

Veritas InfoScale™ Availability 7.0.1 Release Notes - Solaris 10 x64

Platform Release

Veritas InfoScale Availability Release Notes

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

Product version: 7.0.1

Document version: 7.0.1 Rev 1

Legal Notice

Copyright © 2015 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo, and NetBackup 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.

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 Corporation and its licensors, if any.

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

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within the company to answer your questions in a timely fashion.

Our support offerings include the following:

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

For information about our support offerings, you can visit our website at the following URL:

www.veritas.com/support

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.veritas.com/support

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information

- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Technical Support
 - Recent software configuration changes and network changes

Licensing and registration

If your product requires registration or a license key, access our technical support Web page at the following URL:

www.veritas.com/support

Customer service

Customer service information is available at the following URL:

www.veritas.com/support

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Advice about technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

If you want to contact us regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Contents

Technical Support	3	
Chapter 1	About this document 16	
	About this document	16
Chapter 2	About Veritas InfoScale Availability 17	
	About Veritas InfoScale Availability	17
Chapter 3	About Veritas Services and Operations Readiness Tools 18	
	About Veritas Services and Operations Readiness Tools (SORT)	18
Chapter 4	Changes introduced in 7.0.1 23	
	Licensing changes for InfoScale 7.0.1	23
	InfoScale Product keys	23
	The vxlicinstupgrade utility	23
	Changes related to installation and upgrade	24
	Mounting the ISO image	24
	Changes to InfoScale Availability in this release	24
	IMF Support for Oracle VM Server for SPARC Agent	25
	Stronger security with 2048 bit key and SHA256 signature certificates	25
	VMwareDisks agent	25
	Co-existence of SF 6.0.5 and Availability 7.0.1	25
	IMF Support for Oracle VM Server for SPARC Agent	25
	Support for SmartIO caching on SSD devices exported by FSS	25
	Not supported in this release	26
	ApplicationHA is not included in the 7.0.1 Veritas InfoScale product family	26
	Inter Process Messaging (IPM) protocol used for secure communication is not supported	26
	Changes related to documents	26

Chapter 5	System requirements	28
	Supported Solaris operating system	28
	Supported database software	28
	Hardware compatibility list	29
	Number of nodes supported	29
Chapter 6	Fixed Issues	30
	Veritas Cluster Server fixed issues	30
Chapter 7	Known issues	31
	Operational issues for VCS	31
	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]	31
	Some VCS components do not work on the systems where a firewall is configured to block TCP traffic	32
	Stale legacy_run services seen when VCS is upgraded to support SMF [2431741]	32
	Missing characters in system messages [2334245]	33
	After OS upgrade from Solaris 10 update 8 or 9 to Solaris 10 update 10 or 11, Samba server, SambaShare and NetBios agents fail to come online [3321120]	33
	CP server does not allow adding and removing HTTPS virtual IP or ports when it is running [3322154]	34
	CP server does not support IPv6 communication with HTTPS protocol [3209475]	34
	System encounters multiple VCS resource timeouts and agent core dumps [3424429]	34
	Some VCS components do not work on the systems where a firewall is configured to block TCP traffic [3545338]	34
	Issues related to the VCS engine	35
	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 [3856231]	35
	Clusters with VCS 6.0.5 or earlier versions cannot form cross cluster communication (like GCO, STEWARD) with clusters installed with SHA256 signature certificates [3856231]	35
	Extremely high CPU utilization may cause HAD to fail to heartbeat to GAB [1744854]	36
	Missing host names in engine_A.log file (1919953)	36

The <code>ha.cf -cmdtoctf</code> command generates a broken <code>main.cf</code> file [1919951]	36
Character corruption observed when executing the <code>uuidconfig.pl</code> <code>-clus -display -use_llthost</code> command [2350517]	36
Trigger does not get executed when there is more than one leading or trailing slash in the <code>triggerpath</code> [2368061]	36
Service group is not auto started on the node having incorrect value of <code>EngineRestarted</code> [2653688]	37
Group is not brought online if top level resource is disabled [2486476]	37
NFS resource goes offline unexpectedly and reports errors when restarted [2490331]	37
Parent group does not come online on a node where child group is online [2489053]	37
Cannot modify temp attribute when VCS is in LEAVING state [2407850]	38
Oracle service group faults on secondary site during failover in a disaster recovery scenario [2653704]	38
Service group may fail to come online after a flush and a force flush operation [2616779]	38
Elevated <code>TargetCount</code> prevents the online of a service group with <code>hagrp -online -sys</code> command [2871892]	38
Auto failover does not happen in case of two successive primary and secondary cluster failures [2858187]	39
GCO clusters remain in INIT state [2848006]	39
The <code>ha</code> commands may fail for non-root user if cluster is secure [2847998]	39
Startup trust failure messages in system logs [2721512]	40
Running <code>-delete -keys</code> for any scalar attribute causes core dump [3065357]	40
VCS enters into <code>admin_wait</code> state when Cluster Statistics is enabled with load and capacity defined [3199210]	40
Agent reports incorrect state if VCS is not set to start automatically and <code>utmp</code> file is empty before VCS is started [3326504]	40
Site preference fencing policy value fails to set on restart of a site-aware cluster [3380586]	41
VCS crashes if feature tracking file is corrupt [3603291]	41
RemoteGroup agent on versions lower than 6.2 reports service group status as UNKNOWN [3638347]	41
Issues related to the bundled agents	41
The options of the Share resource cannot be updated if the Options attribute is changed when the state of the resource is online [3854084]	42

Entry points that run inside a zone are not cancelled cleanly [1179694]	42
Solaris mount agent fails to mount Linux NFS exported directory	42
The zpool command runs into a loop if all storage paths from a node are disabled	43
Zone remains stuck in down state if tried to halt with file system mounted from global zone [2326105]	43
Process and ProcessOnOnly agent rejects attribute values with white spaces [2303513]	44
The zpool commands hang and remain in memory till reboot if storage connectivity is lost [2368017]	44
Offline of zone resource may fail if <code>zoneadm</code> is invoked simultaneously [2353541]	44
Password changed while using <code>hazonesetup</code> script does not apply to all zones [2332349]	45
RemoteGroup agent does not failover in case of network cable pull [2588807]	45
CoordPoint agent remains in faulted state [2852872]	45
Prevention of Concurrency Violation (PCV) is not supported for applications running in a container [2536037]	45
Share resource goes offline unexpectedly causing service group failover [1939398]	46
Mount agent does not support all scenarios of loopback mounts	46
Invalid Netmask value may display code errors [2583313]	46
Error message is seen for Apache resource when zone is in transient state [2703707]	46
Monitor falsely reports NIC resource as offline when zone is shutting down (2683680)	47
Apache resource does not come online if the directory containing Apache pid file gets deleted when a node or zone restarts (2680661)	47
Application resource running in a container with <code>PidFiles</code> attribute reports offline on upgrade to VCS 6.0 or later [2850927]	47
NFS client reports I/O error because of network split brain [3257399]	49
Mount resource does not support spaces in the <code>MountPoint</code> and <code>BlockDevice</code> attribute values [3335304]	49
IP Agent fails to detect the online state for the resource in an exclusive-IP zone [3592683]	49
SFCache Agent fails to enable caching if cache area is offline [3644424]	50

