
//             CFSQlogckd qlogckd
//             {
//                 CVMCluster cvm_clus
//                 {
//                     CVMVxconfigd cvm_vxconfigd
//                 }
//             }
//         }
//     }
// }
```

- Replace the above lines with the following:

```
cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus
```

```
// resource dependency tree
//
//     group cvm
//     {
//         CFSfsckd vxfsckd
//         {
//             CVMCluster cvm_clus
//             {
//                 CVMVxconfigd cvm_vxconfigd
//             }
//         }
//     }
// }
```

2 For main.cf.london for RVGLogowner in the cvm group:

- Remove the following lines

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

```
cvm_clus requires cvm_vxconfigd
qlogckd requires cvm_clus
vxfsckd requires qlogckd

// resource dependency tree
//
//     group cvm
//     {
//     CFSfsckd vxfsckd
//     {
//     CFSQlogckd qlogckd
//     {
//     CVMCluster cvm_clus
//     {
//     CVMVxconfigd cvm_vxconfigd
//     }
//     }
//     }
//     }
//     }
```

- Replace the above lines with the following:

```
cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus

// resource dependency tree
//
//     group cvm
//     {
//     CFSfsckd vxfsckd
//     {
//     CVMCluster cvm_clus
//     {
//     CVMVxconfigd cvm_vxconfigd
//     }
//     }
//     }
```

- 3 For main.cf.seattle for RVGPrimary agent in the cvm group:

- In the group ORAGrp and for the Oracle resource database, remove the line: `DetailMonitor = 1`
- 4 For main.cf.london for RVGPrimary agent in the cvm group:
- In the group ORAGrp and for the Oracle resource database, remove the line: `DetailMonitor = 1`

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 CFSSMount 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 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.

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

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

Workaround: Use the correct commands:

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

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 Cluster Server Manager (Java Console).

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

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

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

Workaround: Ignore the error.

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

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

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

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

VCS Cluster Configuration wizard issues

VCS Cluster Configuration wizard does not automatically close in Mozilla Firefox [3281450]

You can use the `haappwizard` utility to launch the High Availability wizard to configure application monitoring with Veritas Cluster Server (VCS) on Linux systems. If you configure the utility to launch the wizard in Mozilla Firefox browser, the browser session does not automatically close after the VCS configuration is complete.

Workaround: Use one of the following workarounds:

- Close the Mozilla Firefox browser session once the wizard-based configuration steps are complete.
- Specify a different browser while configuring the `haappwizard` utility.

Configuration inputs page of VCS Cluster Configuration wizard shows multiple cluster systems for the same virtual machine [3237023]

The **Configuration inputs** panel of the VCS Cluster Configuration wizard shows multiple cluster systems for the same virtual machine. This occurs because the value specified for the node in the `SystemList` attribute is different than the one returned by the `hostname` command.

Workaround: Ensure that the value specified for the node in the `SystemList` attribute and the one returned by the `hostname` command is the same.

VCS Cluster Configuration wizard fails to display mount points on native LVM if volume groups are exported [3341937]

On storage selection page of application wizard, mount points mounted on native LVM devices are not shown. If you have one or more native LVM volume groups, and one of them is exported, the application wizard fails to detect mount points configured on these devices.

Workaround: Ensure that you do not have any native volumes exported if you want to configure application that uses native LVM storage.

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.

InfoScale Enterprise: Unable to configure clusters through the VCS Cluster Configuration wizard (3911694)

Unable to configure clusters through the VCS Cluster Configuration wizard in InfoScale Enterprise because the `configure_cluster.response` file that is created has the product type as `Availability` instead of `Enterprise`.

Workaround: In InfoScale Enterprise, configure clusters through the CPI.

Cluster Configuration Wizard fails to configure a cluster due to missing telemetry data (4002133)

The Cluster Configuration Wizard fails to configure a cluster and logs the following error message:

```
<Message>Failed to configure the VCS cluster. Refer to the logs on the system for more details.</Message>
<Debug><![CDATA[CPI ERROR V-9-40-1030 $cfg->{edgeserver_host} and $cfg->{edgeserver_port} should be defined in the responsefile]]></Debug>
```

This issue occurs because the Cluster Configuration Wizard does not yet support the telemetry attributes, `edgeserver_host` and `edgeserver_port`.

Workaround:

Configure the VCS cluster using the InfoScale product installer. Then, use the Veritas High Availability Configuration wizard to configure applications for monitoring.

LLT known issues

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

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

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

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

Rolling upgrade from earlier version to InfoScale 7.4.2 may fail for LLT over UDP configuration in FSS environment (3981917)

In a Flexible Storage Sharing (FSS) environment where LLT is configured over UDP, the rolling upgrade of InfoScale from earlier version to 7.4.2 may fail. The failure may occur if the NIC MTU is set to 9000 and the value of the MTU field in the `/etc/llttab` file is set to – (dash). Due to the mismatch in the MTU values, the remote I/Os initiated by the VxVM layer may get stuck and the upgraded node remains in Joining state.

Workaround:

In an FSS environment where LLT is configured over UDP, before you perform a rolling upgrade from an earlier InfoScale version to 7.4.2, ensure that the MTU value in the `/etc/llttab` file is set to 1500 instead of a – (dash).

That is, instead of setting the value as dash as:

```
link eth1 udp - udp 50000 - 192.168.10.1 -
```

Set the MTU value to 1500 manually:

```
link eth1 udp - udp 50000 1500 192.168.10.1 -
```

If you manually re-plumb (change) the IP address on a network interface card (NIC) which is used by LLT, then LLT may experience heartbeat loss and the node may panic (3188950)

With the LLT interfaces up, if you manually re-plumb the IP address on the NIC, then the LLT link goes down and LLT may experience heartbeat loss. This situation may cause the node to panic.

Workaround: Do not re-plumb the IP address on the NIC that is currently used for LLT operations. Take down the stack before you re-plumb the IP address for the LLT interface.

A network restart of the network interfaces may cause heartbeat loss for the NIC interfaces used by LLT

A network restart may cause heartbeat loss of the network interfaces configured LLT. LLT configured for UDP or LLT configured for RDMA may experience loss of heartbeat between the interfaces, which may cause the node to panic.

Workaround: Recommendations before you restart the network:

- Assess the effect of a network restart on a running cluster that is using LLT over RDMA or LLT over UDP.
- Do not use the network restart functionality to add or configure a new NIC to the system.
- If you are using the network restart functionality, make sure that the LLT interfaces are not affected.
- Increase the `llt-peerinact` time to a higher value to allow network restart to complete within that time.
Run the `# lltconfig -T peerinact:6000` command to increase the `peerinact` time to 1 minute.

Performance degradation occurs when RDMA connection between nodes is down [3877863]

In clusters communicating over RDMA connections, when you reboot cluster nodes, services come online and nodes in the cluster communicate over LLT links. But, sometimes, the RDMA connections between nodes do not come back online. This affects node performance. Such cases are typically seen with 8 node clusters and beyond. To check the status of RDMA links, run the `lltstat -nvvr` configured command on each node to check whether the status of TxRDMA and RxRDMA link is Down.

Workaround: You can either manually restart the stack on all the nodes or run CPI to restart cluster nodes. Terms of use for this information are found in Legal Notices.

After configuring LLT over UDP using IPV6, one of the configured link may show DOWN status for `lltstat` command [3916374]

On configuring LLT over UDP using IPV6, one of the configured links shows down status. However, the IP address for this link can successfully be pinged.

This happens due to the default route settings for ipv6. In case of UDP, the Linux Kernel's UDP APIs rely on the Linux networking stack to send data. In some cases the networking stack uses a different source IP address than the correct one to send packets on a link.

Workaround:

Set the proper static routes so that the networking stack can use the correct network interfaces.

Use the following commands to specify that the destinations with address <IPv6 address> are directly reachable through interface <network interface> on each node.

```
ip route add <Destination_IPV6 address> via <Destination_IPV6 address>  
dev <network interface>
```

For example:

```
ip route add ef80::21a:64ff:fe93:1a92 via ef80::21a:64ff:fe93:1a92  
dev eth2
```

When using FSS over RDMA links during heavy IO, LLT may face link fluctuations [3907179]

On RHEL 7.3, when using FSS over RDMA links during heavy IO, in a rare case, llc may face issues with link fluctuations, and the system may panic. However, the panicked node comes up and join the cluster automatically.

Workaround: Reboot the panicked node.

The LLT window may drop to a very low value in CVM/FSS or CFS environment [3914954]

Sometimes, in CVM/FSS or CFS environment, due to the LLT adaptive window feature, the LLT window drops to a very low value and it affects performance.

Workaround: In such a case, disable the adaptive window feature and set the adaptive window manually to an optimal value, with which minimal retransmits are seen based on llstat output.

When using response files for LLT configuration over UDP, the nodes become unresponsive (3946836)

On SLES, when LLT is configured over UDP using response file, and the low-priority heartbeat Link uses public IP, the nodes become unresponsive. This happens as the public IP gets deleted during the link configuration. This leads to failed network connectivity making the nodes unresponsive.

Workaround: When using response files for LLT configuration over UDP, do not use public IPs for configuring the low-priority heartbeat links.

LLT causes node to panic during TCP connection failure when incomplete packets are received (3944294)

When LLT is configured using TCP and the TCP network experience problems, incomplete packets might be delivered to LLT during the receive operation. LLT cannot handle the partially received data and it panics the receiving node with the stack trace similar to:

```
Call Trace:
dump_stack+0x19/0x1b
panic+0xe3/0x1f2
llt_tcp_recv.part.6+0x32/0x551
llt_tcp_recv+0xc8/0xd0
llt_process_socket+0x2a4/0x6f0
llt_tdlv_thread+0x32c/0x710
```

Workaround: There is no workaround for this issue. However, when the node panics, the receiver node is not visible to the sender node for sending further packets. LLT intentionally panics the node to avoid any data corruption or data loss or node hangs. GAB reconfiguration then handles the change in cluster membership to inform all clients about the node failure to take necessary action. When the panicked node reboots, it automatically joins back the cluster.

I/O fencing known issues

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

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.

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).

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.

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 the `VRTSvxfen` RPM, 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, `vxfsd`, 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.

VCS fails to take virtual machines offline while restarting a physical host in RHEV and KVM environments (3320988)

In RHEV and KVM environments, the virtualization daemons `vdsm` and `libvirt` required to operate virtual machines are stopped before VCS is stopped during a reboot of the physical host. In this scenario, VCS cannot take the virtual machine resource offline and therefore the resource fails to stop. As a result, LLT, GAB and fencing fail to stop. However, the virtual network bridge is removed leading to the loss of cluster interconnects and causing a split-brain situation.

Workaround: If the virtual network bridge is not assigned to any virtual machine, remove the virtual bridge and configure LLT to use the physical interface. Alternatively, before initiating a reboot of the physical host, stop VCS by issuing the `hastop -local` command. The `-evacuate` option can be used to evacuate the virtual machines to another physical host.

Fencing may panic the node while shut down or restart when LLT network interfaces are under Network Manager control [3627749]

When the LLT network interfaces are under Network Manager control, then shutting down or restarting a node may cause fencing race resulting in a panic. On RHEL, VCS requires that LLT network interfaces are not put under Network Manager control, as it might cause problems when a node is shut down or restarted. During shutdown, the Network Manager service might stop before the VCS shutdown scripts are called. As a result, fencing race is triggered and the losing sub-cluster panics.

Workaround: Either exclude the network interfaces to be used by LLT from Network Manager control or disable the Network Manager service before configuring LLT. Please refer to the Red Hat documentation to do the same.

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

After you remove a Coordinator disk used by fencing or fencing disk group by running the `vxddmpadm 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 `vxddmpadm include dmpnodename=<dmp_disk/node>` command to again enable the dmp disk. This disk will show up in subsequent `vxfenconfig -l` output.

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 *Cluster Server Administrator's Guide*.

Storage Foundation and High Availability known issues

This section describes the known issues in this release of Storage Foundation and High Availability (SFHA). These known issues apply to Veritas InfoScale Enterprise.

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.

Installer exits upgrade to 5.1 RP1 with Rolling Upgrade error message (1951825, 1997914)

Installer exits upgrade to 5.1 RP1 with Rolling Upgrade error message, if protocol version entries are present in `/etc/gabtab` and `/etc/vxfenmode` files. Installer program may exit with either one of the following error messages during upgrade from 5.1 to 5.1 RP1:

```
SF51 is installed. Rolling upgrade is only supported from 5.1 to
higher version for the products
```

Or

```
To do rolling upgrade, VCS must be running on <node>.
```

Workaround: If the protocol version entries are present in `/etc/gabtab` and `/etc/vxfenmode` files, then installer detects it as Rolling Upgrade (RU). If you are not attempting RU, and doing full upgrade to 5.1 RP1, remove the protocol version entries from these two files for installer to proceed with regular upgrade.

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.

Process start-up may hang during configuration using the installer (1678116)

After you have installed a Storage Foundation product, some Veritas Volume Manager processes may hang during the configuration phase.

Workaround: Kill the installation program, and rerun the configuration.

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/rac1lgl1.DB_NAME
(No such file or directory).
SFORA vxreptadm ERROR V-81-8847 Cannot get filename from sid
for 'rac1lgl1', rc=-1.
SFORA vxreptadm ERROR V-81-6550 Could not connect to repository
database.
VxVM vxdg ERROR V-5-1-582 Disk group SNAP_rac1ldgl1: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac1ldgl1.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac1ldgl1 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.

Storage Foundation Cluster File System High Availability known issues

This section describes the known issues in this release of Storage Foundation Cluster File System High Availability (SFCFSA). These known issues apply to the following products:

- Veritas InfoScale Storage
- Veritas InfoScale Enterprise

During disk group creation with ‘-o same_enckey=yes’, disk reattach fails if a disk comes online after losing connectivity (4003890)

To use a single key for all the encrypted volumes in a disk group, the value of the `same_enckey` option must be set to yes during disk group creation. The `same_enckey` option sets an additional feature configuration for the disk group. After creating the disk group, if a disk comes online after losing connectivity, the disk reattach operation is performed. This operation does not allow any additional configuration on disk group and causes the disk reattach operation to fail.

Workaround:

Perform one of the following tasks to reattach the disk after it comes online:

Deport the disk group and import it again.

```
# vxdg deport diskGroupName
```

```
# vxdg importOptions import diskGroupName
```

- 1 Deport the disk group .

```
# vxdg deport diskGroupName
```

- 2 Import the disk group.

```
# vxdg importOptions import diskGroupName
```

If it is not feasible to deport and then import a disk group due to an ongoing application load, perform the following steps:

- 1 Obtain the private region length of the disk.

```
# vxdisk list diskName | grep "private:" | awk '{print $4}' | cut
-d= -f2
```

- 2 Initialize the disk.

```
# vxdisk -f init diskName privlen=privateRegionLength
```

Here, *privateRegionLength* is the value that is obtained in step 1.

- 3 Reattach the disk.

```
# vxdg -g diskGroupName -k adddisk diskName
```

- 4 Restart the volumes in the disk group.

```
# vxrecover
```

Transaction hangs when multiple plex-attach or add-mirror operations are triggered on the same volume (3969500)

In an FSS or a CVM environment where the `vol_intent_lock` tunable is set to 1, an I/O count is taken on a volume each time a 'plex attach' or 'add mirror' operation is triggered. During this timeframe, if the same volume undergoes one more 'plex attach' or 'add mirror' operation, another transaction is triggered. This transaction waits for I/O count quiesce or drain. However, in some cases, the I/O count that is taken by previous 'plex attach' or 'add mirror' operation may not be handled correctly. In such a situation, the second 'plex attach' or 'add mirror' operation hangs and subsequently fails with the following transaction timeout error:

```
vxvm:vxconfigd: V-5-1-8011 Internal transaction failed:
Transaction aborted waiting for io drain.
```

Workaround

1. Restart the master node.
2. Set the value of the `vol_intent_lock` tunable to 0 from the new master by using the following command:

```
# vxtune vol_intent_lock 0
```

In an FSS environment, creation of mirrored volumes may fail for SSD media [3932494]

In an FSS environment where SSD devices are used from Storage Access Layer (SAL), the creation of mirrored volumes may fail if `vxconfigd` is restarted on the master node.

This issue occurs because the Mediatype attribute for a device is inconsistently propagated from the kernel during `vxconfigd` startup.

Workaround: Before creating a disk group, set the media type attribute to SSD

```
vxdisk set -f diskname mediatype=ssd
```

Mount command may fail to mount the file system (3913246)

For a file system that was earlier mounted on cluster nodes the first `Mount` command may fail with the following error:

```
UX:vxfs mount.vxfs: ERROR: V-3-28543: Cannot be mounted until it has been cle
Please run "fsck -t vxfs -y /dev/vx/rdisk/<DiskGroup_Name>/<Volume_Name>" befo
Please refer to fsck_vxfs man page for details.
```

At the same time, the following error message may appear in the system log:

```
vxfs: msgcnt 463 mesg 021: V-2-21: vx_fs_init - /dev/vx/dsk/<DiskGroup_Name>/
file system validation failure.
```

These are generic messages and the `Mount` command may fail due to multiple reasons.

Workaround:

1. Verify all the per node logs to check for a dirty log, if any.
2. Mount the file system with `delayfsck` mount option.

Notes:

- This workaround is applicable only if a dirty log is found.
- Step 2 is applicable only for disk layout version 11 or later. For disk layout version prior to 11, you must run the `Full fsck` command before the file system is mounted.

- Even if you have mounted the file system with `delayfsck` mount option, you must run the `Full fsck` command at a later point in time. You may plan the application downtime and then run the `Full fsck` command.

After the local node restarts or panics, the FSS service group cannot be online successfully on the local node and the remote node when the local node is up again (3865289)

When all the nodes that are contributing storage to a shared Flexible Storage Sharing (FSS) DG leave the cluster, the CVMVoIDG resources and their dependent resources such as CFSSMount will be FAULTED. When the nodes rejoin the cluster, the resources/service groups will still remain in the FAULTED or OFFLINE state.

Workaround:

The FAULT on these resources should be manually CLEARED and the OFFLINED resources or service groups should be manually ONLINED.

- To clear the fault on the resource, use the following command:

```
# hares -clear <res> [-sys <system>]
```

- To bring the individual OFFLINED resource to the ONLINE state, use the following command:

```
# hares -online [-force] <res> -sys <system>
```

- To bring all the OFFLINED resource under a service group to the ONLINE state, use the following command:

```
# hagrps -online [-force] <group> -any [-clus <cluster> | -localclus]
```

In the FSS environment, if DG goes to the `dgdisable` state and deep volume monitoring is disabled, successive node joins fail with error 'Slave failed to create remote disk: retry to add a node failed' (3874730)

In the Flexible Storage Sharing (FSS) environment, if deep monitoring is not enabled for the volume used for the file system, the CVMVoIDg agent is able to detect fault and deport the disabled DG. Any new node joining to the cluster fails with error:

```
# /opt/VRTS/bin/vxclustadm -v nodestate
state: out of cluster
reason: Slave failed to create remote disk: retry to add a node failed
```

Workaround:

Enable deep monitoring for the resource using the '-D' option during adding the service group:

```
# cfsmntadm add -D <dgname> <volname> <mountpoint>all=cluster
```

If you have created the service group, use the below command to enable the deep monitoring of volumes:

```
# hares -modify <res_name> CVMVolumeIoTest <vol_list>
```

DG creation fails with error "V-5-1-585 Disk group punedatadg: cannot create: SCSI-3 PR operation failed" on the VSCSI disks (3875044)

If the disks that do not support SCSI3 PR are used to create the shared disk group, the operation fails as the data disk fencing functionality cannot be provided on such disks. The operation fails with error:

```
VxVM vxdg ERROR V-5-1-585 Disk group <DGNAME>: cannot create: SCSI-3
PR operation failed
```

Workaround:

If you still want to allow such disks to be part of shared disk group, disable the data disk fencing functionality in the cluster by running the command on all the nodes in the cluster:

```
# vxdctl scsi3pr off
```

After the disabling process, take caution that it may not protect the disks against the ghost I/Os from nodes that are not part of the cluster.

CVMVOLDg agent is not going into the FAULTED state. [3771283]

In CVMVOLDg monitor script we are not able to parse a variable and hence the volume does not go into the disabled state. This is the reason why the CVMVOLDg agent is not going into the FAULTED state.

Workaround:

Enable CVMVOLIOTEST on the volume for the resource to go into FAULTED state, using the following commands:

```
# haconf -makerw
```

```
# hares -modify test_vol_dg CVMVolumeIoTest testvol
```

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

On CFS, SmartIO is caching writes although the cache appears as nocache on one node (3760253)

On CFS, SmartIO is caching writes although the `sfcache list` output shows the cache in `nocache` mode on one node. The OS `mount` command also shows the file systems as unmounted. This issue is due to a known bug that is documented in the Linux `mount` manual page. The `/etc/mtab` file and the `/proc/mounts` file, which are expected to have entries for all the mounted file systems, do not match. When the `sfcache list` command displays the list of file systems that are mounted in writeback mode, `sfcache list` refers to the `/etc/mtab` entries for the mount status of the file systems. As a result, `sfcache list` may sometimes show a writeback enabled file system as unmounted while in reality the file system is still mounted. The `/proc/mounts` file correctly shows the file systems as mounted.

Workaround:

Verify that the file system is mounted through the contents of the `/proc/mounts` file.

tail -f run on a cluster file system file only works correctly on the local node [3741020]

When you use the `tail -f` command(1M) to monitor a file on a cluster file system, changes to the file made on remote nodes are not detected. This is due to the `tail` command now utilizing `inotify`. Veritas is currently unable to support `inotify` with a cluster file system due to GPL restrictions.

Workaround:

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

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.

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.

Storage Foundation for Oracle RAC known issues

This section describes the known issues in this release of Storage Foundation for Oracle RAC (SFRAC). These known issues apply to Veritas InfoScale Enterprise.

Oracle RAC known issues

This section lists the known issues in Oracle RAC.

Oracle Grid Infrastructure installation may fail with internal driver error

The Oracle Grid Infrastructure installation may fail with the following error:

```
[INS-20702] Unexpected Internal driver error
```

Workaround:

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

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

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

Storage Foundation Oracle RAC issues

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

Oracle database or grid installation using the product installer fails (4004808)

The product installer does not support installation of Oracle 12cR2 and 19c. As a result, if you use the product installer for Oracle database or grid software installation, it fails.

Workaround:

Install Oracle database or grid infrastructure using Oracle installer instead of the Product Installer.

ASM configuration fails if OCR and voting disk volumes are configured on VxFS or CFS for Oracle 19c during the grid installation (4003844)

If you configure OCR and voting disk volumes on VxFS or CFS for Oracle 19c during Grid installation, ASM fails to start when the ASM configuration assistant is invoked via the `asmca` command. This issue has been reported to Oracle (Bug id 28726240).

Workaround:

If you plan to use ASM, ensure that you configure OCR and voting disk volumes on ASM while installing the 19c Grid.

CSSD configuration fails if OCR and voting disk volumes are located on Oracle ASM (3914497)

The Veritas installer fails to configure CSSD if OCR and voting disk volumes are located on Oracle ASM. This is because the installer does not support the configuration of CSSD with OCR and voting disk volumes on Oracle ASM.

Workaround: Configure the CSSD resource manually.

For instructions, see *Section: Installation and upgrade of Oracle RAC* in the *Storage Foundation for Oracle RAC Configuration and Upgrade Guide* document.

When you upgrade to SF Oracle RAC 7.1, VxFS may fail to stop (3872605)

When you upgrade to SF Oracle RAC 7.1, VxFS may fail to stop. This is because the reference count holds on VxFS while the system unregisters AMF and unmounts the file system.

Workaround:

Before upgrading, disable AMF and set `AMF_START=0` in the `/etc/sysconfig/amf` file.

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 details, refer to the TechNote at:

https://www.veritas.com/content/support/en_US/article.100003972

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.

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

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

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

