

# Veritas Storage Foundation™ and High Availability Solutions Release Notes

AIX

5.1 Service Pack 1 Rolling Patch 1

# Storage Foundation and High Availability Solutions Release Notes 5.1 Service Pack 1 Rolling Patch 1

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: 5.1 SP1 RP1

Document version: 5.1SP1RP1.1

## Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

Symantec, the Symantec logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

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 Symantec 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. SYMANTEC 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.

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

Symantec Corporation  
350 Ellis Street  
Mountain View, CA 94043  
<http://www.symantec.com>

# Technical Support

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

Symantec's support offerings include the following:

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

For information about Symantec's support offerings, you can visit our Web site at the following URL:

[www.symantec.com/business/support/index.jsp](http://www.symantec.com/business/support/index.jsp)

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

## Contacting Technical Support

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

[www.symantec.com/business/support/contact\\_techsupp\\_static.jsp](http://www.symantec.com/business/support/contact_techsupp_static.jsp)

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

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

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

## Licensing and registration

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

[www.symantec.com/business/support/](http://www.symantec.com/business/support/)

## Customer service

Customer service information is available at the following URL:

[www.symantec.com/business/support/](http://www.symantec.com/business/support/)

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

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

## Support agreement resources

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

Asia-Pacific and Japan [customercare\\_apac@symantec.com](mailto:customercare_apac@symantec.com)

Europe, Middle-East, and Africa [semea@symantec.com](mailto:semea@symantec.com)

North America and Latin America [supportsolutions@symantec.com](mailto:supportsolutions@symantec.com)

## Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

[docs@symantec.com](mailto:docs@symantec.com)

## About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

# Contents

|                         |  |    |
|-------------------------|--|----|
| Technical Support ..... | 4  |    |
| Chapter 1               | About Veritas Storage Foundation and High Availability Solutions .....                               | 11 |
|                         | Introduction .....   | 11 |
|                         | Changes introduced in 5.1 SP1 RP1 .....  | 12 |
|                         | Oracle VM Server 2.0 for SPARC .....   | 12 |
|                         | About rolling back the RP installation .....   | 12 |
|                         | Use the installrp or uninstallrp script with the -version option to determine product versions ..... | 12 |
|                         | VxVM and ASM co-existence enablement (2194492) .....   | 13 |
|                         | IMF support for DB2 agent .....  | 13 |
|                         | System requirements .....  | 15 |
|                         | Supported AIX operating systems .....  | 16 |
|                         | Supported VCS agents .....   | 16 |
|                         | Database requirements .....  | 17 |
|                         | Recommended memory and swap space .....  | 18 |
|                         | List of products .....   | 18 |
|                         | Fixed issues .....   | 18 |
|                         | Veritas Volume Manager: Issues fixed in 5.1 SP1 RP1 .....  | 19 |
|                         | Veritas File System: Issues fixed in 5.1 SP1 RP1 .....   | 21 |
|                         | Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP1 .....                    | 22 |
|                         | Veritas Cluster Server: Issues fixed in 5.1 SP1 RP1 .....  | 22 |
|                         | Known issues .....   | 24 |
|                         | Issues related to installation .....   | 24 |
|                         | Veritas Storage Foundation known issues .....  | 25 |
|                         | Veritas Volume Manager known issues .....  | 30 |
|                         | Veritas Volume Replicator known issues .....   | 35 |
|                         | Veritas File System known issues .....   | 41 |
|                         | Veritas Cluster Server known issues .....  | 45 |
|                         | Veritas Storage Foundation for Databases (SFDB) tools known issues .....                             | 56 |
|                         | Veritas Storage Foundation for Oracle RAC known issues .....   | 58 |
|                         | Software limitations .....   | 60 |

|           |  |    |
|-----------|--|----|
|           | Veritas Storage Foundation software limitations .....  | 60 |
|           | Veritas Volume Manager software limitations .....  | 60 |
|           | Veritas Volume Replicator software limitations .....   | 61 |
|           | Veritas Cluster Server software limitations .....  | 63 |
|           | Veritas Storage Foundation for Databases tools software<br>limitations .....                       | 71 |
|           | List of patches .....  | 72 |
|           | Downloading the 5.1 SP1 RP1 archive .....  | 73 |
|           | About the installrp script .....   | 73 |
|           | The installrp script options .....   | 73 |
| Chapter 2 | Installing the products for the first time .....   | 77 |
|           | Installing the Veritas software using the script-based installer .....                             | 77 |
|           | Installing Veritas software using the Web-based installer .....                                    | 78 |
|           | Starting the Veritas Web-based installer .....   | 79 |
|           | Obtaining a security exception on Mozilla Firefox .....  | 79 |
|           | Installing products with the Veritas Web-based installer .....                                     | 79 |
|           | Upgrading Veritas product with the Veritas Web-based<br>installer .....                            | 80 |
| Chapter 3 | Upgrading to 5.1 SP1 RP1 .....   | 81 |
|           | Prerequisites for upgrading to 5.1 SP1 RP1 .....   | 81 |
|           | Supported upgrade paths .....  | 81 |
|           | Upgrading from 5.1 SP1 to 5.1 SP1 RP1 .....  | 82 |
|           | Performing a full upgrade to 5.1 SP1 RP1 on a cluster .....  | 82 |
|           | Upgrading to 5.1 SP1 RP1 on a standalone system .....  | 90 |
|           | Performing a rolling upgrade using the script-based<br>installer .....                             | 92 |
|           | Verifying software versions .....  | 94 |
| Chapter 4 | Rolling back Veritas Storage Foundation and High<br>Availability Solutions .....                   | 95 |
|           | About rolling back Veritas Storage Foundation and High Availability<br>Solutions 5.1 SP1 RP1 ..... | 95 |
|           | Prerequisites for rolling back .....   | 95 |
|           | Committing files prior to rolling back .....   | 96 |
|           | Rolling back using the uninstallrp script .....  | 96 |
|           | Rolling back manually .....  | 97 |
|           | Rolling back Storage Foundation or Storage Foundation and High<br>Availability manually .....      | 97 |

|   |     |
|---|-----|
| Rolling back Storage Foundation Cluster File System<br>manually ..... | 99  |
| Rolling back Storage Foundation for Oracle RAC manually .....         | 101 |
| Rolling back Veritas Cluster Server manually .....                    | 103 |
| Rolling back Dynamic Multi-Pathing manually .....                     | 106 |



# About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [Introduction](#)
- [Changes introduced in 5.1 SP1 RP1](#)
- [System requirements](#)
- [List of products](#)
- [Fixed issues](#)
- [Known issues](#)
- [Software limitations](#)
- [List of patches](#)
- [Downloading the 5.1 SP1 RP1 archive](#)
- [About the installrp script](#)

## Introduction

This document provides information about the Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 Rolling Patch 1 release.