RemoteGroup agent may stop working on upgrading the remote cluster in secure mode [3648886]	50
(Solaris 10 x64) Disk may not be visible on VM even after the VMwareDisks resource is online [3838644]	50
(Solaris 10 x64) Virtual machine may hang when the VMwareDisks resource is trying to come online [3849480]	50
(Solaris 10 x64) VCS process may crash after upgrade [3860302]	50
Issues related to the VCS database agents	51
Netsnr agent monitoring can't detect tnslsnr running on Solaris if the entire process name exceeds 79 characters [3784547]	51
The ASMInstAgent does not support having pfile/spfile for the ASM Instance on the ASM diskgroups	51
VCS agent for ASM: Health check monitoring is not supported for ASMInst agent	51
NOFAILOVER action specified for certain Oracle errors	52
ASMInstance resource monitoring offline resource configured with OHASD as application resource logs error messages in VCS logs [2846945]	52
Oracle agent fails to offline pluggable database (PDB) resource with PDB in backup mode [3592142]	52
Clean succeeds for PDB even as PDB staus is UNABLE to OFFLINE [3609351]	53
Second level monitoring fails if user and table names are identical [3594962]	53
Monitor entry point times out for Oracle PDB resources when CDB is moved to suspended state in Oracle 12.1.0.2 [3643582]	53
Oracle agent fails to online and monitor Oracle instance if threaded_execution parameter is set to true [3644425]	53
Issues related to the agent framework	54
Agent may fail to heartbeat under heavy load [2073018]	54
Agent framework cannot handle leading and trailing spaces for the dependent attribute (2027896)	54
The agent framework does not detect if service threads hang inside an entry point [1442255]	54
IMF related error messages while bringing a resource online and offline [2553917]	55
Delayed response to VCS commands observed on nodes with several resources and system has high CPU usage or high swap usage [3208239]	55

CFSSMount agent may fail to heartbeat with VCS engine and logs an error message in the engine log on systems with high memory load [3060779]	56
Logs from the script executed other than the agent entry point goes into the engine logs [3547329]	57
Issues related to Intelligent Monitoring Framework (IMF)	57
Registration error while creating a Firedrill setup [2564350]	57
IMF does not fault zones if zones are in ready or down state [2290883]	58
IMF does not detect the zone state when the zone goes into a maintenance state [2535733]	58
IMF does not provide notification for a registered disk group if it is imported using a different name (2730774)	58
Direct execution of <code>linkamf</code> displays syntax error [2858163]	58
Error messages displayed during reboot cycles [2847950]	58
Error message displayed when ProPCV prevents a process from coming ONLINE to prevent concurrency violation does not have I18N support [2848011]	59
AMF displays StartProgram name multiple times on the console without a VCS error code or logs [2872064]	59
VCS engine shows error for cancellation of reaper when Apache agent is disabled [3043533]	59
Terminating the <code>imfd</code> daemon orphans the <code>vxnotify</code> process [2728787]	60
Agent cannot become IMF-aware with agent directory and agent file configured [2858160]	60
ProPCV fails to prevent a script from running if it is run with relative path [3617014]	60
Issues related to global clusters	60
The engine log file receives too many log messages on the secure site in global cluster environments [1919933]	61
Application group attempts to come online on primary site before fire drill service group goes offline on the secondary site (2107386)	61
Issues related to the Cluster Manager (Java Console)	61
Global Service group cannot switch to the remote site from Java GUI [3857634]	61
Some Cluster Manager features fail to work in a firewall setup [1392406]	61
VCS Cluster Configuration wizard issues	62
IPv6 verification fails while configuring generic application using VCS Cluster Configuration wizard [3614680]	62

Browser shows 404 error and wizard fails to launch when VCS is installed with Jumpstart or upgraded with Live upgrade [3626253]	62
LLT known issues	63
LLT port stats sometimes shows recvcnt larger than recvbytes (1907228)	63
I/O fencing known issues	63
The cpsadm command fails after upgrading CP server to 6.0 or above in secure mode (2846727)	63
Delay in rebooting Solaris 10 nodes due to vxfen service timeout issues (1897449)	64
CP server repetitively logs unavailable IP addresses (2530864)	65
Fencing port b is visible for few seconds even if cluster nodes have not registered with CP server (2415619)	65
The cpsadm command fails if LLT is not configured on the application cluster (2583685)	65
When I/O fencing is not up, the svcs command shows VxFEN as online (2492874)	66
In absence of cluster details in CP server, VxFEN fails with pre-existing split-brain message (2433060)	66
The vxfenswap utility does not detect failure of coordination points validation due to an RSH limitation (2531561)	67
Fencing does not come up on one of the nodes after a reboot (2573599)	67
Common product installer cannot setup trust between a client system on release version 5.1SP1 and a server on release version 6.0 or later [3226290]	68
Hostname and username are case sensitive in CP server (2846392)	68
Server-based fencing comes up incorrectly if default port is not mentioned (2403453)	68
Secure CP server does not connect from localhost using 127.0.0.1 as the IP address (2554981)	68
Unable to customize the 30-second duration (2551621)	69
CoordPoint agent does not report the addition of new disks to a Coordinator disk group [2727672]	69
Fencing may show the RFSM state as replaying for some nodes in the cluster (2555191)	69
The vxfenswap utility deletes comment lines from the <code>/etc/vxfemode</code> file, if you run the utility with <code>hacli</code> option (3318449)	69

When you configure CP server only for HTTPS-based communication, the `engine_A.log` displays a misleading message (3321101) 70

The `vxfsentsthdw` utility may not run on systems installed with partial SFHA stack [3333914] 70

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) 70

The `vxfsenconfig -l` command output does not list Coordinator disks that are removed using the `vxdsmpadm exclude dmpnodename=<dmp_disk/node>` command [3644431] 71

CP server configuration fails while setting up secure credentials for CP server hosted on an SFHA cluster (2621029) 71

Coordination point server-based fencing may fail if it is configured on 5.1SP1RP1 using 6.0.1 coordination point servers (3226290) 72

The CoordPoint agent faults after you detach or reattach one or more coordination disks from a storage array (3317123) 72

The upper bound value of FaultTolerance attribute of CoordPoint agent should be less than the majority of the coordination points. (2846389) 73

The `vxfsenswap` utility deletes comment lines from the `/etc/vxfsemode` file, if you run the utility with `hacli` option (3318449) 73

GAB known issues 73

 While deinitializing GAB client, "gabdebug -R GabTestDriver" command logs refcount value 2 (2536373) 73

 Cluster panics during reconfiguration (2590413) 74

 Cannot run `pfiles` or `truss` files on `gablogd` (2292294) 74

 While deinitializing GAB client, "gabdebug -R GabTestDriver" command logs refcount value 2 (2536373) 74

Chapter 8

InfoScale Availability software limitations in 7.0.1 75

Limitations related to bundled agents 76

 Volumes in a disk group start automatically irrespective of the value of the `StartVolumes` attribute in VCS [2162929] 76