```
# installer -start node1 node2
```

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 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.

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.

Rolling upgrade not supported for upgrades from SF Oracle RAC 5.1 SP1 with fencing configured in `dmpmode`.

Rolling upgrade is not supported if you are upgrading from SF Oracle RAC 5.1 SP1 with fencing configured in `dmpmode`. This is because fencing fails to start after the system reboots during an operating system upgrade prior to upgrading SF Oracle RAC.

The following message is displayed:

```
VxVM V-0-0-0 Received message has a different protocol version
```

Workaround: Perform a full upgrade if you are upgrading from SF Oracle RAC 5.1 SP1 with fencing configured in `dmpmode`.

"Configuration must be ReadWrite : Use haconf -makerw" error message appears in VCS engine log when `hastop -local` is invoked (2609137)

A message similar to the following example appears in the `/var/VRTSvcs/log/engine_A.log` file when you run the `hastop -local` command on any system in a SF Oracle RAC cluster that has CFSMount resources:

```
2011/11/15 19:09:57 VCS ERROR V-16-1-11335 Configuration must be
ReadWrite : Use haconf -makerw
```

The `hastop -local` command successfully runs and you can ignore the error message.

Workaround: There is no workaround for this issue.

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.

vxdisk resize from slave nodes fails with "Command is not supported for command shipping" error (3140314)

When running the `vxdisk resize` command from a slave node for a local disk, the command may fail with the following error message:

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

Workaround: Switch the master to the node to which the disk is locally connected and run the `vxdisk resize` on that node.

CVR configurations are not supported for Flexible Storage Sharing (3155726)

Cluster Volume Replicator (CVR) configurations are not supported in a Flexible Storage Sharing environment.

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
```

SG_IO ioctl hang causes disk group creation, CVM node joins, and storage connects/disconnects, and vxconfigd to hang in the kernel (3193119)

In RHEL 5.x, the `SG_IO ioctl` process hangs in the kernel. This causes disk group creation and CVM node joins to hang. The `vxconfigd` thread hangs in the kernel during storage connects/disconnects and is unresponsive.

Workaround: This issue is fixed in RHEL 6.3. Upgrade to RHEL 6.3.

vxdbg adddisk operation fails when adding nodes containing disks with the same name (3301085)

On a slave node, when using the `vxdbg adddisk` command to add a disk to a disk group, and if the device name already exists in the disk group as disk name (disk media name), the operation fails with the following message:

```
VxVM vxdbg ERROR V-5-1-599 Disk disk_1: Name is already used.
```

Workaround: Explicitly specify the disk media name, which is different from the existing disk media name in the disk group, when running the `vxdbg adddisk` command on the slave node.

For example:

```
# vxdbg -g diskgroup adddisk dm1=diskname1 dm2=diskname2 dm3=diskname3
```

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

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

```
VxVM vxdbg 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.

vxconfigstore is unable to restore FSS cache objects in the pre-commit stage (3461928)

While restoring a Flexible Storage Sharing (FSS) disk group configuration that has cache objects configured, the following error messages may display during the pre-commit phase of the restoration:

```
VxVM vxcache ERROR V-5-1-10128 Cache object meta-data update error
VxVM vxcache ERROR V-5-1-10128 Cache object meta-data update error
VxVM vxvol WARNING V-5-1-10364 Could not start cache object
VxVM vxvol ERROR V-5-1-11802 Volume volume_name cannot be started
VxVM vxvol ERROR V-5-1-13386 Cache object on which Volume volume_name
is constructed is not enabled
VxVM vxvol ERROR V-5-1-13386 Cache object on which Volume volume_name
is constructed is not enabled
```

The error messages are harmless and do not have any impact on restoration. After committing the disk group configuration, the cache object and the volume that is constructed on the cache object are enabled.

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.

Intel SSD cannot be initialized and exported (3584762)

Initializing an Intel SSD with the Flexible Storage Sharing (FSS) export option may fail with the following error message:

```
VxVM vxedpart ERROR V-5-1-10089 partition modification failed: Device
or resource busy
```

Workaround:

Initialize the private region of the SSD disk to 0 and retry the disk initialization operation.

For example:

```
# dd if=/dev/zero of=/dev/vx/dmp/intel_ssd0_0 bs=4096 count=1

# vxdisksetup -i intel_ssd0_0 export
```

VxVM may report false serial split brain under certain FSS scenarios (3565845)

In a Flexible Storage Sharing (FSS) cluster, as part of a restart of the master node, internal storage may become disabled before network service. Any VxVM objects on the master node's internal storage may receive I/O errors and trigger an internal transaction. As part of this internal transaction, VxVM increments serial split brain (SSB) ids for remaining attached disks, to detect any SSB. If you then disable the network service, the master leaves the cluster and this results in a master takeover. In such a scenario, the master takeover (disk group re-import) may fail with a false split brain error and the `vxsplitlines` output displays 0 or 1 pools.

For example:

```
Syslog: "vxvm:vxconfigd: V-5-1-9576 Split Brain. da id is 0.2,
while dm id is 0.3 for dm disk5mirr
```

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
```

Storage Foundation for Databases (SFDB) tools known issues

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

Clone operations fail for instant mode snapshot (3916053)

For Oracle version 12.2.0.1.0, cloning a container database (CDB) fails for “instant mode” snapshots.

