

Veritas InfoScale™ 7.3.1 Release Notes - AIX

Last updated: 2018-08-31

Legal Notice

Copyright © 2018 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
500 E Middlefield Road
Mountain View, CA 94043

<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:

doc.feedback@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	12
	About this document	12
Chapter 2	Changes introduced in 7.3.1	14
	Changes related to installation and upgrades	14
	Change in upgrade path	14
	Changes related to the Cluster Server engine	14
	VCS stop timeout	15
	256-bit encryption for enhanced security	15
	Changes related to Cluster Server agents	15
	New attributes for Cluster Server agents	15
	Changes related to Veritas File System	15
	New option [-i] included in <code>fsadm</code> command to exclude the actively used files during file system reorganization	15
	Delayed allocation support extended to clustered file systems	16
	Support for migrating Oracle database from Oracle ASM to Veritas File System (VxFS)	17
	Changes related to replication	17
	Veritas Volume Replicator Performance Improvements	17
Chapter 3	System requirements	19
	VCS system requirements	19
	Supported AIX operating systems	20
	Storage Foundation for Databases features supported in database environments	20
	Storage Foundation memory requirements	21
	Supported database software	21
	Supported hardware and software	22
	Number of nodes supported	22
	Required attributes of LUNs for DMP devices	22

Chapter 4	Known Issues	24
	Issues related to installation and upgrade	24
	Switch fencing in enable or disable mode may not take effect if VCS is not reconfigured [3798127]	25
	In an upgraded cluster, security configuration may fail while importing <code>VCS_SERVICES</code> file. [3708929]	25
	During an upgrade process, the <code>AMF_START</code> or <code>AMF_STOP</code> variable values may be inconsistent [3763790]	25
	Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups (2574731)	26
	If you have a shared (system) WPAR configured, when you install, upgrade, or uninstall any Veritas product, the filesets in the WPAR are not synchronized correspondingly (3313690)	26
	NetBackup 6.5 or older version is installed on a VxFS file system (2056282)	26
	The <code>VRTSvxvm</code> fileset fails to install on a few cluster nodes because the template file is corrupted (2348780)	27
	After a locale change restart the <code>vxconfig</code> daemon (2417547, 2116264)	28
	The <code>uninstallmr</code> script must not be used after you install InfoScale 7.3.1 on AIX (3931949)	28
	Resource faults during Rolling upgrade due to perl changes (3930605)	28
	Echo or print messages defined in any of the <code>.rc</code> file leads to failure in CPS configuration (3944034)	28
	Storage Foundation known issues	29
	Dynamic Multi-Pathing known issues	29
	Veritas Volume Manager known issues	30
	Veritas File System known issues	42
	Replication known issues	46
	After the product upgrade on secondary site, replication may fail to resume with "Secondary SRL missing" error [3931763]	46
	<code>vradmin repstatus</code> command reports secondary host as "unreachable"(3896588)	47
	RVGPrimary agent operation to start replication between the original Primary and the bunker fails during failback (2036605)	47

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]	48
In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon (1672410, 1672417)	49
vxassist relayout removes the DCM (145413)	49
vradmin functionality may not work after a master switch operation [2158679]	49
Cannot relayout data volumes in an RVG from concat to striped-mirror (2129601)	49
vradmin verifydata operation fails when replicating between versions 5.1 and 6.0 or later (2360713)	50
vradmin verifydata may report differences in a cross-endian environment (2834424)	51
vradmin verifydata operation fails if the RVG contains a volume set (2808902)	51
Bunker replay does not occur with volume sets (3329970)	51
SmartIO does not support write-back caching mode for volumes configured for replication by Volume Replicator (3313920)	51
During moderate to heavy I/O, the vradmin verifydata command may falsely report differences in data (3270067)	52
The vradmin repstatus command does not show that the SmartSync feature is running [3343141]	52
While vradmin commands are running, vradmin may temporarily lose heartbeats (3347656, 3724338)	52
Write I/Os on the primary logowner may take a long time to complete (2622536)	53
DCM logs on a disassociated layered data volume results in configuration changes or CVM node reconfiguration issues (3582509)	53
After performing a CVM master switch on the secondary node, both links detach (3642855)	54
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)	54

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)	54
DCM plex becomes inaccessible and goes into DISABLED(SPARSE) state in case of node failure. (3931775)	55
Cluster Server known issues	56
Operational issues for VCS	56
Issues related to the VCS engine	57
Issues related to the bundled agents	64
Issues related to the VCS database agents	68
Issues related to the agent framework	72
Cluster Server agents for Volume Replicator known issues	76
Issues related to Intelligent Monitoring Framework (IMF)	76
Issues related to global clusters	79
Issues related to the Cluster Manager (Java Console)	80
VCS Cluster Configuration wizard issues	80
LLT known issues	80
I/O fencing known issues	81
Storage Foundation and High Availability known issues	87
Cache area is lost after a disk failure (3158482)	87
In an IPv6 environment, db2icrt and db2idrop commands return a segmentation fault error during instance creation and instance removal (1602444)	88
Oracle 11gR1 may not work on pure IPv6 environment (1819585)	89
Not all the objects are visible in the VOM GUI (1821803)	89
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)	90
A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a dynamic storage tiering placement policy (1880081)	90
Upgrading operating system Technology Levels along with Storage Foundation using an alternate disk fails (2162945)	90
Storage Foundation Cluster File System High Availability known issues	91
In an FSS environment, creation of mirrored volumes may fail for SSD media [3932494]	91
Mount command may fail to mount the file system (3913246)	91

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)	92
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)	93
DG creation fails with error "V-5-1-585 Disk group punedatadg: cannot create: SCSI-3 PR operation failed" on the VSCSI disks (3875044)	93
Write back cache is not supported on the cluster in FSS scenario [3723701]	94
CVMVOLDg agent is not going into the FAULTED state. [3771283]	94
CFS commands might hang when run by non-root (3038283)	94
Inode access and modification times are not getting updated on the primary node when a file owned by the primary node is accessed from a secondary node (2170318)	95
The fsappadm subfilemove command moves all extents of a file (3258678)	95
Certain I/O errors during clone deletion may lead to system panic. (3331273)	96
Panic due to null pointer de-reference in vx_bmap_lookup() (3038285)	96
In a CFS cluster, that has multi-volume file system of a small size, the fsadm operation may hang (3348520)	96
Storage Foundation for Oracle RAC known issues	96
Oracle RAC known issues	96
Storage Foundation Oracle RAC issues	97
Storage Foundation for Databases (SFDB) tools known issues	103
Sometimes SFDB may report the following error message: SFDB remote or privileged command error (2869262)	104
SFDB commands do not work in IPV6 environment (2619958)	104
The database clone operation using the vxsfadm -o clone(1M) command fails (3313715)	104
In an off-host scenario, a clone operation may fail with an error message (3313572)	105
When you attempt to move all the extents of a table, the dbdst_obj_move(1M) command fails with an error (3260289)	105

Attempt to use SmartTier commands fails (2332973)	106
Attempt to use certain names for tiers results in error (2581390)	106
Clone operation failure might leave clone database in unexpected state (2512664)	106
Clone command fails if PFILE entries have their values spread across multiple lines (2844247)	107
Clone fails with error "ORA-01513: invalid current time returned by operating system" with Oracle 11.2.0.3 (2804452)	107
Data population fails after datafile corruption, rollback, and restore of offline checkpoint (2869259)	108
Flashsnap clone fails under some unusual archivelog configuration on RAC (2846399)	108
Database Storage Checkpoints created by using <code>dbed_ckptcreate</code> may not be visible after upgrading to 7.3.1 (2626248)	108
Cloning of a container database may fail after a reverse resync commit operation is performed (3509778)	109
If one of the PDBs is in the read-write restricted state, then cloning of a CDB fails (3516634)	109
Cloning of a CDB fails for point-in-time copies when one of the PDBs is in the read-only mode (3513432)	110
If a CDB has a tablespace in the read-only mode, then the cloning fails (3512370)	110
If any SFDB installation with authentication setup is upgraded to 7.3.1, the commands fail with an error (3644030)	110
Error message displayed when you use the <code>vxsfadm -a oracle -s filesnap -o destroyclone</code> command (3901533)	111