 Programs using networked services may stop responding if the host is disconnected 76

 Volume agent clean may forcibly stop volume resources 76

False concurrency violation when using PidFiles to monitor application resources	76
Online for Oracle VM Server for SPARC resource fails [2517350]	77
Zone agent registered to IMF for Directory Online event	77
Oracle VM Server for SPARC resource calls clean entry point when primary domain is gracefully shut down	77
Application agent limitations	77
Campus cluster fire drill does not work when DSM sites are used to mark site boundaries [3073907]	78
On Solaris 10, the online operation of IP resource may fail if <code>ifconfig -a</code> returns an error [3609861]	78
Mount agent reports resource state as OFFLINE if the configured mount point does not exist [3435266]	78
Limitations related to VCS engine	78
Loads fail to consolidate and optimize when multiple groups fault [3074299]	78
Preferred fencing ignores the forecasted available capacity [3077242]	79
Failover occurs within the SystemZone or site when BiggestAvailable policy is set [3083757]	79
Load for Priority groups is ignored in groups with BiggestAvailable and Priority in the same group [3074314]	79
Cluster configuration wizard limitations	79
Environment variable used to change log directory cannot redefine the log path of the wizard [3609791]	79
Cluster configuration wizard takes long time to configure a cluster on Solaris systems [3582495]	79
Limitations related to the VCS database agents	80
DB2 RestartLimit value [1234959]	80
Sybase agent does not perform qrmutil based checks if Quorum_dev is not set (2724848)	80
Pluggable database (PDB) online may timeout when started after container database (CDB) [3549506]	80
Systems in a cluster must have same system locale setting	80
Limitations with DiskGroupSnap agent [1919329]	80
Cluster Manager (Java console) limitations	81
Cluster Manager (Java Console) version 5.1 and lower cannot manage VCS 6.0 secure clusters	81
Cluster Manager does not work if the hosts file contains IPv6 entries	81
VCS Simulator does not support I/O fencing	81
Limited support from Cluster Manager (Java console)	81

Port change required to connect to secure cluster [2615068] 82

Limitations related to LLT 82

 Limitation of LLT support over UDP using alias IP [3622175] 82

Limitations related to I/O fencing 82

 Preferred fencing limitation when VxFEN activates RACER node
 re-election 82

 Uninstalling VRTSvxvm causes issues when VxFEN is configured
 in SCSI3 mode with dmp disk policy (2522069) 82

 Stopping systems in clusters with I/O fencing configured 83

 Node may panic if HAD process is stopped by force and then node
 is shut down or restarted [3640007] 83

Limitations related to global clusters 84

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] 84

On Solaris 10 x64, if you upgrade from 6.0.5 or earlier releases to 7.0.1
with security configured, you need to upgrade both sites of
GCO 84

Chapter 9 **Documentation** 85

Veritas InfoScale documentation 85

Documentation set 85

About this document

This chapter includes the following topics:

- [About this document](#)

About this document

This document provides important information about InfoScale Availability version 7.0.1 for Solaris. Review this entire document before you install or upgrade VCS.

This is "Document version: 7.0.1 Rev 1" of the *InfoScale Availability Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Veritas website at:

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

About Veritas InfoScale Availability

This chapter includes the following topics:

- [About Veritas InfoScale Availability](#)

About Veritas InfoScale Availability

Veritas InfoScale™ Availability helps keep organizations' information available and critical business services up and running with a robust software-defined approach. Organizations can innovate and gain cost benefits of physical and virtual across commodity server deployments. Maximum IT service continuity is ensured at all times, moving resiliency from the infrastructure layer to the application layer.

About Veritas Services and Operations Readiness Tools

This chapter includes the following topics:

- [About Veritas Services and Operations Readiness Tools \(SORT\)](#)

About Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) helps optimize the end to end experience for Veritas products.

Veritas SORT is a web-based application that automates and simplifies some of the most time-consuming administrative tasks. It helps you identify risks in your datacenters and improve operational efficiency, enabling you to manage the complexity that is associated with datacenter architectures and scale.

[Table 3-1](#) lists major tasks that SORT can help you accomplish during the whole product life cycle.

Table 3-1 SORT Service Tasks and Offerings

Task	SORT service offerings
Understand products	<ul style="list-style-type: none"> ■ Product features Display feature lists of different versions of the products, and the feature history track ■ Products and Platforms lookups Provide available products and versions for a specific platform, and supported platforms of products and versions. It also offers the quick access to other resources on SORT. ■ Future platform and feature plans Forecast the platforms, features, databases and applications to be supported and no longer supported by products, including InfoScale and the previous SFHA product family. ■ Videos List the related videos for SORT, Veritas products, and product features.

Table 3-1 SORT Service Tasks and Offerings (*continued*)

Task	SORT service offerings
Prepare for installations and upgrades	<ul style="list-style-type: none"> <li data-bbox="798 326 1220 413">■ Installation and Upgrade checklists Display system requirements including memory, disk space, and architecture. <li data-bbox="798 421 1220 560">■ SORT Data Collectors - Installation and Upgrade custom reports Create reports that determine if you're ready to install or upgrade a Veritas InfoScale products. <li data-bbox="798 569 1220 743">■ Array-specific Module Finder List the latest Array Support Libraries (ASLs) and Array Policy Modules (APMs) for UNIX servers, and Device Driver Installers (DDIs) and Device Discovery Layers (DDLs) for Windows servers. <li data-bbox="798 751 1220 873">■ Documentation Display Veritas products documents, including manual pages, product guides, and compatibility lists. <li data-bbox="798 881 1220 986">■ HCL Central Display the hardware compatibility of InfoScale products and previous SFHA products. <li data-bbox="798 994 1220 1081">■ SCL Central Display the database support matrix of InfoScale products and SFHA products. <li data-bbox="798 1090 1220 1177">■ Patch finder List downloadable patches of InfoScale products.

Table 3-1 SORT Service Tasks and Offerings (*continued*)

Task	SORT service offerings
Identify risks and get server-specific recommendations	<ul style="list-style-type: none"> <li data-bbox="798 326 1220 444">■ Risk Assessment check lists Display configuration recommendations based on your Veritas InfoScale products and platforms. <li data-bbox="798 447 1220 621">■ SORT Data Collectors - Risk Assessment custom reports Create reports that analyzes your system and recommends system availability, storage usage, performance, and best practices. <li data-bbox="798 624 1220 746">■ Patch notifications Receive automatic email notifications about patch updates. To use this function, you need to log into SORT first. <li data-bbox="798 749 1220 838">■ Error code descriptions and solutions Display descriptions and solutions for thousands of error codes.