The cloning fails with the following error message:

```
SFDB vxsfadm ERROR V-81-0564 Oracle returned error.
Reason: ORA-01157: cannot identify/lock data file 5 - see DBWR trace file
ORA-01110: data file 5: '/data/DB12R2/pdbseed/system01.dbf'
```

Workaround: There is no workaround for this issue. Alternatively, you can use online or offline mode snapshots.

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.

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.

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 Veritas 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_vmccloneb` 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.

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

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

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

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

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

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

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

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

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


```
# umount /dev/shm
```
- 3 Mount the `/dev/shm` file system with the following options:


```
# mount -t tmpfs shmfs -o size=4096m /dev/shm
```
- 4 Start the database.

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

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

You might encounter an Oracle error such as the following:

```
/opt/VRTSdbed/bin/vxsfadm -s flashsnap -o clone
-a oracle -r dblxx64-16-v1 --flashsnap_name TEST11 --clone_path
```

```

/tmp/testRecoverdb --clone_name clone1
USERNAME:  oragrid
STDOUT:
Retrieving snapshot information ... Done
Importing snapshot diskgroups ... Done
Mounting snapshot volumes ... Done

```

ORA-01513: invalid current time returned by operating system

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

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

Workaround: Retry the cloning operation until it succeeds.

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 workaround at this point of time.

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'

```

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 12.1.0.1 or later, 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 12.1.0.1 or later, 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.

SFDB commands fail when an SFDB installation with authentication configured is upgraded to InfoScale 7.4.3 (3644030)

When you upgrade an SFDB installation in which authentication is configured, the SFDB commands fail, and a message similar to the following is logged:

```
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.

For details, refer to one of the following documents:

- *Veritas InfoScale Storage and Availability Management for Oracle Databases*
- *Veritas InfoScale Storage and Availability Management for DB2 Databases*

Benign message displayed upon execution of vxsfadm -a oracle -s filesnap -o destroyclone (3901533)

You may encounter the following message when you run the vxsfadm -a oracle -s filesnap -o destroyclone command:

Redundant argument in sprintf at /opt/VRTSdbed/lib/perl/DBED/Msg.pm line 170.

Eg:

```
vxsfadm -s filesnap -a oracle -o destroyclone --name file1
--clone_name c1n1
```

Redundant argument in sprintf at /opt/VRTSdbed/lib/perl/DBED/Msg.pm line 170.

Shutting down clone database...

Done

Destroying clone...

Done

You can ignore this message; it does not affect the functionality of InfoScale in any manner.

Application isolation feature known Issues

This section describes the known issues in this release for the Application isolation feature.

These known issues apply to the following product:

- Veritas InfoScale Enterprise

Addition of an Oracle instance using Oracle GUI (dbca) does not work with Application Isolation feature enabled

Addition of an Oracle instance using Oracle GUI (dbca) does not work when the Application Isolation feature is enabled.

Workaround:

You can use the equivalent CLI command for adding the Oracle instance.

Auto reattach of detached plexes may not happen for FSS disk groups when auto-mapping feature is used (3902004)

Auto reattach of detached plexes may not happen for FSS disk groups when you use the auto-mapping feature instead of the explicit disk export.

Workaround:

Use explicit disk export for creating FSS disk group.

CPI is not supported for configuring the application isolation feature (3902023)

You cannot configure the application isolation feature using the Common Product Installer (CPI).

Workaround:

Use manual steps for configuring the application isolation feature.

Refer to *Storage Foundation Cluster File System High Availability Administration Guide*.

Thin reclamation does not happen for remote disks if the storage node or the disk owner does not have the file system mounted on it (3902009)

Thin reclamation does not happen for remote disks if the storage node or the disk owner does not have file system mounted on it.

Workaround:

Mount the file system on the disk owner and perform thin reclamation.

Cloud deployment known issues

This section lists the known issues related to deploying the InfoScale products in cloud environment.

Systems in GCP may get stuck in the LEAVING state when multiple nodes are restarted a cascaded manner

When multiple nodes are restarted a cascaded manner, they may get stuck in the LEAVING state for some time. The CVM_clus resources may go into the UNABLE TO OFFLINE state and the CVM_vxconfigd resources may go into the FAULTED state. (3966612)

Workaround: When a system comes up after restart, if the cluster is inactive, start the cluster by running the `hastart` command manually. Alternatively, restart the nodes gracefully one after the other.

An error occurs during VVR or CVR configuration when alias IPs are assigned to GCP VM instances (3965275)

You need to assign alias IPs for a VVR or a CVR configuration. When alias IPs are assigned to VM instances, Google adds the entry of an alias IP against its host name in the `/etc/hosts` file. Therefore, the wrong IP address is resolved with the host name, and the "vradmind not reachable on cluster peer" error is logged.

Workaround: While configuring replication, use private IPs instead of host names.

In an Azure environment, the systems under InfoScale control may panic due to CPU soft lockup [3929534]

In an Azure environment, after you install any of the InfoScale product, you may observe that the systems panic due to the CPU soft lockup issue.

This issue occurs if you have installed the operating system (RHEL 7.1, 7.2, 7.3, and 7.4) using an image that is available on Azure marketplace.

For the mentioned OS versions, in the OS image available on the Azure marketplace, there is a mismatch between the supported and the available kernel version.

Workaround: Deploy the OS with supported kernel version.

To deploy the OS with supported kernel version, run the following command on the systems where the CPU soft lockup issue is observed:

```
$ az vm create -n VirtualMachineName -g rg name --image  
RedHat:RHEL:7.3:7.3.2017071923 --size VMSize --storage-sku
```

```
StorageAccountType --location Location --authentication-type password
--admin-username username --admin-password password

$ uname -r

3.10.0-514.26.2.el7.x86_64
```

In an Azure environment, an InfoScale cluster node may panic if any of the node is rebooted using Azure portal [3930926]

In an Azure environment, an InfoScale cluster node may panic if any of the node is rebooted using Azure portal.

This issue occurs if you have enabled I/O fencing for the cluster.