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

<http://entsupport.symantec.com/docs/334998>

The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit:

<http://entsupport.symantec.com/docs/330441>

Before installing or upgrading Veritas Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

## Changes introduced in 5.1 SP1 RP1

This section lists the changes in 5.1 SP1 RP1.

### Oracle VM Server 2.0 for SPARC

Storage Foundation High Availability Solutions is now qualified with Oracle VM Server 2.0 for SPARC.

See the *Veritas Storage Foundation and High Availability Solutions 5.1 SP1 Virtualization Guide* for supported use cases.

### About rolling back the RP installation

Use the `uninstallrp` script when you want to remove this Veritas rolling patch (RP) from systems. Once you complete the roll back process, your systems revert to the previous version of the product. You can use the roll back feature with the following products: Storage Foundation, Storage Foundation and High Availability, Veritas Storage Foundation Cluster File System, Storage Foundation for Oracle RAC, Veritas Cluster Server, Symantec VirtualStore, and Veritas Dynamic Multi-Pathing.

See “[About rolling back Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1](#)” on page 95.

### Use the `installrp` or `uninstallrp` script with the `-version` option to determine product versions

To determine a product's version, use the `-version` option with `installrp` or `uninstallrp`. After you install 5.1 SP1 RP1, only the `installrp` and `uninstallrp` scripts can detect 5.1 SP1 RP1 versions.

## VxVM and ASM co-existence enablement (2194492)

The VxVM and ASM co-existence was disabled by default. This co-existence is now enabled. So VxVM will identify ASM disks and display the same accordingly.

## IMF support for DB2 agent

Added Intelligent Monitoring Framework (IMF) capability and support for intelligent resource monitoring for DB2 agent. With IMF, VCS supports intelligent resource monitoring in addition to poll-based monitoring. Poll-based monitoring polls the resources periodically, whereas intelligent monitoring performs asynchronous monitoring. You can enable or disable the intelligent resource monitoring functionality of the agent for DB2.

See the *5.1 SP1 Veritas Cluster Server Administrator's Guide* for more information about:

- IMF notification module functions
- Administering the AMF kernel driver

### Before enabling the agent for IMF

Before you enable the DB2 agent for IMF, ensure that the AMF kernel module is loaded and AMF is configured. For details see the **Administering the AMF kernel driver** section of the *5.1 SP1 Veritas Cluster Server Administration Guide*.

### How the DB2 agent supports intelligent resource monitoring

When an IMF-enabled agent starts up, the agent initializes the asynchronous monitoring framework (AMF) kernel driver. After the resource is in a steady state, the agent registers with the AMF kernel driver, the details of the resource that are required to monitor the resource. For example, the agent for DB2 registers the PIDs of the DB2 processes with the IMF notification module. The agent's 'imf\_getnotification' function waits for any resource state changes. When the AMF kernel driver module notifies the `imf_getnotification` function about a resource state change, the agent framework runs the **monitor** agent function to ascertain the state of that resource. The agent notifies the state change to VCS, which then takes appropriate action. See the *5.1 SP1 Veritas Cluster Server Administrator's Guide* for more information.

### Agent functions for the IMF functionality

If the DB2 agent is enabled for IMF, the agent supports the following additional functions.

### **imf\_init**

This function initializes the DB2 agent to interface with the AMF kernel driver, which is the IMF notification module for the agent for DB2. This function runs when the agent starts up.

### **imf\_getnotification**

This function gets notifications about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.

### **imf\_register**

This function registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into a steady online state.

## **Attributes that enable IMF**

If the agent for DB2 is enabled for IMF, the agent uses the following type-level attributes in addition to the attributes described in DB2 agent installation and configuration guide.

## **IMF**

This resource type-level attribute determines whether the DB2 agent must perform intelligent resource monitoring. You can also override the value of this attribute at the resource level. This attribute includes the following keys:

### **Mode**

Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows:

- 0—Does not perform intelligent resource monitoring
- 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources
- 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources
- 3—Performs intelligent resource monitoring for both online and for offline resources

---

**Note:** The agent for DB2 supports intelligent resource monitoring for online resources only. Hence, Mode should be set to either 0 or 2.

---

Type and dimension: integer-association

Default: 0

### MonitorFreq

This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer.

Default: 1

You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:

- After every (*MonitorFreq*  $\times$  *MonitorInterval*) number of seconds for online resources

### RegisterRetryLimit

If you enable intelligent resource monitoring, the agent invokes the `imf_register` agent function to register the resource with the AMF kernel driver. The value of the `RegisterRetryLimit` key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the `Mode` key changes.

Default: 3

### IMFRegList

An ordered list of attributes whose values are registered with the IMF notification module.

Type and dimension: string-vector

Default: `static str IMFRegList[] = { DB2InstOwner, DB2InstHome }`

---

**Note:** In case of an upgrade to VCS 5.1SP1 RP1, please ensure that the new `Db2udbTypes.cf` file is used which contains the definition of **IMFRegList** as above.

---

## System requirements

This section describes the system requirements for this release

## Supported AIX operating systems

This section lists the supported operating systems for this release of Veritas products.

Product installation scripts verify the required update levels. The installation process terminates if the target systems do not meet the maintenance level requirements.