Table 3-1 SORT Service Tasks and Offerings (*continued*)

Task	SORT service offerings
Improve efficiency	<ul style="list-style-type: none"> <li data-bbox="796 322 1213 499">■ SORT Data Collectors – Product Deployment Report Create custom reports that list your installed Veritas InfoScale products and license keys. Display licenses by products, platforms, server tiers, and systems. <li data-bbox="796 508 1213 656">■ InfoScale Entitlement Calculator Assist you to transit from previous Storage Foundation and High Availability product family to InfoScale products with the new pricing meter. <li data-bbox="796 664 1213 812">■ System Performance Value Unit (SPVU) Calculator Use the calculator to assist you with the pricing meter transition for Storage Foundation and High Availability. <li data-bbox="796 821 1213 968">■ SORT Data Collectors -Vxplorer/Windows Data Collection Collect logs of InfoScale products and SFHA products on Linux, UNIX, and Windows for troubleshooting. <li data-bbox="796 977 1213 1055">■ End of Support Life (EOSL) Provide an overview of the product release date and EOSL information. <li data-bbox="796 1064 1213 1142">■ Operations Manager Add-ons Display downloadable Operations Manager add-ons and documentation. <li data-bbox="796 1150 1213 1246">■ High Availability Agents Provide links for agents of applications, databases, replication and bundled ones. <li data-bbox="796 1255 1213 1394">■ SCOM Packs Display the downloadable Veritas Management Packs (MP) for Microsoft System Center Operations Manager (SCOM).

SORT is available at no additional charge.

To access SORT, go to: <https://sort.veritas.com>

Changes introduced in 7.0.1

This chapter includes the following topics:

- [Licensing changes for InfoScale 7.0.1](#)
- [Changes related to installation and upgrade](#)
- [Changes to InfoScale Availability in this release](#)
- [Support for SmartIO caching on SSD devices exported by FSS](#)
- [Not supported in this release](#)
- [Changes related to documents](#)

Licensing changes for InfoScale 7.0.1

The following sections describe the licensing changes in InfoScale

InfoScale Product keys

There are four new license keys that relate to the new InfoScale products. Each licenses a new range of functionality that is described in the product summaries.

The vxlicinstupgrade utility

If you are using AIX, Solaris and Linux platforms, you can license your product using the installer script. If you used to register keys manually at the command line using the vxlicinst binary, you need to use the vxlicinstupgrade binary instead in the following two circumstances:

- Upgrading from 6.0 or later
- Adding a new license when you transition a smaller product to a larger product

For more information,

Changes related to installation and upgrade

The following changes are introduced to the installation and upgrading of 7.0.1 VCS:

Mounting the ISO image

An ISO file is a disc image that must be mounted to a virtual drive for use. You must have superuser (root) privileges to mount the Veritas InfoScale ISO image.

To mount the ISO image

- 1 Log in as superuser on a system where you want to install Veritas InfoScale.
- 2 Associate the ISO image to a block device:

```
# lofiadm -a <ISO_image_path>  
          <block_device>
```

Where:

<ISO_image_path> is the complete path to the ISO image

<block_device> is the complete path to the block device

- 3 Mount the image:

```
# mount -F hsfs -o ro <block_device> /mnt
```

Changes to InfoScale Availability in this release

The following sections describe the changes in InfoScale Availability in 7.0.1:

IMF Support for Oracle VM Server for SPARC Agent

The VCS Oracle VM Server for SPARC Agent is IMF-aware and uses the AMF kernel driver for asynchronous IMF notifications. The agent also performs regular monitoring of the Oracle VM Server for SPARC.

Stronger security with 2048 bit key and SHA256 signature certificates

Stronger security with 2048 bit key and SHA256 signature certificates VCS in Veritas InfoScale Availability 7.0.1 uses 2048 bit key and SHA256 signature certificates. The vcsauthserver will generate certificates with 2048 bit key and SHA256 signature. The enhancement provides stronger security to VCS users. All the certificates will be updated by default on upgrading to VCS 7.0.1.

VMwareDisks agent

For the Solaris 10 x64 platform, Veritas InfoScale Availability has introduced a new resource type — VMwareDisks, which can monitor and control the disks attached to the VMware Virtual Machines. With the help of VMwareDisks resources, VCS can now support vMotion. These resources are managed by VMwareDisks agent.

Co-existence of SF 6.0.5 and Availability 7.0.1

This release supports the co-existence of SF 6.0.5 and Availability 7.0.1.

IMF Support for Oracle VM Server for SPARC Agent

The VCS Oracle VM Server for SPARC Agent is IMF-aware and uses the AMF kernel driver for asynchronous IMF notifications. The agent also performs regular monitoring of the Oracle VM Server for SPARC.

Support for SmartIO caching on SSD devices exported by FSS

The SmartIO feature in Veritas InfoScale Foundation, Veritas InfoScale Storage and Veritas InfoScale Enterprise products supports the use of SSD devices exported by FSS to provide caching services for applications running on VxVM volumes and VxFS file system. In this scenario, Flexible Storage Sharing (FSS) exports SSDs from nodes that have a local SSD. FSS then creates a pool of the exported SSDs in the cluster. From this shared pool, a cache area is created for those nodes in the cluster that do not have local SSDs. Each cache area is accessible only to that particular node for which it is created. The cache area can be a VxVM cache area or a VxFS cache area. The cache areas can be enabled to support warm or persistent caching across reboots.

For more information, see *Veritas InfoScale Solutions 7.0 SmartIO for Solid-State Drives Solutions Guide*.

Not supported in this release

The following features are not supported in this release but they may be supported in a future release:

- Rolling Upgrade and Phased Upgrade
- Deployment Server
- `-makeresponsefile` option for installer

Note: You can use the response file that is created by operating the installer.

The following features will not be supported by the Veritas InfoScale products:

- Web-based installation
- Oracle RAC 10g Release 2 is not supported in Storage Foundation for Oracle RAC.

ApplicationHA is not included in the 7.0.1 Veritas InfoScale product family

ApplicationHA is a standalone product and is not included in Veritas InfoScale product family. You can use any earlier version.

Inter Process Messaging (IPM) protocol used for secure communication is not supported

From 7.0.1, for CP server and InfoScale Availability clusters, HTTPS is the only supported communication protocol and IPM protocol is not supported. If you upgrade CPS server with IPM-based CP server configured, reconfigure CP server. If you want to upgrade the application clusters that use CP server based fencing to version 6.1 and later, make sure that VCS or SFHA is upgraded to version 6.1 and later first.

Note: The steps to configure CP server in *Cluster Server Administrator's Guide* and *Cluster Server Configuration and Upgrade Guide* might not work because of the disablement of the IPM based communication between CP servers and client.

Changes related to documents

The following changes have been introduced to the documents in 7.0.1:

- The look-and-feel of the documents reflect the new Veritas brand.
- This release introduces a single Release Notes and Installation Guide for the new Veritas InfoScale products.
- The Release Notes and Installation Guide documents for each component are deprecated.
- This release introduces configuration and upgrade guides for each Veritas InfoScale component.
- The software image has an updated document directory structure to reflect the product changes.
- The documents that were titled *Storage Foundation and High Availability Solutions* are renamed to *Veritas Infoscale*. The file names for these documents are changed as well.
- The ApplicationHA and Symantec High Availability Console documents are moved online.

For more information, See "[Documentation](#)" on page 85.

System requirements

This chapter includes the following topics:

- [Supported Solaris operating system](#)
- [Supported database software](#)
- [Hardware compatibility list](#)
- [Number of nodes supported](#)

Supported Solaris operating system

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

[Table 5-1](#) shows the supported operating system for this release.