Chapter 5	Software Limitations	112
	Storage Foundation software limitations	112
	Dynamic Multi-Pathing software limitations	112
	Veritas Volume Manager software limitations	114
	Veritas File System software limitations	115
	SmartIO software limitations	116
	Replication software limitations	117
	VVR Replication in a shared environment	117
	VVR IPv6 software limitations	117
	VVR support for replicating across Storage Foundation versions	117
	Cluster Server software limitations	117
	Limitations related to bundled agents	118
	Limitations related to VCS engine	120

Veritas cluster configuration wizard limitations	120
Limitations related to IMF	121
Limitations related to the VCS database agents	121
Systems in a cluster must have same system locale setting	121
Limitations with DiskGroupSnap agent [1919329]	121
Virtualizing shared storage using VIO servers and client partitions	122
Cluster Manager (Java console) limitations	124
The operating system does not distinguish between IPv4 and IPv6 packet counts	125
A service group that runs inside of a WPAR may not fail over when its network connection is lost	125
Limitations related to LLT	125
Limitations related to I/O fencing	126
Limitations related to global clusters	128
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]	129
Storage Foundation Cluster File System High Availability software limitations	129
cfsmntadm command does not verify the mount options (2078634)	129
Upgrade of secure clusters not supported using native operating system tools	129
Stale SCSI-3 PR keys remain on disk after stopping the cluster and deporting the disk group	129
Unsupported FSS scenarios	130
Storage Foundation for Oracle RAC software limitations	130
Supportability constraints for normal or high redundancy ASM disk groups with CVM I/O shipping and FSS (3600155)	130
Limitations of CSSD agent	130
Oracle Clusterware/Grid Infrastructure installation fails if the cluster name exceeds 14 characters	131
Policy-managed databases not supported by CRSResource agent	131
Health checks may fail on clusters that have more than 10 nodes	131
Cached ODM not supported in Veritas InfoScale environments	131
Storage Foundation for Databases (SFDB) tools software limitations	131
Parallel execution of <code>vxsfadm</code> is not supported (2515442)	131

Creating point-in-time copies during database structural changes is not supported (2496178)	132
Oracle Data Guard in an Oracle RAC environment	132
Operating system limitations	132
Upgrading from AIX 7.1 TL4 SP2 to SP3 can cause agents to fault (3916650)	132
Index	133

Introduction

This chapter includes the following topics:

- [About this document](#)

About this document

This document provides information that is specific to version 7.3.1 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 AIX products:

- *Veritas InfoScale Getting Started Guide*
- *Veritas InfoScale What's New 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.100040420.html

Changes introduced in 7.3.1

This chapter includes the following topics:

- [Changes related to installation and upgrades](#)
- [Changes related to the Cluster Server engine](#)
- [Changes related to Cluster Server agents](#)
- [Changes related to Veritas File System](#)
- [Changes related to replication](#)

Changes related to installation and upgrades

The following changes are introduced to the installation and upgrading of Veritas InfoScale 7.3.1.

Change in upgrade path

You can upgrade to Veritas InfoScale 7.3.1 only if the base version of your currently installed product is 6.1.1 or later.

Changes related to the Cluster Server engine

The following sections describe the changes that are introduced in the Cluster Server engine.

VCS stop timeout

An environment variable, `VCS_STOP_TIMEOUT`, is introduced, which prevents the VCS stop operation from being hung during system shutdown or restart.

Default value (seconds): 0

256-bit encryption for enhanced security

The Cluster Server component now provides 256-bit encryption for enhanced security. The `vcseencrypt` utility, which you use to generate encrypted VCS and agent passwords, now generates 256-bit encrypted passwords.

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

Changes related to Cluster Server agents

The following sections describe the changes that are introduced in the Cluster Server agents.

New attributes for Cluster Server agents

The following new attributes are introduced in this release:

DiskGroup agent

Attribute	Description
Name: DGOptions	Specifies the options for the <code>vxdg import</code> command. The agent uses this attribute only while bringing a DiskGroup resource online.
Type: string	
Dimension: scalar	For more information, see the <i>vxdg(1m)</i> manual page.
	Example: <code>-o noautostart -o updateid</code>

Changes related to Veritas File System

The following changes are introduced to Veritas File System (VxFS) of Veritas InfoScale 7.3.1.

New option `[-i]` included in `fsadm` command to exclude the actively used files during file system reorganization

A new option `[-i]` is added to the `fsadm` command, to address the reorganization of actively used files.

The updated `fsadm` command is as follows:

```
fsadm [-t vxfs] [-e] [-d] [-E] [-D] [-H] [-i] [-r rawdev] mount_point
```

The `[-i]` option indicates that the files that are accessed in last 60 seconds must not be reorganized. During the file system reorganization, if an application is actively accessing a particular file or has accessed it in last 60 seconds, then the file system reorganization does not affect the read-write operation on that file and the file will not be reorganized.

Delayed allocation support extended to clustered file systems

The delayed allocation capability for extending writes on a file system was available for local mounts. This capability is now extended for clustered file systems (CFS mounts).

With this support, depending on an application I/O size, instead of allocating a single block for every write operation a clustered file system will now allocate multiple blocks in a single instance. The delayed allocation thus reduces the file system fragmentation.

When an application I/O is received, the delayed allocation capability enables the clustered file system to split the write operation in to the following sequence:

1. Reserve a disk space

When an application I/O is received, the file system first reserves a disk space and the data is cached.

2. Allocate extents

After a disk space is reserved, a scheduler allocates disk blocks at the background and the file system then combines multiple block allocation requests to allocate extents.

The delayed extent allocation thus helps to avoid file system fragmentation and keeps the extent contiguous even if several files grow at the same time.

The delayed allocation is not dependent on the file system disk layout version and is disabled by default. You can enable delayed allocation using the `vxtunefs` command. You can display the delayed allocation range in the file by using the `fsmap` command.

See the `vxtunefs(1M)` and `fsmap(1M)` manual pages.

Notes:

- Delayed allocation must be disabled in cases where the data must be immediately written to the disk. For example, direct I/O, concurrent I/O, FDD/ODM access, and synchronous I/O.

- Delayed allocation is not supported on memory-mapped files, BSD quotas, and shared mount points in a Cluster File System (CFS).
- When BSD quotas are enabled on a file system, delayed allocation is turned off automatically for that file system.

Support for migrating Oracle database from Oracle ASM to Veritas File System (VxFS)

(Linux and Solaris only) Veritas InfoScale supports real-time migration of:

- (Linux only) Standalone and RAC Oracle databases, hosted on Oracle ASM disks to VxFS systems.
- (Solaris only) Standalone Oracle ASM disks to VxFS system and is mounted on the VxVM disks.

The migration requires a source system on which the primary database is hosted on Oracle ASM disks, and target serves as standby during the migration. That target contains VxVM disks on which the Veritas File System is mounted. The target disks can be on the same host as the primary database or on a different host. In Linux environment, you can migrate multiple instances of database at a time in a RAC environment.

For more information, see *Veritas InfoScale Solution's Guide*.

Changes related to replication

The following changes are introduced to replication of Veritas InfoScale 7.3.1.

Veritas Volume Replicator Performance Improvements

Previously TCP was one of the replication protocols used, where one data structure was used to transport the data. This data structure had eight connections which used the bandwidth. These eight connections can be increased and tuned to 16 or 32 to use more bandwidth. However, with one data structure the bandwidth utilization would have very little effect on latency between the Primary host and Secondary host.