The minimum system requirements for this release are as follows:

For Power 7 processors at one of the following levels:

- AIX 6.1 TL5 with Service Pack 1 or later
- AIX Version 5.3 executing in POWER6 or POWER6+ compatibility at the following levels:
  - TL12 or later
  - TL11 with Service Pack 2 or later
  - TL10 with Service Pack 4 or later

For Power 6 or earlier processors at one of the following levels:

- AIX 6.1 TL2 or later
- AIX 5.3 at one of the following levels:
  - TL7 with SP6 or later
  - TL8 with SP4 or later

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

<http://entsupport.symantec.com/docs/334998>

## Supported VCS agents

Table 1-1 lists the agents for enterprise applications and the software that the agents support.

**Table 1-1** Supported software for the VCS agents for enterprise applications

| Agent | Application                   | Application version     | version          |
|-------|-------------------------------|-------------------------|------------------|
| DB2   | DB2 Enterprise Server Edition | 8.1, 8.2, 9.1, 9.5, 9.7 | AIX 5.3, AIX 6.1 |

**Table 1-1** Supported software for the VCS agents for enterprise applications  
*(continued)*

| Agent  | Application                       | Application version | version          |
|--------|-----------------------------------|---------------------|------------------|
| Oracle | Oracle                            | 11gR1, 10gR2, 11gR2 | AIX 5.3, AIX 6.1 |
| Sybase | Sybase Adaptive Server Enterprise | 12.5.x, 15.x        | AIX 5.3, AIX 6.1 |

See the *Veritas Cluster Server Installation Guide* for the agent for more details.

For a list of the VCS application agents and the software that the agents support, see the [Veritas Cluster Server Agents Support Matrix](#) at Symantec website.

## Database requirements

The following TechNote identifies the most current information on supported databases (Oracle Single Instance, DB2, and Sybase) and operating system combinations:

<http://entsupport.symantec.com/docs/331625>

---

**Note:** Veritas Storage Foundation (SF) and Veritas Storage Foundation Cluster File System (SFCFS) support running Oracle, DB2, and Sybase on VxFS and VxVM.

SF and SFCFS do not support running SFDB tools with DB2 and Sybase.

---

## Additional Oracle support for SF Oracle RAC

**Table 1-2** Oracle RAC versions that SF Oracle RAC supports

| Oracle version          | AIX 5.3 | AIX 6.1 |
|-------------------------|---------|---------|
| 10gR2 10.2 (64-bit)     | Yes     | Yes     |
| 11gR1 11.1 (64-bit)     | Yes     | Yes     |
| 11gR2 11.2.0.2 (64-bit) | No      | Yes     |

---

**Note:** For latest information on support for Oracle database versions with SF Oracle RAC, see the Technical Support Technote:

<http://entsupport.symantec.com/docs/280186>

---

## Recommended memory and swap space

Symantec recommends the following memory and swap space sizes:

- On the system where you run the installation, use the following guidelines for memory minimums when you install on:
  - One to eight nodes, use 1 GB of memory
  - More than eight nodes, use 2 GB of memory or more
- On the system where you run the installation, use the following guidelines for swap space when you install on:
  - One to eight nodes, use (*number of nodes* + 1) x 128 MB of free swap space
  - More than eight nodes, 1 GB of free swap space

## List of products

Apply this patch for the following Veritas Storage Foundation and High Availability products:

- Veritas Storage Foundation (SF)
- Veritas Storage Foundation and High Availability (SFHA)
- Veritas Storage Foundation Cluster File System (SFCFS)
- Veritas Storage Foundation Cluster File System and High Availability (SFCFSHA)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)
- Veritas Volume Manager (VM)
- Veritas File System (FS)
- Veritas Cluster Server (VCS)
- Veritas Dynamic Multi-Pathing (DMP)
- Symantec VirtualStore (SVS)

## Fixed issues

This section describes the issues fixed in this release.

## Veritas Volume Manager: Issues fixed in 5.1 SP1 RP1

This section describes the incidents that are fixed in Veritas Volume Manager in this release.

**Table 1-3** Veritas Volume Manager 5.1 SP1 RP1 fixed issues

| Fixed issues | Description   |
|--------------|---|
| 1426480      | VOLCVM_CLEAR_PR ioctl does not propogate the error returned by DMP to the caller  |
| 1829285      | vxconfigd core dumps while assigning unique native name to a disk   |
| 1869002      | Introduction of Circular buffer at vold level for master-slave communication.   |
| 1940052      | [cvm] Need rendezvous point during node leave completion  |
| 1959513      | Propogate -o noreonline option of diskgroup import to slave nodes   |
| 1970560      | When vxconfigd is idle (which is not shipping the command ) slave dies and command shipping is in progress, vxconfigd core dumped on Master |
| 2015467      | Performance improvement work for NetBackup 6.5.5 on SF 5.1 VxVM mapping provider  |
| 2038928      | creation of pre 5.1 SP1 (older) version diskgroup fails   |
| 2080730      | vxvm/vxdmp exclude file contents after updation should be consistent via vxdiskadm and vxdmpadm   |
| 2082450      | In case of failure, vxdisk resize should display more meaningful error message  |
| 2088007      | Possibility of reviving only secondary paths in DMP   |
| 2105547      | tagmeta info records are not cleaned-up during DGSJ operation and leading to huge delay in DGSJ operation after few iterations              |
| 2129477      | vxdisk reclaim command fails after resize operation.  |
| 2129989      | EVA ASL should report an error message if pref_bit is not set for a LUN   |
| 2133503      | Renaming enclosure results in dmpevents.log reporting Mode for Enclosure has changed from Private to Private                                |
| 2148682      | while shipping a command node hangs in master selection on slave nodes and master update on master node                                     |
| 2149532      | Enabling storage keys with ldata code in DMP  |