Table 5-1 Supported operating system

Operating system	Levels	Chipsets	Supported Product
Solaris 10	Update 9, 10, and 11	x64	Veritas InfoScale Availability

Supported database software

For the latest information on supported database, see the following TechNote: https://www.veritas.com/support/en_US/article.DOC4039

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

Hardware compatibility list

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

https://www.veritas.com/support/en_US/article.TECH230646

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

Number of nodes supported

InfoScale Availability supports cluster configurations with up to 64 nodes.

Fixed Issues

This chapter includes the following topics:

- [Veritas Cluster Server fixed issues](#)

Veritas Cluster Server fixed issues

There is no fixed issue for InfoScale Availability 7.0.1 release on Solaris 10 x64 platform.

Known issues

This chapter includes the following topics:

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

Operational issues for VCS

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.

Note: On Solaris 11 x64, you will not hit this issue if you upgrade from VCS 7.0 to 7.0.1, because VCS 7.0 on Solaris 11 x64 has 2048 bit SHA2 certificates.

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

The following issues may occur if you install and configure VCS on systems where a firewall is installed:

- If you set up Disaster Recovery using the Global Cluster Option (GCO), the status of the remote cluster (cluster at the secondary site) shows as "initing".
- If you configure fencing to use CP server, fencing client fails to register with the CP server.
- Setting up trust relationships between servers fails.

Workaround:

- Ensure that the required ports and services are not blocked by the firewall. Refer to the *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.

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

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

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

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

For example:

```
svccfg -s smf/legacy_run
svc:/smf/legacy_run> listpg *
rc2_d_S7011t    framework    NONPERSISTENT
rc2_d_S92gab   framework    NONPERSISTENT
svc:/smf/legacy_run> delpg rc2_d_S7011t
svc:/smf/legacy_run> delpg rc2_d_S92gab
svc:/smf/legacy_run> exit
```

- Reboot the system.

Missing characters in system messages [2334245]

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

Workaround: No workaround.

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

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

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

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

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

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

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

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

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

<https://support.oracle.com/epmos/faces/BugDisplay?id=15659645>

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

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

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

The following issues may occur if you install and configure VCS on systems where a firewall is installed:

- If you set up Disaster Recovery using the Global Cluster Option (GCO), the status of the remote cluster (cluster at the secondary site) shows as "initing".

- If you configure fencing to use CP server, fencing client fails to register with the CP server.
- Setting up trust relationships between servers fails.

Workaround:

- Ensure that the required ports and services are not blocked by the firewall. Refer to the *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.

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 [3856231]

On Solaris 10 x64, the default security certificates installed with VCS 6.0.5 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.

Clusters with VCS 6.0.5 or earlier versions cannot form cross cluster communication (like GCO, STEWARD) with clusters installed with SHA256 signature certificates [3856231]

Since VCS 7.0.1, the default signature certificates installed on clusters have been upgraded to SHA256, but VCS 6.0.5 or earlier versions on Solaris 10 x64 platform does not recognize SHA256. As a result, clusters with VCS 6.0.5 or earlier versions cannot form cross cluster communication (like GCO, STEWARD) with clusters installed with SHA256 certificates.

Workaround:

Upgrade VCS to 7.0.1 or later versions.

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

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

Missing host names in engine_A.log file (1919953)

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

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

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

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

Character corruption observed when executing the `uidconfig.pl -clus -display -use_!lthost` command [2350517]

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

Workaround: No workaround.

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

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

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

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

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

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

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

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

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

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

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

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

Workaround: Terminate the agent process.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Workaround: No workaround.

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

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

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

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

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

GCO clusters remain in INIT state [2848006]

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

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

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

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

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

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

Workaround

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

Startup trust failure messages in system logs [2721512]

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

Workaround: No workaround.

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

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

Workaround: No workaround.

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

VCS enters into `admin_wait` state when started locally if:

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

Workaround:

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

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

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

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

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

Site preference fencing policy value fails to set on restart of a site-aware cluster [3380586]

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

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

VCS crashes if feature tracking file is corrupt [3603291]

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

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

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

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

Workaround: No workaround.

Issues related to the bundled agents

This section describes the known issues of the bundled agents.

The options of the Share resource cannot be updated if the Options attribute is changed when the state of the resource is online [3854084]

If the `Options` attribute of the `Share` resource is changed when the resource is in `online` state, VCS does not update the options of the `Share` resource dynamically.

Workaround:

Offline the `Share` service group. Then update the `Options` attribute and online the service group.

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

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

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

Solaris mount agent fails to mount Linux NFS exported directory

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

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

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

The workaround for this is to configure `MountOpt` attribute in mount resource and set `vers=3` for it.

Example

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

```
FSType = nfs
MountOpt = "vers=3"
)
```

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

The Solaris Zpool agent runs `zpool` commands to import and export zpools. If all paths to the storage are disabled, the `zpool` command does not respond. Instead, the `zpool` export command goes into a loop and attempts to export the zpool. This continues till the storage paths are restored and `zpool` is cleared. As a result, the offline and clean procedures of Zpool Agent fail and the service group cannot fail over to the other node.

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

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

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

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

To unmount when `VxFSMountLock` is 1

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

To forcefully unmount when `VxFSMountLock` is 1:

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

To unmount when `VxFSMountLock` is 0:

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

To forcefully unmount when `VxFSMountLock` is 0:

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

To halt the zone, use following command:

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

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

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

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

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

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

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

Workaround: There is no recommended workaround for this issue.

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

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

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

Workaround: No workaround.

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

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

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

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

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

Workaround:

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

CoordPoint agent remains in faulted state [2852872]

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

Workaround: After HAD has stopped, reconfigure fencing.

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

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

PCV does not function for an application running inside a local Zone on Solaris

Workaround: No workaround.

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

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

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

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

Mount agent does not support all scenarios of loopback mounts

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

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

Invalid Netmask value may display code errors [2583313]

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

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

Workaround: Make sure you specify a valid Netmask value.

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

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

```
Argument "VCS ERROR V-16-1-10600 Cannot connect to VCS engine\n"
isn't numeric in numeric ge (>=) at /opt/VRTSvcs/bin/Apache/Apache.pm
line 452.
VCS ERROR V-16-1-10600 Cannot connect to VCS engine
```

```
LogInt(halog call failed):TAG:E:20314 <Apache::ArgsValid> SecondLevel
MonitorTimeout must be less than MonitorTimeout.
```

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

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

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

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

Calculate the `ToleranceLimit` value as follows:

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

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

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

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

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

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

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

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

In releases prior to VCS 7.0.1, the value expected for the attribute was the pathname of the PID file including the zone root.

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

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

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

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

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

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

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


NFS client reports I/O error because of network split brain [3257399]

When network split brain occurs, the failing node may take some time to panic. As a result, the service group on the failover node may fail to come online as some of the resources (such as IP resource) are still online on the failing node. The disk group on the failing node may also get disabled but IP resource on the same node continues to be online.