When you reboot any cluster node using Azure portal, the time required for reboot is more than the time taken by the fencing module to detect a network partition and to perform the fencing operation. As a result, the node that is rebooted or any other node in the cluster may panic.

Workaround: Increase the LLT peerinact timeout value to 120 seconds (2 mins). The default LLT peerinact timeout value is 16 seconds.

Increasing the LLT peerinact timeout value delays the fencing race during which time the node reboot operation is complete.

To increase the LLT peerinact timeout, run the following command on any of cluster node that is active:

```
lltconfig -T peerinact:12000
```

If you disable a public IP from the Azure portal, the corresponding AzureIP resource goes into UNKNOWN state [3928222]

When an AzureIP resource is ONLINE, if you disable the corresponding public IP from the Azure portal, the public IP resource ID is no longer associated with the Azure IP configuration. Therefore, the AzureIP resource goes into the UNKNOWN state.

Workaround:

1. On the Azure portal, search for the IP configuration by using the private IP.
2. Delete the IP configuration.
3. Bring the AzureIP resource online by using the following command:

```
# hares -online <azureip_resource_name> -sys <system_name>
```

After rolling upgrade phase 1, xprtld service fails to start on AWS instances (4004450)

When you perform a rolling upgrade of InfoScale 7.4.1 with Patch 1600 to InfoScale 7.4.2 on a RHEL 8.1 system in an AWS environment, the `xprtld` service fails to start after phase 1.

Workaround:

After the upgrade completes, perform the following steps to stop and start the `xprtld` service.

- 1 Stop the `xprtld` service.

```
# systemctl stop xprtld.service
```

- 2 Verify that the `xprtld` service is stopped.

```
# systemctl status xprtld.service
```

- 3 Start the `xprtld` service.

```
# systemctl start xprtld.service
```

- 4 After the service starts, verify that the service status is active (running).

```
# systemctl status xprtld.service
```

Issues related to Veritas InfoScale Storage in Amazon Web Services cloud environments

This section describes the known issues related to Veritas InfoScale Storage in Amazon Web Services cloud environments.

Incorrect media type displayed for AWS EC2 volumes

In AWS EC2 environments, the media type property of a disk cannot be accurately determined due to certain limitations. If media type is not determined correctly, manually set this property so that the disk is utilized to its full potential.

Workaround:

Update the setting manually as follows:

- 1 Set the media type property to SSD.

```
# vxdisk -f set disk_name mediatype=ssd
```

- 2 Verify the update.

```
# vxdisk -p list disk_name | grep ALERTS
ALERTS                               : media_mismatch
```

Inconsistencies in instance store volumes

The instance store volumes attached to an EC2 instance are not consistent across reboot, start, or stop operations on an instance. The attached volumes get attached to other instances once their current running instance is stopped. Thus, a different set of volumes get attached to the instance when the instance is restarted and disks appear invalid when the node starts.

Workaround: Initialize the disks using the following command:

```
# /etc/vx/bin/vxdisksetup -if disk_name
```

Stale remote disks on some nodes after failure of vxdisk unexport operation

When the disk unexport operation (`vxdisk unexport`) fails on a partial set of nodes, stale remote disks are left on a few nodes in the cluster even after the disk shows unexported on the node that has direct connectivity to the disk. If you encounter this issue, manually remove the stale remote disks.

Workaround:

Remove the stale remote disks manually:

1. Export the disk, which has stale remote disks. Run the command on the node that has direct connectivity to the disk.

```
# vxdisk export disk_name
```

2. Unexport the same disk so that the remote disks are removed from all nodes. Run the command on the node that has direct connectivity to the disk.

```
# vxdisk unexport disk_name
```

UDID of AWS volumes not updated after migration

The Unique Device Identifier (UDID) of AWS volumes are not automatically updated when the volume is migrated from one EC2 instance to another. The following command must be executed at the destination EC2 instance after attaching the volume.

Workaround: Update the UDID manually after the migration. Run the following command on the destination EC2 instance after attaching the volume.

```
# vxdisk updateudid disk_name
```

Partial detachment of volumes from AWS console

VxVM volumes detached from the AWS console continue to be listed as VxVM disks.

Workaround:

Perform the following steps to gracefully detach the volumes. Run the commands on the EC2 instance on which the volume is attached before detaching the volume from console.

1. Take the disk offline.

```
# vxdisk offline disk_name
```

2. Remove the disk from Volume Manager.

```
# vxdisk rm disk_name
```

3. Remove the path from the SCSI subsystem.

```
echo 1 > /sys/block/device_name/device/delete
```

Crash dump logs not available when EC2 instances crash

Crash dump logs are not available on the Amazon Web Services (AWS) Elastic Compute Cloud (EC2) instance if the instance panics. This is a known limitation with AWS. Therefore, for scenarios that involve system crashes, dumps and analysis, Veritas recommends that you contact Amazon to obtain the crash dump for the issue.

vxcloudd daemon fails with a core dump when the bucket name on the target exceeds 32 characters (3916980)

The maximum characters allowed for an S3 bucket name on the target is 32 characters. If the bucket name exceeds the limit, the `vxcloudd` daemon fails with a core dump.

Workaround: Restrict the number of characters in the S3 bucket name to 32 characters.

Migration of data to cloud volumes using S3 Connector fails with core dump (3915555)

Enforcing the policy to move data between the local volume and the cloud volume using the `fsppadm enforce` command fails causing the `vxcloud` daemon to terminate with a core dump and restart. This could be due to issues with OpenSSL. All write operations on the cloud volume fail.

Workaround: Run the `fsppadm enforce` command again as follows:

```
# fsppadm enforce mount_path_of_data_volume
```

Container deployment know issues

This section describes the known issues that are related to InfoScale deployments in container environments.

Creation of full sized snapshot of encrypted volume may lead to loss of data accessibility (4010594)

When you create a full sized snapshot of an encrypted volume and the I/O operation happens on the snapshot volume before the snapshot creation is complete, you may observe issue of data inaccessibility for the snapshot volume. However, the data on the original encrypted volume is intact.