**Table 1-3** Veritas Volume Manager 5.1 SP1 RP1 fixed issues (*continued*)

| Fixed issues | Description  |
|--------------|--|
| 2158438      | vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core   |
| 2159947      | Bump up the dmpslab_minsz to 512 elements  |
| 2160199      | Master takeover fails as the upcoming Master could not import shared DG  |
| 2164988      | After upgrading from 5.1 to 5.1 SP1 with rootability enabled, root support may not get retained.                                     |
| 2166682      | checks needed to make sure that a plex is active before reading from it during fsmv mirror read interface                            |
| 2172488      | FMR: with dco version 0 restore operation doesn't sync the existing snapshot mirrors   |
| 2176601      | SRDF-R2 devices are seen in error state when devices are in write-protected mode   |
| 2181631      | Striped-mirror volume cannot be grown across sites with -oallowspansites with DRL  |
| 2181877      | System panic due to absence of KEY_PRIVATE1 storage key in single path iodone  |
| 2183984      | System panics due to race condition while updating DMP I/O statistics  |
| 2188590      | An ilock acquired by a SLAVE node for a read on a DCL object can lead to IO hang when the node becomes MASTER before reading is done |
| 2191693      | 'vxdkmpadm native list' command is not displaying any output nor error   |
| 2194492      | VxVM-ASM co-existence enablement 2062190 vxrootadm split/join operation fails when there is a rvg present in the root/back upgrade   |
| 2199496      | Data Corruption seen with "site mirror" Campus Cluster feature   |
| 2200670      | vxattachd does not recover disks if disk group is not imported   |
| 2201149      | DMP should try all possibilities to service I/O upon receipt of a SCSI illegal request following HBA fault                           |
| 2218706      | Support for MAXCPU on Power 7  |
| 2226813      | VVR: rlinks remain disconnected with UDP protocol if data ports are specified  |
| 2227923      | renaming of enclosure name is not persistent   |

**Table 1-3** Veritas Volume Manager 5.1 SP1 RP1 fixed issues (*continued*)

| Fixed issues | Description                                |
|--------------|--|
| 2234844      | asm2vxfs conversion fails                  |
| 2215216      | vxkprint does not report TP related values |

## Veritas File System: Issues fixed in 5.1 SP1 RP1

This section describes the incidents that are fixed in Veritas File System in this release.

**Table 1-4** Veritas File System fixed issues

| Incident | Description   |
|----------|---|
| 1929221  | vxrepquota truncating username and groupname to 8 characters is addressed.  |
| 2030119  | fspadm core dumps when analysing a badly formatted XML file, is resolved  |
| 2162822  | During online migration from ufs to vxfs, df command returns a non-zero return value.   |
| 2169273  | During online migration, nfs export of the migrating file system leads to system panic  |
| 2177253  | A warning message is displayed when mounting a fs using disk layout version 4 or 5, to indicate that mount of layout versions 4 and 5 are supported only for vxupgrade purposes |
| 2178147  | Linking a IFSOC file now properly calls vx_dotdot_op(), which fixes the cause of a corrupted inode.   |
| 2184528  | fsck no longer fails to repair corrupt directory blocks that have duplicate directory entries.  |
| 2194618  | Link operations on socket files residing on vxfs leads to incorrectly setting fsck flag on the file system  |
| 2215377  | Perf issue due to memory/glm  |
| 2221623  | Fixed a performance loss due to a delxwri_ilst spin lock with the default values for vx_idelxwri_timelag.   |

## Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP1

This section describes the incidents that are fixed in Veritas Storage Foundation Cluster File System in this release.

**Table 1-5** Veritas Storage Foundation Cluster File System fixed issues

| Incident | Description  |
|----------|--|
| 1296491  | Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted   |
| 2149659  | In the presence of file system checkpoints containing snapshot files, truncate operation of a file with shared extent results in hitting f:xted_validate_cuttran:10 or vx_te_mklbtran:1b |
| 2153512  | cfs freeze ioctl hang due to mdele lock not being released during an error condition, is resolved.   |
| 2169538  | The cfsmntadm add command fails, if one host name is a substring of another host name in the list  |
| 2180905  | fsadm -S shared mountpoint gives wrong error message when mount points of veritas filesystem version is other than 8.  |
| 2181833  | "vxfilesnap" gives wrong error message on checkpoint filesystem on cluster   |
| 2184114  | In a large filesystem, stat operation on cluster mount leaves the file system frozen for too long leading to CVMVoldg and CFSSMount agent timeouts.                                      |
| 2203917  | ODM I/O performance bottleneck due to threads waiting in odm_rwsleep_lock() to get lock on odm_iop_table is resolved   |
| 2232554  | System panic in vx_iupdat_clustblks() due to an unprotected inode getting corrupted.   |

## Veritas Cluster Server: Issues fixed in 5.1 SP1 RP1

This section describes the incidents that are fixed in Veritas Cluster Server in this release.

**Table 1-6** Veritas Cluster Server 5.1 SP1 RP1 fixed issues

| Fixed Issues | Description  |
|--------------|--|
| 1999058      | RVGSnapshot: DR Fire Drills support Oracle RAC environments using VVR.   |
| 2011536      | Db2 IMF Integration for NON MPP PRON).   |
| 2179652      | The monitor script of Db2udb do not handle the case when a parameter is undefined, which make an empty value being passed to next level.                                       |
| 2180721      | IPMultiNICB: haipswitch does not support AIX version 6.  |
| 2180759      | Add WorkLoad attribute to WPAR.xml.  |
| 2184205      | Parent service group does not failover in case of online local firm dependency with child service group.   |
| 2185494      | Panic issue related to fp_close().   |
| 2194473      | HAD dumping core while overriding the static attribute to resource level.  |
| 2205556      | DNS Agent: The offline EP does not remove all A/AAAA records if OffDelRR=1 for Multi-home records  |
| 2205563      | DNS Agent: Clean EP does not remove any resource records for OffDelRR=1.   |
| 2205567      | DNS Agent: master.vfd fails to query dns server  |
| 2209337      | RemoteGroup agent crashes if VCSAPI log level is set to non zero value.  |
| 2210489      | cfs.noise.n1 test hit the assert "xtpw_inactive_free:1c xtpw_free is empty!"   |
| 2214539      | When node reboots sometimes the intentonline of group is set to 2 even if group is online elsewhere. This later causes group to consider autostartlist and not doing failover. |
| 2218556      | cpsadm should not fail if llt is not installed/configured on a single node cluster.  |
| 2218565      | MonitorTimeStats incorrectly showing 303 secs Intermittently.  |
| 2219955      | Split-brain occurs even when using VCS Steward.  |
| 2220317      | Application agent clean script fails to work when using PidFiles due to bad use of array.  |
| 2221618      | Fixed an issue where Cluster Manager (Java Console) was not encrypting the "DBAPword" attribute of Oracle Agent.   |
| 2223135      | nfs_sg fail when execute hastop -all.  |