Now in this 7.3.1 release, to improve the bandwidth utilization, TCP connections are tuned to multiple data structures to achieve better performance. Every data structure has one transport connection which is set as default, and available all the time. Currently there are 16 data structures available and enabled which are intended to use in parallel. Hence all the data is transited in parallel through these data structure connections, and distributes the network load. If more bandwidth is

required, you can tune the existing data structure in the multiple of 16. See *Veritas InfoScale™ Replication Administrator's Guide* for more information.

System requirements

This chapter includes the following topics:

- [VCS system requirements](#)
- [Supported AIX operating systems](#)
- [Storage Foundation for Databases features supported in database environments](#)
- [Storage Foundation memory requirements](#)
- [Supported database software](#)
- [Supported hardware and software](#)
- [Number of nodes supported](#)
- [Required attributes of LUNs for DMP devices](#)

VCS system requirements

This section describes system requirements for VCS.

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

VCS supports an environment where a few nodes in the cluster are hosted on LPARs with storage and network connectivity presented to the OS using VIOS. The remaining nodes in the cluster are hosted on physical systems with storage and network connectivity presented to the OS directly. However SCSI3 I/O fencing will be supported in this environment only if storage is available through NPIV in the LPARs. If NPIV is not available in LPARs, non-SCSI3 fencing is supported.

VCS requires that all nodes in the cluster use the same processor architecture and all nodes in the cluster must run the same VCS version. However, the nodes can have different versions of the supported operating system.

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

Table 3-1 Supported operating systems

Operating systems	Levels	Chipsets
AIX 7.2	TL0, TL1	Power 6, Power 7, Power 8
AIX 7.1	TL4	Power 6, Power 7, Power 8

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

Storage Foundation for Databases features supported in database environments

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

Table 3-2 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
Quick I/O	Yes	Yes	Yes	Yes
Cached Quick I/O	Yes	Yes	Yes	Yes
Concurrent I/O	Yes	Yes	Yes	Yes
Storage Checkpoints	Yes	Yes	Yes	Yes
Flashsnap	Yes	Yes	Yes	Yes
SmartTier	Yes	Yes	Yes	Yes

Table 3-2 SFDB features supported in database environments (*continued*)

Storage Foundation feature	DB2	Oracle	Oracle RAC	Sybase
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.

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

For 6.2 and earlier versions: <http://www.veritas.com/docs/000002658>

For 7.0 and later versions: <http://www.veritas.com/docs/000115952>

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 database, see the following TechNote: <http://www.veritas.com/docs/000002658>

Additionally, see the following Oracle support site for information on patches that may be required by Oracle for each release. <https://support.oracle.com>

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/DocumentBrowsing.html?product=InfoScale & Storage Foundation](https://www.veritas.com/content/support/en_US/DocumentBrowsing.html?product=InfoScale&StorageFoundation)

Select 7.3.1 in the **Version** filter and select **Compatibility Lists** under **Document Type**.

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.

In addition to the compatibility list, a Late Breaking News (LBN) is available for the latest updates, patches, and software issues regarding this release: https://www.veritas.com/support/en_US/article.100040109.html

Number of nodes supported

Veritas InfoScale supports cluster configurations up to 64 nodes.

SFHA, SFCFSHA, SF Oracle RAC: Flexible Storage Sharing (FSS) only supports cluster configurations with up to 8 nodes.

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

Required attributes of LUNs for DMP devices

When the `reserve_policy=single_path` and `reserve_lock=yes`, the SCSI-2 reserve may be placed on the device, which affects I/O load balancing and performance. To prevent the impact to load balancing and performance, set the `reserve_policy=no_reserve` and `reserve_lock=no` for the devices that are managed by DMP.

These settings are also required for a cluster set-up.

Set the following attributes for LUNs

- 1 Set the following attributes:
 - If the path has the `reserve_policy` attribute set, change the `reserve_policy` attribute to `no_reserve` for all the paths.

```
# lsattr -El hdisk557 | grep res
```

```
reserve_policy single_path
```

```
Reserve Policy True
```

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

```
hdisk557 changed
```

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

```
# lsattr -El hdisk558 | grep reserve_lock
```

```
reserve_lock yes
```

```
Reserve Device on open True
```

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

```
hdisk558 changed
```

- 2 Reboot the system for the changes to take effect.

Known Issues

This chapter includes the following topics:

- [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](#)

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

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 .

In an upgraded cluster, security configuration may fail while importing VCS_SERVICES file. [3708929]

When the cluster is upgraded from 6.0.1 to 6.2.1, the utility `atutil` is upgraded to 6.2.1, while the file of `VCS_SERVICES` remains as version 6.0.1. When a node is added to a cluster, the security is configured for the newly added node.

The file of `VCS_SERVICES` required by the security configuration is taken from the nodes which are already present in the cluster. In this case, importing information from `VCS_SERVICES` fails.

Workaround: If a new node is to be added, do the following:

- 1 Disable the security.
- 2 Disable the security.
- 3 Enable the security.

When the security is enabled, `VCS_SERVICES` for version 6.2.1 is used and security configurations works.

But since you reconfigure the security on the cluster with new keys, all credentials (certificates) are newly created. And any trust already present needs to be re-established.

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
```

If you have a shared (system) WPAR configured, when you install, upgrade, or uninstall any Veritas product, the filesets in the WPAR are not synchronized correspondingly (3313690)

On AIX, if you have a shared (system) workload partition (WPAR) configured, when you perform an install, upgrade, or uninstall task on any Veritas product by the Veritas product installer, the filesets may not be installed, upgraded, or uninstalled correspondingly.

Workaround: After an install, upgrade, or uninstall task, execute the following command to synchronize your WPAR with global systems:

```
# /usr/sbin/syncwpar -A
```

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 7.3.1, if you unmount all VxFS file systems including the one that hosts the NetBackup binaries (`/usr/openv`), then while upgrading to SF 7.3.1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure filesets `VRTSspbx`, `VRTSat`, and `VRTSiscso`. This causes NetBackup to stop working.

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

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

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

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

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

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

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

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

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

After a locale change restart the vxconfig daemon (2417547, 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."

The `uninstallmr` script must not be used after you install InfoScale 7.3.1 on AIX (3931949)

During the installation of InfoScale 7.3.1 on AIX, the CPI creates the `uninstallmr` script under `/opt/VRTS/install/`. This script is intended to roll back the changes that are related to a minor release of InfoScale, when you upgrade to a major release.

Hence, if you execute the script after you install InfoScale 7.3.1, it stops the product and mistakenly tries to uninstall the product patches, if any.

Workaround: Remove the script to eliminate the possibility of mistakenly using it.

To remove `uninstallmr`, execute the following command:

```
#rm uninstallmr
```

Resource faults during Rolling upgrade due to perl changes (3930605)

During rolling upgrade in 7.3.1, resources can go into FAULTED state due to Perl changes.

During rolling upgrade, after Phase I upgrade on the first sub cluster is complete and Phase I upgrade on the second sub cluster is about to start, if the `VCS Service Groups` is switched back to the first sub cluster, resources go into PARTIAL , STARTING, or FAULTED state.

Workaround: Complete the upgrade process, the resources become online after the upgrade.

Echo or print messages defined in any of the .rc file leads to failure in CPS configuration (3944034)

In Linux or UNIX, you can create `.rc` files (run commands files) for various utilities, which gets invoked when the utility is used.

For example:

- All commands written in `kshrc` file are executed when you use `/bin/ksh` based scripts
- All commands written in `vimrc` file are executed when you use vim editor to view or edit files

If the file `kshrc` file is modified to print any custom messages using `echo` or `print`, then every time a KSH based script is called, all the custom messages from the `kshrc` file are also printed on the console.

For CPS configurations, some KSH based scripts are executed. These scripts gather and print some information from the system. If some custom `echo` or `print` messages are already set in the `kshrc` file, the CPS KSH script not only prints its own output, but also prints the `echo` or `print` messages defined in the `kshrc` file.

This leads to inconsistent output and consequently failure in CPS configuration as it cannot parse the output as expected.

Workaround: Make sure that the `.rc` file does not contain any `echo` or `print` messages.

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).

Migration of root disks under DMP control may fail with "VxVM vxdmadm ERROR V-5-1-15253 bosboot" error [3930243]

In a scenario where DMP is enabled for native devices, migration of root disks to SAN disks may fail with the following error:

```
VxVM vxdmadm ERROR V-5-1-15253 bosboot would not succeed, please run
/usr/sbin/bosboot -vd /dev/hdisk1007 manually to find the cause of failure
VxVM vxdmadm ERROR V-5-1-15251 bosboot check failed
```

During the migration process, DMP attempts to connect to the `/usr/lib/drivers/vxdmp.6.1` file which is a device driver file from an earlier supported DMP version. The file does not exist in the latest release and this causes the failure in the migration process.

Workaround: Perform the following steps to resolve the issue:

1. Find the cause of failure:
`/usr/sbin/bosboot -vd /dev/hdiskx1`
2. Create a dummy DMP driver file
`"echo "1" > /usr/lib/drivers/vxdmp.6.1"`
3. Enable DMP support attribute
`vxdmpadm settune dmp_native_support=on`

vxdmpadm exclude ctrl=emcp command doesn't exclude PowerPath devices properly [3741636]

When devices are under the PowerPath (PP) control, the following command might not exclude all devices from the Veritas Volume Manager (VxVM) view.

```
# vxdmpadm exclude ctrl=emcp
```

Workaround: If devices are under PP control, you are not recommended to exclude devices using the `ctrl` option.

To exclude the devices, use any of the following commands:

```
# vxdmpadm exclude dmpnodename=<dmp-device-name>
```

```
# vxdmpadm exclude path=<path-name>
```

```
# vxdmpadm exclude product=<VID:PID>
```

Veritas Volume Manager known issues

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
```

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 dgname -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 more details, refer to the tech note:

https://www.veritas.com/support/en_US/article.TECH227660.

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

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

Workaround:

To solve this issue, use any of the workarounds:

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

OR

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

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


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.

Mounting CFS under VVR may fail, after rolling upgrade phase 1 on one node. [3764652]

After rolling upgrade on phase 1, if you run `mount cfs mount` all on one of the nodes of the CFS cluster, the mount operation hangs or fails. It's because there is some I/O pending on volumes.

Workaround:

There is no workaround.

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 relay layout operation (3492350)

If a relay layout 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.

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.

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
```

Co-existence check might fail for CDS disks

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

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 dgroupname
```

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
```

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.

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

```
# vxddm 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.

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.

Mksysb restore fails if physical volumes have identical PVIDs (3133542)

When you have multiple paths to the rootvg devices, restoring a `mksysb` backup file fails with the following error:

```
0516-1775 /usr/sbin/varyonvg: Physical volumes diskname1 and diskname2 have identical PVIDs.
```

This error is caused by an issue with IBM AIX.

Workaround:

Contact IBM support to obtain the fix. Refer to IBM APAR IV25286 for more details.

`vxconfigd` daemon hangs when Veritas InfoScale Storage or Veritas InfoScale Enterprise is run on AIX7.2SP1 or any earlier version (3901325)

When Veritas InfoScale Storage or Veritas InfoScale Enterprise is run on AIX7.2SP1 (or any earlier version), a unique product hang issue occurs because of the `vxconfigd` process getting hung. The `vxdisk` command, `vxdg` command, and `vxprint` command do not respond. In a clustered environment, if this issue hits on any one of the nodes, the commands may get hung on entire cluster nodes as well. The `vxconfigd` process that call IBM ODM-related interface functions hangs with the following stack trace

```
(0)> f 2658
pvthread+0A6200 STACK:
Use current context [F000000030019600] of cpu 4
[0065D0B0]kwpar_getmrc+000050 (F1000A0150520420 [??])
[0013EB30]sigprocmask+0004D0 (??, ??, ??)
[00003938]syscall+000230 ()
[D030635C]system+00019C (??)
[D06A6968]clr_disk_odm_owner+0000E8 (??, ??, ??)
[100E3978]devintf_clr_disk_odm_owner+000118 (00280018)
[100E3AE4]devintf_unstamp_odm+0000C4 (30337DD4)
[100CC904]req_disk_op+000264 (30246040, 30189328)
[10048DA8]request_loop+0011E8 ()
[10002BE0]main+001620 (00000002, 2FF22C50)
[10000208]__start+000068 ()
[kdb_read_mem] no real storage @ FFFFFFFF9240
(0)>
```

This occurs because of an issue with the IBM AIX's Object Data Manager (ODM).

Workaround: IBM has fixed this issue by APAR IV80412. All the InfoScale produce nodes are required to have this fix, either by installing IBM's APAR IV80412, or by upgrading the system to AIX7.2 SP2, which already has this APAR fix included in it.

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.

Veritas File System known issues

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

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.

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.

The fsppadm subfilemove command moves all extents of a file [3760225]

This issue occurs when you run the fsppadm subfilemove command from a cluster file system (CFS) secondary node. Then 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(1M) 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
```

dchunk_enable does not get set through vxtunefs in AIX (3551030)

The vxtunefs command uses the oslevel -s command to get OS TL level, but the OS level reports the lowest level of any installed AIX fileset on a system. Therefore, the oslevel command provides inaccurate OS level which may break the conditional logic inside VxFS.

Workaround:

To check which fileset is at a lower level and upgrade it to the recommended level

- 1 This command shows which fileset is at a lower SP level:

```
#oslevel -s -l `oslevel -sq 2>/dev/null | sed -n '1p'`
```

- 2 This command shows which fileset is at a lower TL level :

```
#oslevel -r -l `oslevel -rq 2>/dev/null | sed -n '1p'`
```

- 3 Upgrade the lower TL level fileset.

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

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

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

This issue is verified with the following commands:

- `cp -r`
- `du`

However, this issue might occur with other commands.

Workaround: Run the command from a different directory.

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

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

Workaround: Workaround is not available.

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.

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

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

Workaround: Unexport the file system prior to unmounting.

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

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

Workaround:

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

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

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

You cannot have in-place compress/uncompress rules and relocate compress/uncompress rules in the same policy file. If they are in the same file, file relocation is unpredictable.

Workaround: Create a different policy file for each policy, and enforce the policy as per the required sequence.

The file system may hang when it has compression enabled (3331276)

In a VxFS file system that has compression enabled, a deadlock in page fault handler can lead to the file system hang.

Workaround:

There is no workaround for this issue.

Unaligned large reads may lead to performance issues (3064877)

On AIX, when there are unaligned large reads, there may be a performance degradation.

Workaround:

There is no workaround for this issue.

Replication known issues

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

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

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

Workaround:

To resolve this issue

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

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

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

For example:

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


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

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

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

vxassist relayout removes the DCM (145413)

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

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

vradmin functionality may not work after a master switch operation [2158679]

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

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

Workaround:

To restore vradmin functionality after a master switch operation

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

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

2 Re-enter the command that failed.

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

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

Workaround:

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

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


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


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


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


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


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


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

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

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

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

Workaround: There are two workarounds for this issue.

- Upgrade the hosts running Storage Foundation 5.1 to Storage Foundation 6.0 or later and re-run the `vradmin verifydata` command.
- Follow the offline verification procedure in the "Verifying the data on the Secondary" section of the *Storage Foundation and High Availability Solutions Replication Administrator's Guide*. This process requires ensuring that the secondary is up-to-date, pausing replication, and running the `vradmin syncrvg` command with the `-verify` option.

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

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

vradmin verifydata operation fails if the RVG contains a volume set (2808902)

In a VVR environment, the `vradmin verifydata` command fails with the following error if the replicated volume group (RVG) contains any volume set:

```
Message from Primary:
VxVM VVR vxrsync ERROR V-5-52-2009 Could not open device
/dev/vx/dsk/vvrdg/<volname> due to: stat of raw character volume path
failed
```

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 7.3.1 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*.

The `vradmin repstatus` command does not show that the SmartSync feature is running [3343141]

In a Volume Replicator (VVR) environment, after you start the initial synchronization with the `vradmin -a startrep` command with file system mounted on the primary data volumes, the `vradmin repstatus` command does not show that the SmartSync feature is running. This is an only issue with the output of the `vradmin repstatus` command.

Workaround:

To confirm that SmartSync is running, enter:

```
vxrlink status rlink
```

While `vradmin` commands are running, `vradmind` may temporarily lose heartbeats (3347656, 3724338)

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

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

Workaround:

To resolve this issue:

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

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

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

Write I/Os on the primary logowner may take a long time to complete (2622536)

Under a heavy I/O load, write I/Os on the Volume Replicator (VVR) primary logowner take a long time to complete.

Workaround:

There is no workaround for this issue.

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.

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 `vxrvrg -g dg -P snap_prefix snapdestroy rvg` command. Clear the application service group and bring it back online manually.

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

Issue 1:

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

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

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

Issue 2:

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

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

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

Workaround: The following workarounds resolve these issues.

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

For example:

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

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

For example:

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

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=dcn  
<local_disks_across_different_nodes>
```

2. Create RVG on the data volume.

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.

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

Connecting to start the database outside VCS control, using sqlplus takes more than 10 minutes to respond after pulling the public network cable. [704069]

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*.

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

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

Workaround: No workaround.

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.

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 -cmdtofcf command generates a broken main.cf file [1919951]

The `hacf -cmdtofcf` 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 -cmdtofcf` command.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Workaround: Terminate the agent process.

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

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

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

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

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

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

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.

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

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

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

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

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

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

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

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

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

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).

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

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

Workaround: No workaround.

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.
- WPAR users and 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 WPAR and non-root users are validated.

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.

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

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

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

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

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

Workaround: Set the LinkTestRatio attribute to 0.

Bringing the LPAR resource offline may fail [2418615]

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

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

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

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

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

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

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

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

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

Workaround:

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

CoordPoint agent remains in faulted state [2852872]

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

Workaround: After HAD has stopped, reconfigure fencing.

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

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

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

Workaround: No workaround.

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

If a WPAR is already present on the system at the time of VCS installation, and this WPAR or an application running inside this WPAR needs to be monitored using

VCS, then VCS does not monitor the application running in that WPAR. This is because the VCS packages/files are not visible inside that WPAR.

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

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

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

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

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

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

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

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

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

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

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.

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

WPAR-aware agents cannot run in a non-shared WPAR [3313698]

Non-shared WPAR has writable `/usr` file system and `/opt` file system local to WPAR. The common product installer installs VCS packages in `/opt/VRTSvcs` and libraries in `/usr/lib` in a global environment. As VCS packages cannot be synchronized with a local copy of `/usr` and `/opt` on a non-shared WPAR, they are not available to the non-shared WPAR. Therefore in absence of the VCS packages, agents which are configured to monitor applications inside non-shared WPAR cannot run.

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

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

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

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

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

Workaround: Online the cache area using `sfcache` command

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

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

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

Workaround: Restart the RemoteGroup agent.

Issues related to the VCS database agents

This section describes the known issues about VCS database agents.

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. You can run the following commands to check if the Management DB is running on a node:

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

Run the following commands to migrate the Management DB to another node:

```
# /oracle/12102/app/gridhome/bin/srvctl relocate mgmtdb -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.

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

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

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

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 (`vpdb`) hangs. Due to this, the monitor entry point in PDB gets timed out and there is no issue found in oracle-12.1.0.1.0 .

Workaround: No workaround.

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

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

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

Workaround: Disable the threaded execution feature as it is not supported on Oracle 12C.

Issues related to the agent framework

This section describes the known issues about the agent framework.

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

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

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

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

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

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

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

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

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

Workaround: No workaround.

Delayed response to VCS commands observed on nodes with several resources and system has high CPU usage or high swap usage [3208239]

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

Some of the commands are mentioned below:

- `# hares -online`
- `# hares -offline`
- `# hagrp -online`
- `# hagrp -offline`

■ # hares -switch

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

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

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

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

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

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

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

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

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

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

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

VCS engine must receive periodic heartbeat from CFSMount agent to ensure that it is running properly on the system. The heartbeat is decided by AgentReplyTimeout attribute. Due to high CPU usage or memory workload (for example, swap usage

greater than 85%), agent may not get enough CPU cycles to schedule. This causes heartbeat loss with VCS engine and as a result VCS engine terminates the agent and starts the new agent. This can be identified with the following error message in the engine log:

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

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

```
# hatype -modify CFSSMount NumThreads 1
```

Even after the above command if CFSSMount 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 7.3.1 release.

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

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

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

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

Issues related to Intelligent Monitoring Framework (IMF)

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.

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.

Process offline monitoring registrations through the AMF program freezes Cluster Server (VCS) nodes on some service packs of AIX 7.1 and 6.1 versions [3540463]

Process offline monitoring registrations with Asynchronous Monitoring Framework (AMF) freezes Cluster Server nodes on AIX 7.1 TL3 SP3 or SP4 and AIX 6.1 TL9 SP3 or SP4 versions.

Workaround: For resolve this issue, install the IBM APAR IV63274 on all nodes. Refer to the following TechNote for details: <http://www.veritas.com/docs/000022623>.

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).

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

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.

LLT known issues

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

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

With each received packet, LLT increments the following variables:

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

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

This does not impact the LLT functionality.

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 steps to add a static route using the `smitty` utility:

- 1 Run the command: `# smitty tcpip`
- 2 Select **IPV6 Configuration** from the menu and press Enter.
- 3 Select `IPV6 Static Routes` and then `Add an IPV6 Static Route` from the menu.
- 4 Specify the value for *** IPV6 DESTINATION Address**, *** IPV6 GATEWAY Address**, and `Network Interface` fields to set source and destination address.

I/O fencing known issues

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

CP server repetitively logs unavailable IP addresses (2530864)

If coordination point server (CP server) fails to listen on any of the IP addresses that are mentioned in the `vxcps.conf` file or that are dynamically added using the command line, then CP server logs an error at regular intervals to indicate the failure. The logging continues until the IP address is bound to successfully.

```
CPS ERROR V-97-51-103 Could not create socket for host
10.209.79.60 on port 14250
CPS ERROR V-97-1400-791 Coordination point server could not
open listening port = [10.209.79.60]:14250
Check if port is already in use.
```

Workaround: Remove the offending IP address from the listening IP addresses list using the `rm_port` action of the `cpsadm` command.

See the *prod_ug* for more details.

Fencing port b is visible for few seconds even if cluster nodes have not registered with CP server (2415619)

Even if the cluster nodes have no registration on the CP server and if you provide coordination point server (CP server) information in the `vxfenmode` file of the cluster nodes, and then start fencing, the fencing port b is visible for a few seconds and then disappears.

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

The cpsadm command fails if LLT is not configured on the application cluster (2583685)

The `cpsadm` command fails to communicate with the coordination point server (CP server) if LLT is not configured on the application cluster node where you run the `cpsadm` command. You may see errors similar to the following:

```
# cpsadm -s 10.209.125.200 -a ping_cps
CPS ERROR V-97-1400-729 Please ensure a valid nodeid using
environment variable
CPS_NODEID
CPS ERROR V-97-1400-777 Client unable to communicate with CPS.
```

However, if you run the `cpsadm` command on the CP server, this issue does not arise even if LLT is not configured on the node that hosts CP server. The `cpsadm`

command on the CP server node always assumes the LLT node ID as 0 if LLT is not configured.

According to the protocol between the CP server and the application cluster, when you run the `cpsadm` on an application cluster node, `cpsadm` needs to send the LLT node ID of the local node to the CP server. But if LLT is unconfigured temporarily, or if the node is a single-node VCS configuration where LLT is not configured, then the `cpsadm` command cannot retrieve the LLT node ID. In such situations, the `cpsadm` command fails.

Workaround: Set the value of the `CPS_NODEID` environment variable to 255. The `cpsadm` command reads the `CPS_NODEID` variable and proceeds if the command is unable to get LLT node ID from LLT.

In absence of cluster details in CP server, VxFEN fails with pre-existing split-brain message (2433060)

When you start server-based I/O fencing, the node may not join the cluster and prints error messages in logs similar to the following:

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1043
Detected a preexisting split brain. Unable to join cluster.
```