Workaround: Configure the preonline trigger for the service groups containing DiskGroup resource with reservation on each system in the service group:

- 1 Copy the preonline_ipc trigger from
`/opt/VRTSvcs/bin/sample_triggers/VRTSvcs to`
`/opt/VRTSvcs/bin/triggers/preonline/ as T0preonline_ipc:`

`# cp /opt/VRTSvcs/bin/sample_triggers/VRTSvcs/preonline_ipc`
`/opt/VRTSvcs/bin/triggers/preonline/T0preonline_ipc`
- 2 Enable the preonline trigger for the service group.

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

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

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

Workaround: No workaround.

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

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

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

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

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

Workaround: Online the cache area using `sfcache` command

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

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

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

Workaround: Restart the RemoteGroup agent.

(Solaris 10 x64) Disk may not be visible on VM even after the VMwareDisks resource is online [3838644]

The disks that are attached during VMwareDisks resource online operation may not be visible to the VM user through OS commands. Due to Solaris operating system behavior, the hot plugged disks may not be visible immediately through the commands and hence user is unable to locate those disks.

Workaround: Run the command `devfsadm -Cv` on the virtual machine to rescan devices.

(Solaris 10 x64) Virtual machine may hang when the VMwareDisks resource is trying to come online [3849480]

If VMwareDisks resource attempts to attach a disk to a virtual machine that is already attached to some other virtual machine, then the attach operation may hang, causing the virtual machine to hang. As a result, the VM may miss LLT heartbeats, and may get isolated in the network.

Workaround: Ensure that the disks used by a virtual machine are not attached or used by any other virtual machine outside the VCS cluster.

(Solaris 10 x64) VCS process may crash after upgrade [3860302]

During upgrade process, installer does a forceful stop of VCS. If notifier resource is configured outside of the ClusterService group, the notifier process may remain

running during upgrade. When VCS is started post-upgrade, notifier process tries to communicate with HAD daemon. But, since the notifier process is running with an older version, this may cause HAD process to crash.

Workaround:

- Configure NotifierMngr resource under ClusterService group.
- If notifier resource is configured in a different service group, do a clean offline of the resource before starting upgrade.

Issues related to the VCS database agents

This section describes the known issues about VCS database agents.

Netlsnr agent monitoring can't detect tnslnr running on Solaris if the entire process name exceeds 79 characters [3784547]

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

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

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

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

Workaround:

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

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

The ASMInst agent does not support health check monitoring.

Workaround: Set the MonitorOption attribute to 0.

NOFAILOVER action specified for certain Oracle errors

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

See the *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.

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

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

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

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

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

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

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

Clean succeeds for PDB even as PDB 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 may fail to heartbeat under heavy load [2073018]

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

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

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

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

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

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

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

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

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

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

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

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

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

Workaround: No workaround.

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

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

Some of the commands are mentioned below:

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

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

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

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

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

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

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

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

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

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

CFSSMount 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, CFSSMount 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 CFSSMount 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 Symantec support team.

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

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

Workaround: No workaround.

Issues related to Intelligent Monitoring Framework (IMF)

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

Registration error while creating a Firedrill setup [2564350]

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

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

During Firedrill operations, VCS may log error messages related to IMF registration failure in the engine log. This happens because in the firedrill service group, there is a second 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 fault zones if zones are in ready or down state [2290883]

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

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

Workaround: Offline the zones and then restart.

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

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

Workaround: No workaround.

IMF does not provide notification for a registered disk group if it is imported using a different name (2730774)

If a disk group resource is registered with the AMF and the disk group is then imported using a different name, AMF does not recognize the renamed disk group and hence does not provide notification to DiskGroup agent. Therefore, the DiskGroup agent keeps reporting the disk group resource as offline.

Workaround: Make sure that while importing a disk group, the disk group name matches the one registered with the AMF.

Direct execution of `linkamf` displays syntax error [2858163]

Bash cannot interpret Perl when executed directly.

Workaround: Run `linkamf` as follows:

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

Error messages displayed during reboot cycles [2847950]

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

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

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

Workaround: No workaround.

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

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

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

Workaround: No Workaround.

AMF displays StartProgram name multiple times on the console without a VCS error code or logs [2872064]

When VCS AMF prevents a process from starting, it displays a message on the console and in syslog. The message contains the signature of the process that was prevented from starting. In some cases, this signature might not match the signature visible in the PS output. For example, the name of the shell script that was prevented from executing will be printed twice.

Workaround: No workaround.

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

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

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

This is an expected behavior and not an issue.

Workaround: No workaround.

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

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

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

For example:

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

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

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

Issues related to global clusters

This section describes the known issues about global clusters.

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

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

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

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

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

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

Issues related to the Cluster Manager (Java Console)

This section describes the known issues about Cluster Server Manager (Java Console).

Global Service group cannot switch to the remote site from Java GUI [3857634]

This issue occurs because local user does not have permission on the remote site.

Workaround:

Add the local user to the remote site, for example, `root@FQDN_localsystem:`

```
# hauser -add root@sys1.example.com -priv Administrator
```

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.

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

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

Workaround:

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

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

- 3 Run the following commands:

For VCS:

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

For ApplicationHA

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

LLT known issues

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

LLT port stats sometimes shows rcvcnt larger than rcvbytes (1907228)

With each received packet, LLT increments the following variables:

- rcvcnt (increment by one for every packet)
- rcvbytes (increment by size of packet for every packet)

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

This does not impact the LLT functionality.

I/O fencing known issues

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

The cpsadm command fails after upgrading CP server to 6.0 or above in secure mode (2846727)

The `cpsadm` command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTSaf package is not removed from the system, the `cpsadm` command loads the old security libraries present on the system. As the installer runs the `cpsadm` command on the CP server to add or upgrade the SFCFSHA cluster (application cluster), the installer also fails.

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

To resolve this issue

- 1 Rename `cpsadm` to `cpsadmbin`:

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

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

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

- 3 Change the permissions of the new file to 775:

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

Delay in rebooting Solaris 10 nodes due to vxfen service timeout issues (1897449)

When you reboot the nodes using the `shutdown -i6 -g0 -y` command, the following error messages may appear:

```
svc:/system/vxfen:default:Method or service exit
timed out. Killing contract 142
svc:/system/vxfen:default:Method "/lib/svc/method/vxfen stop"
failed due to signal Kill.
```

This error occurs because the vxfen client is still active when VCS attempts to stop I/O fencing. As a result, the vxfen stop service times out and delays the system reboot.

Workaround: Perform the following steps to avoid this vxfen stop service timeout error.

To avoid the vxfen stop service timeout error

- 1 Stop VCS. On any node in the cluster, run the following command:

```
# hastop -all
```

- 2 Reboot the systems:

```
# shutdown -i6 -g0 -y
```


CP server repetitively logs unavailable IP addresses (2530864)

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

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

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

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

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

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

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

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

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

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

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

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

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

When I/O fencing is not up, the `svcs` command shows VxFEN as online (2492874)

Solaris 10 SMF marks the service status based on the exit code of the start method for that service. The VxFEN start method executes the `vxfen-startup` script in the background and exits with code 0. Hence, if the `vxfen-startup` script subsequently exits with failure then this change is not propagated to SMF. This behavior causes the `svcs` command to show incorrect status for VxFEN.

Workaround: Use the `vxfenadm` command to verify that I/O fencing is running.