Workaround:

Restart the `vxencryptd` daemon running on system using the following steps:

- 1 `systemctl stop vxvm-encrypt`
- 2 `systemctl start vxvm-encrypt`

After restoring space-optimized snapshot to new PVC, mount on restored PVC may fail if the source snapshot volume is detached (4012858)

CSI snapshot created using block level space-optimized (SO) snapshot functionality may get detached if:

- The underlying cache volume becomes full due to intensive writes on the source of the SO snapshot or the snapshot itself
- Autogrow fails due to some reason

When a PVC is restored from such SO snapshot where data from the source is not synchronized completely, the restored volume may not have complete data copy and may redirect the data read operations to the source SO snapshot.

When such restored PVCs are mounted, the FS metadata read operations that are redirected to the source fails if the source SO snapshot is detached. As a result, fsck may get invoked on the restored PVC.

Workaround:

To address this issue, veritas recommends one of the following actions:

- Use CSI space-optimized snapshot functionality for read-intensive applications
- Use full-instant snapshot or CSI clone functionality for write-intensive or read-write-update applications
- Manually set the values of the configurable parameters like `cachesize` to an appropriate value based on the application workload while creating CSI `volumesnapshotclass` object

CSI controller pod remains in 'Terminating' state in case of graceful node shutdown or power-off (4011482)

CSI controller pod does not failover when the node is powered off, rebooted, or shutdown gracefully using CLI. This results in the CSI controller pod remaining in the 'Terminating' state. This is an expected behavior.

Workaround:

To address this issue, veritas recommends one of the following actions:

- If a node must be kept shut down for certain period, to ensure availability, use the following command to drain the node before shutting it down:

```
# kubectl drain <node_name> --force --ignore-daemonsets  
--delete-local-data
```
- If you intend to delete the node from the Kubernetes cluster, delete the node object. In such case, you need not drain the node manually.

CSI node pods does not get rescheduled on other worker nodes when its parent node is drained (4011384)

After evacuating Kubernetes node using the `kubectl drain <HOST> --ignore-daemonsets -force` command, InfoScale CSI node pod that is running on evacuated node do not get rescheduled.

This issue happens because the `kubectl drain <HOST> --ignore-daemonsets -force` command delete the pods that are not managed by any controller. As the InfoScale CSI nodes are managed by daemonset controllers, these nodes cannot be deleted or rescheduled even with `--force` option.

Workaround: No workaround.

While restoring a snapshot, PVC goes into pending state after rebooting all nodes except the master node in the cluster (4014525)

While creating a snapshot in a single control-plane Kubernetes cluster environment, if the VxVM master is abruptly shut down, the associated volume goes into DISABLED or DETACHED state. If a snapshot of such volume in DISABLED or DETACHED state is created, the PVC goes into pending state.

Workaround: Delete the snapshot volume using the `vxedit` command. Kubernetes automatically reattempts to create a volume snapshot again.

InfoScale cluster on the worker node panics when network split occurs and the master node is disconnected from all the worker nodes. (4018090)

In a single master server configuration in Kubernetes environment, if the master server gets disconnected from the cluster due to LLT link failure or gets rebooted due to any other reason, all worker nodes in the cluster are rebooted. This disconnection does not cause data corruption but results in application downtime.

Workaround: Resolve the network issues between the master and the slave nodes in Kubernetes cluster.

When a node in the Kubernetes cluster is shutdown gracefully, fencing is not invoked (4011482)

When a Kubernetes node is shutdown gracefully, the kubelet running on that node also exits gracefully without draining the already running pods on the node. It does not wait for the pod to move to terminated state as that may disturb the normal system shutdown process. Due to this uncertainty, the cluster does not restart the statefulset pods that are scheduled on the node that was shut down, rather it keeps the stateful pods in terminating state till the node restarts.

This is an expected behaviour.

Workaround: To address this issue, veritas recommends one of the following actions:

- If a node must be kept shut down for certain period, to ensure availability, use the following command to drain the node before shutting it down:

```
# kubectl drain <node_name> --force --ignore-daemonsets
```
- If you intend to delete the node from the Kubernetes cluster, delete the node object. In such case, you need not drain the node manually.

In case of storage failure, application I/Os to the mountpoint inside container fails and pod goes into CreateContainerConfigError or Error state (4011219, 4014758, 4015259)

If application pod is using PVC for persistent storage and if storage failure occurs, then the application I/Os to the corresponding mountpoint inside the container fails and the pod goes into the CreateContainerConfigError or Error state. This issue occurs because Kubernetes and the CSI plugin does not detect any storage failure to reschedule the pod. Therefore, even if the storage failure is resolved, the application pod remains in the same state as the old disabled filesystem does not get remounted.

Workaround: Manually restart the application pod after the storage failure is resolved.

Pod gets stuck in 'Terminating' state and file system remains mounted (4017013)

If stateful applications mount a large number of InfoScale PVCs on a node, then during operations like node drain, all the volumes are unmounted in parallel. In such a case, some of the pods may remain in the 'Terminating' state and prevent the InfoScale volumes from getting unmounted.

Workaround: Manually unmount the InfoScale volumes from the host and then forcefully delete the pod that is stuck in the terminating state.

When more than 20 containers are started simultaneously on a worker node, some of these containers may fail to start (4015069)

When more than 20 containers are started simultaneously, the start-up activity may happen in multiple availability containers at the same time and may limit the availability of the required resources. This condition may cause the HAD to be unresponsiveness and may prevent some of the containers to start.

Workaround: Start containers in batches, only a few containers at a time. Do not start many containers simultaneously.

While stopping InfoScale on a Kubernetes worker node, the InfoScale cluster gets stuck in the LEAVING state (4018544)

The containers running on a worker node may not be a part of the InfoScale cluster configuration on that host. When a InfoScale cluster is stopped, the containers on the worker nodes continue to use the CVM storage and prevents CVM from getting offline. Consequently, the InfoScale cluster gets stuck in the LEAVING state.

Workaround: Manually stop the containers on the worker nodes before stopping the InfoScale cluster.