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
operation failed.
CPS ERROR V-97-1400-446 Un-authorized user cpsclient@sys1,
domaintype vx; not allowing action
```

The `vxfend` daemon on the application cluster queries the coordination point server (CP server) to check if the cluster members as seen in the GAB membership are registered with the CP server. If the application cluster fails to contact the CP server due to some reason, then fencing cannot determine the registrations on the CP server and conservatively assumes a pre-existing split-brain.

Workaround: Before you attempt to start VxFEN on the application cluster, ensure that the cluster details such as cluster name, UUID, nodes, and privileges are added to the CP server.

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

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

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

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

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

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

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

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

Workaround: Start VxFEN again after some time.

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

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

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

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

When you configure fencing in customized mode and do not provide default port, fencing comes up. However, the `vxfenconfig -l` command output does not list the port numbers.

Workaround: Retain the "port_https=<port_value>" setting in the `/etc/vxfenmode` file, when using customized fencing with at least one CP server. The default port value is 443.

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

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

Workaround:

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

The vxfenswap utility deletes comment lines from the /etc/vxfemode file, if you run the utility with hacli option (3318449)

The vxfenswap utility uses RSH, SSH, or hacli protocol to communicate with peer nodes in the cluster. When you use vxfenswap to replace coordination disk(s) in disk-based fencing, vxfenswap copies `/etc/vxfenmode` (local node) to `/etc/vxfenmode` (remote node).

With the hacli option, the utility removes the comment lines from the remote `/etc/vxfenmode` file, but, it retains comments in the local `/etc/vxfenmode` file.

Workaround: Copy the comments manually from local `/etc/vxfenmode` to remote nodes.

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 fileset, 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, `vxfsend`, invokes some of the fencing scripts on the node. Each of these scripts has a timeout value of 120 seconds. If any of these scripts fails, fencing configuration fails on that node.

Some of these scripts use `cpsadm` commands to communicate with CP servers. When the node comes up, `cpsadm` commands try to connect to the CP server using VIPs for a timeout value of 60 seconds. So, if the multiple `cpsadm` commands that are run within a single script exceed the timeout value, then the total timeout value exceeds 120 seconds, which causes one of the scripts to time out. Hence, I/O fencing does not come up on the client node.

Note that this issue does not occur with IPM-based communication between CP server and client clusters.

Workaround: Fix the CP server.

The `vxfsenconfig -l` command output does not list Coordinator disks that are removed using the `vxdsmpadm exclude dsmpnodename=<dsmp_disk/node>` command [3644431]

After you remove a Coordinator disk used by fencing or fencing disk group by running the `vxdsmpadm exclude dsmpnodename=<dsmp_disk/node>` command, the removed disk is not listed in the `vxfsenconfig -l` command output.

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

Workaround: Run the `vxdsmpadm include dsmpnodename=<dsmp_disk/node>` command to again enable the dsmp disk. This disk will show up in subsequent `vxfsenconfig -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 `vxfsenwap` 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 `vxfsenwap` 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*.

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

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

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.

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.

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

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

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

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

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

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

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

Workaround:

To resolve this known issue

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

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

An error message is received when you perform off-host clone for RAC and the off-host node is not part of the CVM cluster (1834860)

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

```
Cannot open file /etc/vx/vxdba/rac11g1/.DB_NAME
(No such file or directory).
SFORA vxreptadm ERROR V-81-8847 Cannot get filename from sid
for 'rac11g1', rc=-1.
SFORA vxreptadm ERROR V-81-6550 Could not connect to repository
database.
VxVM vxdbg ERROR V-5-1-582 Disk group SNAP_rac11dgl: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac11dgl.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac11dgl failed.
```

Workaround: Currently there is no workaround 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.

Upgrading operating system Technology Levels along with Storage Foundation using an alternate disk fails (2162945)

Upgrading the operating system Technology Levels (TL) along with Storage Foundation using an alternate disk fails occasionally with the following error:

```
alt_disk_copy: 0505-224 ATTENTION:
An error occurred during installation of
one or more software components.
Modifying ODM on cloned disk.
Building boot image on cloned disk.
forced unmount of /alt_inst/var/adm/ras/platform
forced unmount of /alt_inst/var
umount: error unmounting /dev/alt_hd2: Device busy
0505-144 alt_disk_install: Unable to unmount alt_inst filesystems.
```

No issues have been observed with Storage Foundation in the cause of the failure.

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

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 CVMVolDg 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 vxldg 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:

```
# vxctl 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.

Write back cache is not supported on the cluster in FSS scenario [3723701]

Write back cache is not supported in FSS scenario on Cluster file system. When the Write back is enabled, for example, node N1 and N2 both have its own SSD and they are using each other's SSD as remote cache. Then it may cause data corruption and the recovery is not possible on cluster.

Workaround: This issue has been fixed.

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 CMMVolumeIoTest testvol

# haconf -dump -makero
```

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.

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

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

Workaround: There is no workaround for this issue.

The `fsppadm 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.

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

Node fails to join the cluster after installation or upgrade to specific operating system versions

After you install or upgrade to operating system versions AIX 6.1 TL8 or 7.1 TL2 with Oracle RAC 11g Release 2 and later versions, Oracle Clusterware successfully starts on only one node in the cluster. The remaining nodes fail to join the cluster as the daemons `CRSD` and `EVMD` are in an intermediate state. For more information, see the metalink document: 1528452.1

Workaround: Install the following AIX HIPER efixes:

```
AIX 6.1 TL08 SP01 - APAR IV35888
AIX 7.1 TL02 SP01 - APAR IV35893
```

Storage Foundation Oracle RAC issues

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

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.

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

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

- CVM master node panics
- Private network failure

As a result, the ASM disk groups get dismounted.

Workaround: See to the Oracle metalink document: 1581684.1

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

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

For more information, see the following Technote:

<http://www.veritas.com/docs/000010309>

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.

Process offline monitoring issues with Asynchronous Monitoring Framework [3540463]

Process offline monitoring registrations with Asynchronous Monitoring Framework (AMF) freezes Cluster Server nodes on AIX 7.1 TL3 SP3/SP4 and AIX 6.1 TL9 SP3/SP4 versions.

Workaround: Install the IBM APAR [IV63274](#) on all nodes to fix this issue.

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.

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

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

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

- On galaxy: 192.168.12.1
- On nebula: 192.168.12.2

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

CVMVolDg agent may fail to deport CVM disk group

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

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

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.

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
```

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

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

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

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

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

Workaround:

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

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.

Storage Foundation for Databases (SFDB) tools known issues

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

Sometimes SFDB may report the following error message: SFDB remote or privileged command error (2869262)

While using SFDB tools, if you attempt to run commands, such as `dbed_update` then you may observe the following error:

```
$ /opt/VRTSdbed/bin/dbed_update
No repository found for database faildb, creating new one.
SFDB vxsfadm ERROR V-81-0450 A remote or privileged command could not
be executed on swpa04
```

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

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

Workaround: There is no workaround for this issue.

SFDB commands do not work in IPV6 environment (2619958)

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

Workaround:

There is no workaround at this point of time.