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

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

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

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

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

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

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

The `vxfereswap` utility runs the `vxferesconfig -o modify` command over RSH or SSH on each cluster node for validation of coordination points. If you run the `vxfereswap` 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, `vxfereswap` 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 `vxfereswap` 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 `vxfereswap` 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/vxferes/vxferes.log` file:

```
VXFEN vxferesconfig ERROR V-11-2-1007 Vxferes already configured
```

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

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

Workaround: Start VxFEN again after some time.

Common product installer cannot setup trust between a client system on release version 5.1SP1 and a server on release version 6.0 or later [3226290]

The issue exists because the VCS 5.1SP1 release version does not support separate directories for truststores. However, VCS version 6.0 and later support separate directories for truststores. Because of this mismatch in support for truststores, you cannot set up trust between client systems and servers.

Workaround: Set up trust manually between the coordination point server and client systems using the `cpsat` or `vcsat` command so that the servers and client systems can communicate in a secure mode.

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

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

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

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

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

Workaround: Retain the "port=<port_value>" setting in the `/etc/vxfenmode` file, when using customized fencing with at least one CP server. The default port value is 14250.

Secure CP server does not connect from localhost using 127.0.0.1 as the IP address (2554981)

The `cpsadm` command does not connect to the secure CP server on the localhost using 127.0.0.1 as the IP address

Workaround: Connect the secure CP server using any of the virtual IPs that is configured with the CP server and is plumbed on the local node.

Unable to customize the 30-second duration (2551621)

When the vxcperv process is not able to bind to an IP address during startup, it attempts to bind to that IP address at an interval of 30 seconds. This interval is not configurable.

Workaround: There is no workaround for this issue.

CoordPoint agent does not report the addition of new disks to a Coordinator disk group [2727672]

The LevelTwo monitoring of the CoordPoint agent does not report a fault even if the constituent of a coordinator disk group changes due to addition of new disks in the coordinator disk group

Workaround: There is no workaround for this issue.

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

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

Workaround:

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

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

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

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

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

When you configure CP server only for HTTPS-based communication, the `engine_A.log` displays a misleading message (3321101)

The `engine_A.log` file displays the following message when you configure CP server only for HTTPS-based communication but not for IPM-based communication.

```
No VIP for IPM specified in /etc/vxcps.conf
```

Workaround: Ignore the message.

The `vxfcntlsthdw` utility may not run on systems installed with partial SFHA stack [3333914]

The `vxfcntlsthdw` utility runs if the SFHA stack and VCS are fully installed with properly configured SF and VxVM. It also runs if the entire SFHA stack and VCS are not installed. However, partial installs where SF is installed and configured but VCS is not installed is not supported. The utility will display an error with the `-g` or `-c` options.

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

When a client node goes down, for reasons such as node panic, I/O fencing does not come up on that client node after node restart (3341322)

This issue happens when one of the following conditions is true:

- Any of the CP servers configured for HTTPS communication goes down.
- The CP server service group in any of the CP servers configured for HTTPS communication goes down.
- Any of the VIPs in any of the CP servers configured for HTTPS communication goes down.

When you restart the client node, fencing configuration starts on the node. The fencing daemon, `vxftend`, invokes some of the fencing scripts on the node. Each of these scripts has a timeout value of 120 seconds. If any of these scripts fails, fencing configuration fails on that node.

Some of these scripts use `cpsadm` commands to communicate with CP servers. When the node comes up, `cpsadm` commands try to connect to the CP server using VIPs for a timeout value of 60 seconds. So, if the multiple `cpsadm` commands that are run within a single script exceed the timeout value, then the total timeout value

exceeds 120 seconds, which causes one of the scripts to time out. Hence, I/O fencing does not come up on the client node.

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

Workaround: Fix the CP server.

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

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

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

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

CP server configuration fails while setting up secure credentials for CP server hosted on an SFHA cluster (2621029)

When you configure CP server using the `configure_cps.pl` utility, the configuration fails while setting up secure credentials for CP server that is hosted on an SFHA cluster. You may see the following error:

```
Creating softlink to credential directory /etc/VRTScps/db/CPSEVER
on node nodename.
Unable to connect to node nodename using /usr/bin/ssh.
Please configure ssh communication and retry. Exiting.
```

Workaround: You can use any of the following options:

- Before running the `configure_cps.pl` utility, change the default shell for root user to either KSH or bash.
- Perform the following steps after running the `configure_cps.pl` utility on each node of the cluster:
 - Manually remove the old credential directory or softlink. For example:

```
# rm -rf /var/VRTSvcs/vcsauth/data/CPSEVER
```

- Create a new soft-link to the shared location of the credential directory:

```
# ln -s path_of_CP_server_credential_directory \  
/var/VRTSvcs/vcsauth/data/CPSERVER
```

- Start the CPSSG service group:

```
# hagrps -online CPSSG -any
```

Coordination point server-based fencing may fail if it is configured on 5.1SP1RP1 using 6.0.1 coordination point servers (3226290)

The 5.1SP1 installer (CPI) cannot set up trust between a 5.1SP1 client and a 6.0 or later server, because there are no separate directories for truststores in the 5.1SP1. When trust cannot be setup, the 5.1SP1 installer cannot configure 5.1SP1 clients to work with 6.0 or later CPS in secure mode.

Workaround:

Set up trust manually between the CPS and clients using the `cpsat` or the `vcsat` command. After that, CPS and client will be able to communicate properly in the secure mode.

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

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

Workaround: Run the `vxfsnwap` 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 `vxfsnwap` 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.

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.

GAB known issues

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

While deinitializing GAB client, "`gabdebug -R GabTestDriver`" command logs `refcount` value 2 (2536373)

After you unregister the `gtx` port with `-nodeinit` option, the `gabconfig -C` command shows `refcount` as 1. But when forceful `deinit` option (`gabdebug -R GabTestDriver`) is run to deinitialize GAB client, then a message similar to the following is logged.

```
GAB INFO V-15-1-20239
Client GabTestDriver with refcount 2 forcibly deinitd on user request
```

The `refcount` value is incremented by 1 internally. However, the `refcount` value is shown as 2 which conflicts with the `gabconfig -C` command output.

Workaround: There is no workaround for this issue.

Cluster panics during reconfiguration (2590413)

While a cluster is reconfiguring, GAB broadcast protocol encounters a race condition in the sequence request path. This condition occurs in an extremely narrow window which eventually causes the GAB master to panic.

Workaround: There is no workaround for this issue.

Cannot run pfiles or truss files on gablogd (2292294)

When pfiles or truss is run on gablogd, a signal is issued to gablogd. gablogd is blocked since it has called an `gab ioctl` and is waiting for events. As a result, the pfiles command hangs.

Workaround: None.

While deinitializing GAB client, "gabdebug -R GabTestDriver" command logs refcount value 2 (2536373)

After you unregister the gtx port with `-nodeinit` option, the `gabconfig -C` command shows `refcount` as 1. But when forceful `deinit` option (`gabdebug -R GabTestDriver`) is run to deinitialize GAB client, then a message similar to the following is logged.

```
GAB INFO V-15-1-20239
Client GabTestDriver with refcount 2 forcibly deinitd on user request
```

The `refcount` value is incremented by 1 internally. However, the `refcount` value is shown as 2 which conflicts with the `gabconfig -C` command output.

Workaround: There is no workaround for this issue.

InfoScale Availability software limitations in 7.0.1

This chapter includes the following topics:

- [Limitations related to bundled agents](#)
- [Limitations related to VCS engine](#)
- [Cluster configuration wizard limitations](#)
- [Limitations related to the VCS database agents](#)
- [Systems in a cluster must have same system locale setting](#)
- [Limitations with DiskGroupSnap agent \[1919329\]](#)
- [Cluster Manager \(Java console\) limitations](#)
- [Limitations related to LLT](#)
- [Limitations related to I/O fencing](#)
- [Limitations related to global clusters](#)
- [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 \[3812313\]](#)
- [On Solaris 10 x64, if you upgrade from 6.0.5 or earlier releases to 7.0.1 with security configured, you need to upgrade both sites of GCO](#)

Limitations related to bundled agents

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.

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

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

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

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

```
/etc/nsswitch.conf
```

Volume agent clean may forcibly stop volume resources

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

False concurrency violation when using PidFiles to monitor application resources

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

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

Online for Oracle VM Server for SPARC resource fails [2517350]

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

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

Workaround: None.

Zone agent registered to IMF for Directory Online event

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

Oracle VM Server for SPARC resource calls clean entry point when primary domain is gracefully shut down

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

Workaround: No workaround.

Application agent limitations

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

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

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

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

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

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

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

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

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

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

Limitations related to VCS engine

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

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

Workaround: No workaround.

Preferred fencing ignores the forecasted available capacity [3077242]

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

Workaround: No workaround.

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

Failover always occurs within the SystemZone 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.

Cluster configuration wizard limitations

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

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

Workaround: No workaround.

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

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

Workaround: No workaround.

Limitations related to the VCS database agents

DB2 RestartLimit value [1234959]

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

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

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

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

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

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

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

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

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.
- Symantec recommends that you use the Gold configuration for the DiskGroupSnap resource.

Cluster Manager (Java console) limitations

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

Cluster Manager (Java Console) version 5.1 and lower cannot manage VCS 6.0 secure clusters

Cluster Manager (Java Console) from versions lower than VCS 5.1 cannot be used to manage VCS 6.0 secure clusters. Symantec recommends using the latest version of Cluster Manager.

See the *Cluster Server Configuration and Upgrade Guide* for instructions on upgrading Cluster Manager.

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".

Limited support from Cluster Manager (Java console)

Features introduced in VCS 6.0 may not work as expected with Java console. However, CLI option of the simulator supports all the VCS 6.0 features. You are recommended to use Veritas Operations Manager (VOM) since all new features are already supported in VOM. However, Java console may continue to work as expected with features of releases prior to VCS 6.0.

Port change required to connect to secure cluster [2615068]

In order to connect to secure cluster, the default port must be changed from 2821 to 14149. You must choose **Advanced settings** in the **Login** dialog box and change **IP: 2821** to **IP: 14149** for secure cluster login.

Limitations related to LLT

This section covers LLT-related software limitations.

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.

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

When VxFEN is configured in SCSI3 mode with dmp disk policy, the DMP nodes for the coordinator disks can be accessed during system shutdown or fencing arbitration. After uninstalling VRTSvxvm package, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm package is uninstalled, if VxFEN attempts to access DMP devices during shutdown or fencing arbitration, the system panics.

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.

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

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

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

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

or

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

- 2 Restart or shut down the node.

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

Workaround: No workaround.

Limitations related to global clusters

- Cluster address for global cluster requires resolved virtual IP.
The virtual IP address must have a DNS entry if virtual IP is used for heartbeat agents.
- Total number of clusters in a global cluster configuration can not exceed four.
- Cluster may not be declared as faulted when Symm heartbeat agent is configured even when all hosts are down.
The Symm agent is used to monitor the link between two Symmetrix arrays. When all the hosts are down in a cluster but the Symm agent is able to see the replication link between the local and remote storage, it would report the heartbeat as ALIVE. Due to this, DR site does not declare the primary site as faulted.

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

In global clusters, when you install or upgrade VCS to 7.0.1 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.

On Solaris 10 x64, if you upgrade from 6.0.5 or earlier releases to 7.0.1 with security configured, you need to upgrade both sites of GCO

On Solaris 10 x64, if you upgrade from 6.0.5 or earlier releases to 7.0.1 with security configured, you need to upgrade both sites of GCO so that communication can work across them. Upgrading only one site to 7.0.1 will break the GCO even if you do setuptrust.

Documentation

This chapter includes the following topics:

- [Veritas InfoScale documentation](#)
- [Documentation set](#)

Veritas InfoScale documentation

Veritas InfoScale documentation is available in the Adobe Portable Document Format (PDF) on the product media or with the downloaded software.

See the release notes for information on documentation changes in this release.

The documentation is available in the `/docs` directory on the product media.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The documents are updated periodically for errors or corrections. The latest documentation is available on the Services and Operations Readiness Tools (SORT) website.

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

You need to specify the product and the platform and apply other filters for finding the appropriate document.

Documentation set

The Veritas InfoScale Availability documentation includes a installation guide, release notes, a configuration guide for VCS and additional documents such as administration and agent guides.

Note: The GNOME PDF Viewer is unable to view Veritas InfoScale Availability documentation. You must use Adobe Acrobat to view the documentation.

Veritas InfoScale Availability documentation

[Table 9-1](#) lists the documentation for InfoScale Availability.

Table 9-1 Veritas InfoScale Availability documentation

Document title	File name	Description
<i>InfoScale Availability Installation Guide</i>	infoscale_availability_install_701_sol.pdf	Provides information on how to install the InfoScale Availability products.
<i>InfoScale Availability Release Notes</i>	infoscale_availability_notes_701_sol.pdf	Provides release information such as system requirements, changes, fixed incidents, known issues, and limitations of InfoScale Availability .
<i>Veritas InfoScale Getting Started Guide</i>	infoscale_getting_started_70_sol.pdf	Provides a high-level overview of installing VCS products using the script-based installer. The guide is useful for new users and returning users that want a quick refresher.
<i>Cluster Server Configuration and Upgrade Guide</i>	vcs_config_70_sol.pdf	Provides information required to configure and upgrade the component.
<i>Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents_70_sol.pdf	Provides information about bundled agents, their resources and attributes, and more related information.
<i>Cluster Server Generic Application Agent Configuration Guide</i>	vcs_gen_agent_70_unix.pdf	Provides notes for installing and configuring the generic Application agent.
<i>Cluster Server Agent for DB2 Installation and Configuration Guide</i>	vcs_db2_agent_70_sol.pdf	Provides notes for installing and configuring the DB2 agent.
<i>Cluster Server Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_agent_70_sol.pdf	Provides notes for installing and configuring the Oracle agent.
<i>Cluster Server Agent for Sybase Installation and Configuration Guide</i>	vcs_sybase_agent_70_sol.pdf	Provides notes for installing and configuring the Sybase agent.