The database clone operation using the `vxsfadm -o clone(1M)` command fails (3313715)

In an Oracle RAC environment, while you try to bring up a cloned database instance on a remote RAC node using the `SECONDARY_HOST` parameter in snapshot configuration, the database clone operation fails. Additionally, the following error message occurs:

```
[oracle@rac-v01 ~]$ vxsfadm -s flashsnap -a oracle -o clone
--flashsnap_name sn115 --clone_path /cloneoracle --clone_name cln709
--secondary_host rac-v02
```

SFDB vxsfadm ERROR V-81-0602 Remote execution failed:

SFDB vxsfadm ERROR V-81-0000 Another instance of vxsfadm is running

Workaround: Avoid using the `SECONDARY_HOST` parameter in a snapshot configuration. Additionally, perform the cloning operation locally on the RAC node where you need the cloned instance to be brought up.

In an off-host scenario, a clone operation may fail with an error message (3313572)

A clone operation may fail with the following error due to restricted process resource limits in effect for the root user.

```
ORA-00283: recovery session canceled due to errors
ORA-01110: data file 5: '/flash_snap/oracle/oradata/run/soe.dbf'
ORA-01157: cannot identify/lock data file 5 - see DBWR trace file
ORA-01110: data file 5: '/flash_snap/oracle/oradata/run/soe.dbf'
```

All off-host operations especially the clone operations are routed through the `vxdbd` daemon, which is currently unable to support the per-user process resource limits that are set for the non-root users. Thus, all operations that are routed through `vxdbd`, inherit the resource limits set for the root user. If these limits are restrictive, then the operation may fail.

Workaround: Set the resource limit for the root user to a maximum range such that it is close to the Oracle database's requirement.

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:

```
fspadm: 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_vmcloneb` fails if you put in multiple lines for parameter.

Workaround:Edit the PFILE to arrange the text so that the parameter values are on a single line. If the database uses a spfile and some parameter values are spread across multiple lines, then use the Oracle commands to edit the parameter values such as they fit in a single line.

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 archive log configuration on RAC (2846399)

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

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

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

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

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

Database Storage Checkpoints created by using dbed_ckptcreate may not be visible after upgrading to 7.3.1 (2626248)

After upgrading from a 5.0 release to 7.3.1, the Database Storage Checkpoints created earlier using dbed_ckptcreate may not be migrated.

Workaround: Perform the following steps to make the old Database Storage Checkpoints visible.

To resolve the issue

- 1 Remove the new repository.
 - Examine the contents of the /var/vx/vxdba/rep_locfile to determine the location of the 7.3.1 repository.

- Remove the `.sfae` directory specified as the `location` attribute.
- 2 Remove the repository location file: `/var/vx/vxdba/rep_loc`.
- 3 Create a symlink `/var/vx/vxdba/<SID>/sfdb_rept` pointing to the `.sfdb_rept` directory created in the same location as the `.sfae` directory removed earlier.

```
$ ln -s <location>/sfdb_rept /var/vx/vxdba/<SID>/sfdb_rept
```

This step creates a symlink to the old repository.

- 4 Import repository data by running the `dbed_update` command.

This step imports the data from the old repository.

The old Database Storage Checkpoints are now visible.

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

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

Workaround: There is no workaround for this issue.

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

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

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

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

Workaround: There is no workaround for this issue.

If any SFDB installation with authentication setup is upgraded to 7.3.1, the commands fail with an error (3644030)

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

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

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

```
Action: Verify that the prodhost is reachable. If it is, verify
that
```

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

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

Error message displayed when you use the vxsfadm -a oracle -s filesnap -o destroyclone command (3901533)

The vxsfadm -a oracle -s filesnap -o destroyclone command gives with the following error message:

```
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 cln1
Redundant argument in sprintf at /opt/VRTSdbed/lib/perl/DBED/Msg.pm line 170.
Shutting down clone database... Done
Destroying clone... Done
```

Workaround: This message can be ignored. It does not affect the functionality in any manner.

Software Limitations

This chapter includes the following topics:

- [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](#)
- [Operating system limitations](#)

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 5-1 describes the DMP tunable parameters and the new values.

Table 5-1 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

- Issue the following commands:

```
# vxddmpadm settune dmp_restore_interval=60

# vxddmpadm settune dmp_path_age=120
```

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

```
# vxddmpadm gettune dmp_restore_interval

# vxddmpadm gettune dmp_path_age
```

DMP support in AIX virtualization environment (2138060)

DMP does not support exporting paths to the same LUN through both vSCSI and NPIV interfaces.

DMP treats the same LUN seen through vSCSI and NPIV interfaces as two separate LUNs, because the behavior of the LUN at the VIOC level is different due to the intermediate SCSI interface at the VIOS level for vSCSI devices.

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.

This issue is only applicable to non-root devices.

Veritas Volume Manager software limitations

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

MPIO device names shown in error state (3169587)

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

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

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

Workaround:

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

To resolve the issue with MPIO devices in error state

- 1 Remove the following file:

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

- 2 Reset the `vxconfigd` daemon:

```
# vxconfigd -kr reset
```

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.

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

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

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

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

Veritas File System software limitations

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

Recommended limit of number of files in a directory

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

The shell cannot handle 64-bit inode numbers inside the .checkpoint directory when uniqueino is enabled

Due to a limitation with the AIX operating system, the shell cannot handle the 64-bit inode numbers inside the `.checkpoint` directory when the `uniqueino` mount option is enabled. Some shell functions such as auto-complete and globs, for example `rm *`, do not function properly in the `.checkpoint` directory. This also affects 32-bit applications that try to read the contents of the `.checkpoint` directory or any of its subdirectories. This does not affect any 64-bit applications.

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

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

Limitations with delayed allocation for extending writes feature

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

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

FlashBackup feature of NetBackup 7.5 (or earlier) does not support disk layout Version 8, 9, or 10

The FlashBackup feature of NetBackup 7.5 (or earlier) does not support disk layout Version 8, 9, or 10.

SmartIO software limitations

The following are the SmartIO software limitations in this release.

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

VVR Replication in a shared environment

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

VVR IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

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

VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 6.1 and the prior major releases of Storage Foundation (6.0 and 6.0.1). Replication between versions is supported for disk group versions 170, 180, and 190 only. Both the Primary and Secondary hosts must be using a supported disk group version.

Cluster Server software limitations

These software limitations apply to the following products:

- Veritas InfoScale Availability
- Veritas InfoScale Enterprise

Limitations related to bundled agents

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

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

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

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

```
/etc/netsvd.conf
```

Volume agent clean may forcibly stop volume resources

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

False concurrency violation when using PidFiles to monitor application resources

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

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

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

Volumes in a disk group are started automatically when the disk group is imported, irrespective of the value of the `StartVolumes` attribute in VCS. This behavior is

observed if the value of the system-level attribute `autostartvolumes` in Veritas Volume Manager is set to On.

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

WPAR agent registered to IMF for Directory Online event

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

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.

Live Partition Mobility (LPM) of management LPAR is not supported

Live Partition Mobility (LPM) of management LPAR is not supported.

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

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

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

Limitations related to VCS engine

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

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

Workaround: No workaround.

Preferred fencing ignores the forecasted available capacity [3077242]

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

Workaround: No workaround.

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

Veritas cluster configuration wizard limitations

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 IMF

- If a process is registered with IMF for offline monitoring, IMF may not detect the process being executed if the length of the process and related arguments exceed 70 characters. In case of ProPCV, IMF may not be able to prevent the process from coming online if the length of the process and related arguments exceeds 70 characters. This limitation affects Application agent and Process agent. Refer to the *Cluster Server Bundled Agents Reference Guide* for more information. [2768558]

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.

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.

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.

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.

Virtualizing shared storage using VIO servers and client partitions

In an Advanced POWER™ Virtualization (APV) environment, AIX uses the VIO Server to monitor and manage the I/O paths for the virtualized client partitions. At a very high level, the VIO server provides a partition's access to storage that is external to the physical computer. The VIO server encapsulates the physical hardware into virtual adapters called virtual SCSI adapters (server adapter). On the client side, you can create virtual adapters (client adapters) that map to the server adapter and enable a partition to connect to external storage.

The VIO server provides similar mechanisms to share limited networking resources across partitions. Refer to the manual that came with your system to help set up partitions, and to configure and use the various components such as VIO server and HMC, which are integral parts of IBM's APV environment.

The minimum patch level for using VIO servers with VCS is: version 2.1.3.10-FP-23 and later.

Supported storage

Refer to the IBM data sheet:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/home.html>

Disk Restrictions

When using VCS in combination with VIO servers and their client partitions, you need to ensure that no reservations are placed on the shared storage. This enables client partitions on different systems to access and use the same shared storage.

- If the shared storage is under MPIO control, set the `reserve_policy` attribute of the disk to `no_reserve`.
- If the shared storage is not under MPIO control, look up the array documentation to locate a similar attribute to set on the disk.

Internal testing on EMC disks shows that this field maps as the `reserve_lock` attribute for EMC disks. In this case, setting it to `no` achieves the same result.

Accessing the same LUNs from Client Partitions on different Central Electronics Complex (CEC) modules

This section briefly outlines how to set shared storage so that it is visible from client partitions on different CEC modules.

With the VIO server and client partitions set up and ready, make sure that you have installed the right level of operating system on the client partitions, and that you have mapped the physical adapters to the client partitions to provide access to the external shared storage.

To create a shareable diskgroup, you need to ensure that the different partitions use the same set of disks. A good way to make sure that the disks (that are seen from multiple partitions) are the same is to use the disks serial numbers, which are unique.

Run the following commands on the VIO server (in non-root mode), unless otherwise noted.

Get the serial number of the disk of interest:

```
$ lsdev -dev hdisk20 -vpd
    hdisk20
    U787A.001.DNZ06TT-P1-C6-T1-W500507630308037C-
    L401 0401A00000000 IBM FC 2107

Manufacturer.....IBM
Machine Type and Model.....2107900
Serial Number.....7548111101A
EC Level.....131
Device Specific.(Z0).....10
Device Specific.(Z1).....0100
...
```

Make sure the other VIO server returns the same serial number. This ensures that you are viewing the same actual physical disk.

List the virtual SCSI adapters.

```
$ lsdev -virtual | grep vhost
vhost0 Available Virtual SCSI Server Adapter
vhost1 Available Virtual SCSI Server Adapter
```

Note: Usually vhost0 is the adapter for the internal disks. vhost1 in the example above maps the SCSI adapter to the external shared storage.

Prior to mapping hdisk20 (in the example) to a SCSI adapter, change the reservation policy on the disk.

```
$ chdev -dev hdisk20 -attr reserve_policy=no_reserve
hdisk20 changed
```

For hdisk20 (in the example) to be available to client partitions, map it to a suitable virtual SCSI adapter.

If you now print the reserve policy on hdisk20 the output resembles:

```
$ lsdev -dev hdisk20 attr reserve_policy
value
no_reserve
```

Next create a virtual device to map hdisk20 to vhost1.

```
$ mkvdev -vdev hdisk20 -vadapter vhost1 -dev mp1_hdisk5
mp1_hdisk5 Available
```

Finally on the client partition run the cfgmgr command to make this disk visible via the client SCSI adapter.

You can use this disk (hdisk20 physical, and known as mp1_hdisk5 on the client partitions) to create a diskgroup, a shared volume, and eventually a shared file system.

Perform regular VCS operations on the clients vis-a-vis service groups, resources, resource attributes, etc.

Cluster Manager (Java console) limitations

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

Cluster Manager does not work if the hosts file contains IPv6 entries

VCS Cluster Manager fails to connect to the VCS engine if the /etc/hosts file contains IPv6 entries.

Workaround: Remove IPv6 entries from the /etc/hosts file.

VCS Simulator does not support I/O fencing

When running the Simulator, be sure the UseFence attribute is set to the default, "None".

The operating system does not distinguish between IPv4 and IPv6 packet counts

In a dual-stack configuration, when you use packet counts and the IPv6 network is disabled, the NIC agent might not detect a faulted NIC. It might not detect a fault because while the IPv6 network is down its packet count still increases. The packet count increases because the operating system does not distinguish between the packet counts for IPv4 and IPv6 networks. The agent then concludes that the NIC is up. If you are using the same NIC device for IPv4 as well as IPv6 resources, set PingOptimize to 0 and specify a value for the NetworkHosts attribute for either the IPv6 or the IPv4 NIC resource. [1061253]

A service group that runs inside of a WPAR may not fail over when its network connection is lost

For a WPAR configuration when the WPAR root is on NFS, the WPAR service group may not fail over if the NFS connection is lost. This issue is due to an AIX operating system limitation. [1637430]

Limitations related to LLT

This section covers LLT-related software limitations.

LLT over IPv6 UDP cannot detect other nodes while Veritas InfoScale tries to form a cluster (1907223)

LLT over IPv6 requires link-local scope multicast to discover other nodes when Veritas InfoScale tries to form a cluster. If multicast networking is undesirable, or unavailable in your environment, use the address of the peer nodes to eliminate the need for the multicast traffic.

Workaround: Add the set-addr entry for each local link into the /etc/lldtab file. You add the entry to specify the address of the peer nodes that are available on the corresponding peer links. For example, you add the following lines into the lldtab file to specify the set-addr entry for a node. In this example, the node's IPv6 address is fe80::21a:64ff:fe92:1d70.

```
set-addr 1 link1 fe80::21a:64ff:fe92:1d70
set-arp 0
```

LLT does not start automatically after system reboot (2058752)

After you reboot the systems, if you had not completed the terminal setting procedure, LLT does not start automatically and does not log any error messages. You can manually start LLT using the `/etc/init.d/llt.rc` command.

If you reinstall a system, when the system reboots a message appears on the system console to set the terminal setting if you have not already done so. LLT does not start until you complete the terminal setting procedure.

Workaround: To resolve the LLT startup issue

- 1 After you reboot the systems, open the system console using any available method, for example, from HMC.
- 2 On the console, go to the terminal setting menu, and set the terminal of your choice.
- 3 Select the **Task Completed** menu option.

Limitation of LLT support over UDP 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, LLT may not work properly.

Workaround: Do not use alias IP addresses for LLT over UDP.

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.

Limitation with RDAC driver and FAStT array for coordinator disks that use raw disks

For multi-pathing to connected storage, AIX uses the RDAC driver for FAStT arrays. Since it is an active/passive array, only the current active path is exposed to clients. The I/O fencing driver, vxfen, can use only a single active path and has no foreknowledge of the passive paths to the coordinator disks on an array. If the single active path fails, all nodes in the cluster lose access to the coordinator disks.

The loss of the path to the coordinator disks can potentially go unnoticed until a reboot, split brain, or any other reason that leads to a cluster membership change occurs. In any of these conditions, the cluster cannot form, and all nodes panic to prevent data corruption. No data loss occurs.

Workaround: Use DMP and specify paths to coordinator disks as DMP paths rather than raw disks to avoid this limitation.

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 fileset, the DMP module will no longer be

loaded in memory. On a system where VRTSvxvm fileset 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.

```
# hstop -local -force
```

or

```
# hstop -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 7.3.1 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.

Upgrade of secure clusters not supported using native operating system tools

This release does not support the upgrade of secure clusters using native operating system tools such as Alternate Disk Installation (ADI) and Network Install Manager Alternate Disk Migration (NIMADM).

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:

- For Oracle RAC 11g Release 2 and later versions: 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.

Operating system limitations

Upgrading from AIX 7.1 TL4 SP2 to SP3 can cause agents to fault (3916650)

Upgrading from AIX 7.1 TL4 SP2 to SP3 can cause agents to fault and may result in an application or system outage. This affects the physical NICs only.

To fix this issue, you must download and apply the following epackage from IBM Support Portal:

```
IV92499: CORE DUMPS IN MULTITHREADED PROCESS USING OLSON TZ
```

Index

K

Known issues
SFCFS 91

V

VxFS Limitations
software 115