**Table 1-6** Veritas Cluster Server 5.1 SP1 RP1 fixed issues (*continued*)

| Fixed Issues | Description   |
|--------------|---|
| 2238968      | LLT: disable the fastpath in LLT and make it optional.  |
| 2241419      | halogin does not work in secure environment where Root broker is not VCS node.  |
| 2244148      | Fixed an issue with Netlsnr agent where not specifying the container name would result into core dump if debug logs were enabled. |

## Known issues

This section covers the known issues in this release.

### Issues related to installation

This section describes the known issues during installation and upgrade.

#### **EULA changes (2161557)**

The locations for all EULAs have changed.

The English EULAs now appear in */product\_dir/EULA/en/product\_eula.pdf*

The EULAs for Japanese and Chinese now appear in those language in the following locations:

The Japanese EULAs appear in */product\_dir/EULA/ja/product\_eula.pdf*

The Chinese EULAs appear in */product\_dir/EULA/zh/product\_eula.pdf*

#### **NBU 6.5 or older version is installed on a VxFS file system (2056282)**

NBU 6.5 or older version is installed on a VxFS file system. Before upgrading to SF5.1, user umounts all VxFS file systems including the one which hosts NBU binaries(/usr/opensv). While upgrading SF5.1, installer fails to check out NBU is installed on the same machine and uninstalls the shared infrastructure packages: VRTSspb, VRTSsat, VRTSsisco which cause NBU not working.

**Workaround:** Before you umount the VxFS file system which hosts NBU, copy the two files /usr/opensv/netbackup/bin/version and /usr/opensv/netbackup/version to /tmp directory. After you umount the NBU file system, manually copy these two version files from /tmp to their original path.

If the path doesn't exist, make the same directory path with the command: `mkdir -p /usr/openv/netbackup/bin` and `mkdir -p /usr/openv/netbackup/bin`. Run the installer to finish the upgrade process. After upgrade process is done, remove the two version files and their directory paths.

How to recover systems already affected by this issue: Manually install VRTSpxb, VRTSat, VRTSisco packages after the upgrade process is done.

### **During product migration the installer overestimates disk space use (2088827)**

The installer displays the space that all the product and patches needs. During migration some are already installed and during migration some are removed. This releases disk space. The installer then claims more space than it actually needs.

**Workaround:** Run the installer with `-nospacecheck` option if the disk space is less than that installer claims but more than actually required.

### **The VRTSacclib is deprecated (2032052)**

The VRTSacclib is deprecated. For installation, uninstallation, and upgrades, note the following:

- Fresh installs: Do not install VRTSacclib.
- Upgrade: Uninstall old VRTSacclib and install new VRTSacclib.
- Uninstall: Ignore VRTSacclib.

### **Ignore VRTSgms request to boot during installation (2143672)**

During installation, you may see this error which you can ignore.

```
VRTSgms: old driver is still loaded...  
VRTSgms: You must reboot the system after installation...
```

## Veritas Storage Foundation known issues

This section describes the known issues in this release of Veritas Storage Foundation (SF).

### **db2exp may frequently dump core (1854459)**

If a host is configured to an SFM central server with DB2 version 9.x, then the command-line interface db2exp may frequently dump core.

**Workaround:** There is a hotfix patch available for this issue. Contact Symantec Technical Support for the hotfix patch.

## In an IPv6 environment, db2icrt and db2idrop commands return a segmentation fault error during instance creation and instance removal (1602444)

When using IBM DB2 `db2icrt` command to create a DB2 database instance on a pure IPv6 environment, the `db2icrt` command returns segmentation fault error message. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2icrt -a server -u db2fen1 db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 4700: 26182 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -addfcm -i ${INSTNAME?}
```

The `db2idrop` command also returns segmentation fault, but the instance is removed successfully after the `db2idrop` command is issued. For example:

```
$ /opt/ibm/db2/V9.5/instance/db2idrop db2inst1
/opt/ibm/db2/V9.5/instance/db2iutil: line 3599: 7350 Segmentation fault
$ {DB2DIR?}/instance/db2isrv -remove -s DB2_${INSTNAME?} 2> /dev/null
```

```
DBI1070I Program db2idrop completed successfully.
```

This happens on DB2 9.1, 9.5, and 9.7.

This issue has been identified as an IBM issue. Once IBM has fixed this issue, then IBM will provide a hotfix for this segmentation problem.

At this time, you can communicate in a dual-stack to avoid the segmentation fault error message until IBM provides a hotfix.

### To communicate in a dual-stack environment

- ◆ Add an IPv6 hostname as an IPv4 loopback address to the `/etc/hosts` file. For example:

```
127.0.0.1 swlx20-v6
```

Or

```
127.0.0.1 swlx20-v6.punipv6.com
```

where `127.0.0.1` is the IPv4 loopback address.

where `swlx20-v6` and `swlx20-v6.punipv6.com` is the IPv6 hostname.

## AT Server crashes when authenticating unixpwd user multiple times (1705860)

There is a known issue in the AIX kernel code that causes 'getgrent\_r' function to corrupt the heap. This issue is present in AIX 5.3 and AIX 6.1 Refer to IBM's Web site for more information:

<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52585>

AT uses `getgrent_r` function to get the groups of the authenticated user.

IBM has released the fix as a patch to `fileset bos.rte.libc`. There are different patches available for different version of `bos.rte.libc`. You need to check the version of `bos.rte.libc` (For example: `lslpp -l grep bos.rte.libc`) and apply the appropriate IBM patch:

- For version 6.1.3.1:  
<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52959>  
For the fix:  
<ftp://ftp.software.ibm.com/aix/efixes/iz52959/>
- For version 6.1.2.4:  
<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52720>  
For the fix:  
<ftp://ftp.software.ibm.com/aix/efixes/iz52720/>
- For version 6.1.2.5 :  
<http://www-01.ibm.com/support/docview.wss?uid=isg1IZ52975>  
For the fix:  
<ftp://ftp.software.ibm.com/aix/efixes/iz52975/>

There are IBM patches for only certain version of `bos.rte.libc` that are available. If your system has a different `bos.rte.libc` version, you may have to upgrade to a higher version where the fix is available. If your version is not available, you may have to contact IBM.

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

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

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

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

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

## Not all the objects are visible in the SFM GUI (1821803)

After upgrading SF stack from 5.0MP3SP1 RP1 to 5.1, the volumes are not visible under the Volumes tab and the shared diskgroup is discovered as Private and Deported under the Disgroup tab in the SFM GUI.

**Workaround:**

To resolve this known issue

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

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

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

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

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

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

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

## DB2 databases are not visible from the SFM Web console (1850100)

If you upgraded to SF 5.1, DB2 databases will be not visible from the SFM web console.

This will be fixed in the SF 5.1 Patch 1 release.

**Workaround:** Reinstall is required for SFM DB2-Hotfix (HF020008500-06.sfa), if the host is upgraded to SF 5.1 Use the deployment framework and reinstall the hotfix for DB2 (HF020008500-06.sfa) on the managed host.

### To resolve this issue

- 1 In the Web GUI, go to **Settings > Deployment**.
- 2 Select **HF020008500-06 hotfix**.
- 3 Click **Install**.
- 4 Check the **force** option while reinstalling the hotfix.

## A volume's placement class tags are not visible in the Veritas Enterprise Administrator GUI when creating a dynamic storage tiering placement policy (1880622)

A volume's placement class tags are not visible in the Veritas Enterprise Administrator (VEA) GUI when you are creating a dynamic storage tiering (DST) placement policy if you do not tag the volume with the placement classes prior to constructing a volume set for the volume.

**Workaround:** To see the placement class tags in the VEA GUI, you must tag the volumes prior to constructing the volume set. If you already constructed the volume set before tagging the volumes, restart `vxsvc` to make the tags visible in the GUI.

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

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

```
alt_disk_copy: 0505-224 ATTENTION:  
An error occurred during installation of  
one or more software components.  
Modifying ODM on cloned disk.  
Building boot image on cloned disk.  
forced unmount of /alt_inst/var/adm/ras/platform
```

```
forced unmount of /alt_inst/var  
umount: error unmounting /dev/alt_hd2: Device busy  
0505-144 alt_disk_install: Unable to unmount alt_inst filesystems.
```

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

## Veritas Volume Manager known issues

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

### **The vxcdsconvert utility is not supported for EFI disks (2064490)**

The vxcdsconvert utility is not supported for EFI disks.

### **Subpaths be marked as DISABLED after lun failover/failback in array on 6.1TL6(2242364)**

It may be possible that secondary paths of a AP/F array goes into disable state in case of change of ownership of the paths. This happens only when ownership change is triggered from outside of the dmp (or from second node of SFHA cluster). The restore daemon should bring the disabled path back to online state or one can run vxctl enable command to bring the disable path back to online.

### **Machine having CPUs >128 may get panicked after uninstallrp (2246837)**

Intermittent failures or system crashes might occur if VRTSvxvm level is rolledback to 5.1 SP1 on a system having more than 128 CPUs.

It is recommended to maintain the VRTSvxvm version as 5.1SP1RP1 or 5.1SP1P2

### **The cluster may hang if a node goes down (1835718)**

The cluster may hang if a node goes down while one array is disabled or offline in a mirror=enclosure configuration.

This may occur, if a node panics or loses power while one array of a mirror=enclosure configuration is offline or disabled, then the cluster, fencing, I/O loads, and VxVM transactions hang.

There is no workaround at this time.

## After installing Volume Manager, you may be prompted to reinstall it (1704161)

If you remove pre-5.1 Volume Manager packages and then install 5.1 Volume Manager without using the product installer, the following message is displayed:

```
The Volume Manager appears to be installed already. You should use
vxdiskadm to add more disks to the system. Installation with vxinstall
will attempt to reinstall the Volume Manager from the beginning.
Depending upon how your system is currently configured, a
reinstallation may fail and could leave your system unusable.
```

```
Are you sure you want to reinstall [y,n,q,?] (default: n)
```

### Workaround

When you are prompted to reinstall, enter y.

Note: This message is not displayed if you install Volume Manager with the product installer.

## vxconvert failures if PowerPath disks are formatted as simple disks (857504)

If a PowerPath disk is formatted as a simple disk (a foreign device), then the vxconvert utility may fail during conversion of LVM to VxVM. To view the format of the disk, use the vxdisk list command. This issue may also occur if the /etc/vx/darecs file contains an hdiskpower disk entry. This entry may be present if PowerPath disks were configured as foreign disks in Storage Foundation 4.0, and the entry was not changed after subsequent upgrades.

## Veritas Volume Manager (VxVM) might report false serial split brain under certain scenarios (1834513)

VxVM might detect and report a false serial split brain when all of the following conditions are met:

- One or more arrays that provide the shared storage for the cluster are being powered off
- At the same time when the arrays are being powered off, an operation that requires an internal transaction is initiated (such as VxVM configuration commands)

In such a scenario, disk group import will fail with a split brain error and the vxsplitlines output will show 0 or 1 pools.

### Workaround:

### To recover from this situation

- 1 Retrieve the disk media identifier (`dm_id`) from the configuration copy:

```
# /etc/vx/diag.d/vxprivutil dumpconfig device-path
```

The `dm_id` is also the serial split brain id (`ssbid`)

- 2 Use the `dm_id` in the following command to recover from the situation:

```
# /etc/vx/diag.d/vxprivutil set device-path ssbid=dm_id
```

### **vxdisk -f init can overwrite some of the public region contents (1190117)**

If a disk was initialized by a previous VxVM version or defined with a smaller private region than the new default of 32Mb, then the public region data will be overridden. The work around is to specify explicitly the length of `privoffset`, `puboffset`, `publen`, `privlen` while initializing the disk.

### **The relayout operation fails when there are too many disks in the disk group. (2015135)**

The attempted relayout operation on diskgroup containing approximately more than 300 Luns/disks may fail with error 'Cannot setup space

### **Co-existence check might fail for CDS disks**

In Veritas Volume Manager (VxVM) 5.1 SP1 RP1, VxVM introduces the ability to support Cross-platform Data Sharing (CDS) on disks larger than 1 TB. VxVM uses the SUN VTOC Table to initialize the `cdsdisk` layout on devices up to 1 TB. VxVM uses the GUID Partition Table (GPT) to initialize the `cdsdisk` layout on devices larger than 1 TB.

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

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

**Workaround:** There is no workaround for this issue.

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

If a path loses connectivity to the array, the path is marked with the NODE\_SUSPECT flag. After the connectivity is restored, the restore daemon detects that the path is restored when the restore daemon probes the paths. The restore daemon clears the NODE\_SUSPECT flag and makes the path is available for I/O. The restore daemon probes the paths at the interval set with the tunable parameter `dmp_restore_interval`. If you set the `dmp_restore_interval` parameter to a high value, the paths are not available for I/O until the next interval.

## **Node is not able to join the cluster with high I/O load on the array with VCS (2124595)**

When the array has a high I/O load, the DMP database exchange between master node and joining node takes a longer time. This situation results in VCS resource online timeout, and then VCS stops the join operation.

### **Workaround:**

Increase the online timeout value for the HA resource to 600 seconds. The default value is 300 seconds.

### **To set the OnlineTimeout attribute for the HA resource type CVMCluster**

- 1 Make the VCS configuration to be ReadWrite:

```
# haconf -makerw
```

- 2 Change the OnlineTimeout attribute value of CVMCluster:

```
# hatype -modify CVMCluster OnlineTimeout 600
```

- 3 Display the current value of OnlineTimeout attribute of CVMCluster:

```
# hatype -display CVMCluster -attribute OnlineTimeout
```

- 4 Save and close the VCS configuration:

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

## Changes in enclosure attributes are not persistent after an upgrade to VxVM 5.1 SP1 RP1 (2082414)

The Veritas Volume Manager (VxVM) 5.1 SP1 RP1 includes several array names that differ from the array names in previous releases. Therefore, if you upgrade from a previous release to VxVM 5.1 SP1 RP1, changes in the enclosure attributes may not remain persistent. Any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 5.1 SP1 RP1. Manually reconfigure the enclosure attributes to resolve the issue.

[Table 1-7](#) shows the Hitachi arrays that have new array names.

**Table 1-7** Hitachi arrays with new array names

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

In addition, the Array Support Library (ASL) for the enclosures XIV and 3PAR now converts the cabinet serial number that is reported from Hex to Decimal, to correspond with the value shown on the GUI. The persistence of the enclosure name is achieved with the `/etc/vx/array.info` file, which stores the mapping between cabinet serial number and array name. Because the cabinet serial number has changed, any enclosure attribute set for these arrays may be reset to the default value after an upgrade to VxVM 5.1 SP1 RP1. Manually reconfigure the enclosure attributes to resolve the issue.

The cabinet serial numbers are changed for the following enclosures:

- IBM XIV Series arrays
- 3PAR arrays

## DS4K series array limitations

In case of DS4K array series connected to AIX host(s), when all the paths to the storage are disconnected and reconnected back, the storage does not get discovered

automatically. To discover the storage, run the `cfgmgr OS` command on all the affected hosts. After the `cfgmgr` command is run, the DMP restore daemon brings the paths back online automatically in the next path restore cycle. The time of next path restore cycle depends on the restore daemon interval specified (in seconds) by the tunable `dmp_restore_interval`.

```
# vxddmpadm gettune dmp_restore_interval
          Tunable                Current Value  Default Value
-----
dmp_restore_interval            300           300
```

On DS4K array series connected to AIX host(s) DMP is supported in conjunction with RDAC. DMP is not supported on DS4K series arrays connected to AIX hosts in MPIO environment.

### **vxconfigd hang with path removal operation while IO is in-progress (1932829)**

In AIX with system firmware version SF240\_320, `vxdisk scandisks` (device discovery) takes a long time when a path is disabled from the switch or from the array.

**Workaround:**

To resolve this issue, upgrade the system firmware version to SF240\_382.

### **EFI labelled disks of size greater than 2TB size; will not be supported. (2270880)**

On Solaris 10, if the size of EFI labelled disk is greater than 2TB, the disk capacity will be truncated to 2TB when it is initialized under Veritas Volume Manager.

**Workaround:** There is no workaround for this issue.

## **Veritas Volume Replicator known issues**

This section describes the known issues in this release of Veritas Volume Replicator (VVR).

### **vradmin syncvol command compatibility with IPv6 addresses (2075307)**

The `vradmin syncvol` command does not work with the compressed form of IPv6 addresses. In IPv6 environments, if you run the `vradmin syncvol` command and identify the target host using compressed form of the IPv6 address, the command fails with following error message:

```
# vradmin -s -full syncvol vol1 fe80::221:5eff:fe49:ad10:dgl:vol1
VxVM VVR vradmin ERROR V-5-52-420 Incorrect format for syncvol.
```

Also, if you run the `vradmin addsec` command and you specify the Secondary host using the compressed IPv6 address, the `vradmin syncvol` command also fails – even if you specify the target as `hostname`.

**Workaround:** When you use the `vradmin addsec` and `vradmin syncvol` commands, do not specify compressed IPv6 addresses; instead, use hostnames.

## **RVGPrimary agent operation to start replication between the original Primary and the bunker fails during failback (2054804)**

The RVGPrimary agent initiated operation to start replication between the original Primary and the bunker fails during failback – when migrating back to the original Primary after disaster recovery – with the error message:

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

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

### **Workaround:**

#### **To resolve this issue**

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

## **Bunker replay did not occur when the Application Service Group was configured on some of the systems in the Primary cluster, and ClusterFailoverPolicy is set to "AUTO" (2047724)**

The time that it takes for a global cluster to fail over an application service group can sometimes be smaller than the time that it takes for VVR to detect the configuration change associated with the primary fault. This can occur in a bunkered, globally clustered configuration when the value of the `ClusterFailoverPolicy` attribute is `Auto` and the `AppGroup` is configured on a subset of nodes of the primary cluster.

This causes the RVGPrimary online at the failover site to fail. The following messages appear the VCS engine log:

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname could not be
imported on bunker host hostname. Operation failed with error 256
and message VxVM VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote
server unreachable... Timestamp VCS ERROR V-16-2-13066 (hostname)
Agent is calling clean for resource(RVGPrimary) because the resource
is not up even after online completed.
```

#### **Workaround:**

##### **To resolve this issue**

- ◆ When the configuration includes a bunker node, set the value of the `OnlineRetryLimit` attribute of the RVGPrimary resource to a non-zero value.

## **Interrupting the vradmin syncvol command may leave volumes open (2063307)**

Interrupting the `vradmin syncvol` command may leave volumes on the Secondary site in an open state.

**Workaround:** On the Secondary site, restart the `in.vxrsyncd` daemon. Enter the following:

```
# /etc/init.d/vxrsyncd.sh stop

# /etc/init.d/vxrsyncd.sh start
```

## **The RVGPrimary agent may fail to bring the application service group online on the new Primary site because of a previous primary-elect operation not being run or not completing successfully (2043831)**

In a primary-elect configuration, the RVGPrimary agent may fail to bring the application service groups online on the new Primary site, due to the existence of previously-created instant snapshots. This may happen if you do not run the `ElectPrimary` command to elect the new Primary or if the previous `ElectPrimary` command did not complete successfully.

**Workaround:** Destroy the instant snapshots manually using the `vxrvg -g dg -P snap_prefix snapdestroy rvg` command. Clear the application service group and bring it back online manually.

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

### Issue 1:

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

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

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

### Issue 2:

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

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

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

**Workaround:** The following workarounds resolve these issues.

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

For example:

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

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

For example:

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

## Veritas product 5.0MP3 Rolling Patch 2 required for replication between 5.0 MP3 and 5.1 SP1 (1800600)

In order to replicate between Primary sites running Veritas product 5.0 MP3 and Secondary sites running Veritas product 5.1 SP1, or vice versa, you must install the Veritas product 5.0MP3 Rolling Patch 2 on the nodes using 5.0MP3. This patch resolves several outstanding issues for replicating between versions.

## In an IPv6-only environment RVG, data volumes or SRL names cannot contain a colon

Issue: After upgrading VVR to an IPv6-only environment in 5.1 release, `vradmin` commands may not work when a colon is specified in the RVG, data volume(s) and/or SRL name. It is also possible that after upgrading VVR to an IPv6-only environment, `vradmin createpri` may dump core when provided with RVG, volume and/or SRL names containing a colon in it.

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

## `vradmin` commands might fail on non-logowner node after logowner change (1810827)

When VVR is used for replicating shared disk groups in an SFCFS or SFRAC environment consisting of three or more nodes, a logowner change event might, in rare instances, render `vradmin` commands unusable on some or all of the cluster nodes. In such instances, the following message appears in the "Config Errors:" section of the output of the `vradmin repstatus` and `vradmin printrvg` commands:

```
vradmind not reachable on cluster peer
```

In addition, all other `vradmin` commands (except `vradmin printvol`) fail with the error:

```
"VxVM VVR vradmin ERROR V-5-52-488 RDS has configuration error related to the master and logowner."
```

This is due to a defect in the internal communication sub-system, which will be resolved in a later release.

**Workaround:** Restart `vradmind` on all the cluster nodes using the following commands:

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

## While vradmin changeip is running, vradmind may temporarily lose heart beats (2162625)

This issue occurs when you use the `vradmin changeip` command to change the host name or IP address set in the Primary and Secondary RLINKs. While the `vradmin changeip` command runs, `vradmind` may temporarily lose heart beats, and the command terminates with an error message.

### Workaround:

#### To resolve this issue

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

```
export IPM_HEARTBEAT_TIMEOUT
IPM_HEARTBEAT_TIMEOUT=120
```

- 2 Restart `vradmind` to put the new `IPM_HEARTBEAT_TIMEOUT` value into affect. Enter the following:

## If using VEA to create a replicated data set fails, messages display corrupt strings in the Japanese locale (1726499, 1377599)

When using VEA to create a replicated data set, because the volumes do not have a DCM log on all nodes, the message window displays corrupt strings and unlocalized error messages.

**Workaround:** There is no workaround for this issue.

## vxassist relayout removes the DCM (2162522)

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

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

## vxassist and vxresize operations do not work with layered volumes that are associated to an RVG (2162579)

This issue occurs when you try a resize operation on a volume that is associated to an RVG and has a striped-mirror layout.

**Workaround:****To resize layered volumes that are associated to an RVG**

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

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

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

```
# vxvol -g diskgroup dis vol
```
- 5 Resize the volumes. In this example, the volume is increased to 10 GB. Enter the following:  

```
# vxassist -g diskgroup growto vol 10G
```
- 6 Associate the data volumes to the RVG. Enter the following:  

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

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

## Veritas File System known issues

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

### **VxFS read ahead can cause stalled I/O on all write operations (1965647)**

Changing the `read_ahead` parameter can lead to frozen I/O. Under heavy load, the system can take to several minutes to recover from this state.

**Workaround:** There is no workaround for this issue.

## Asynchronous cached ODM requests do not use the cache (2010139)

Asynchronous cached ODM requests do not use the cache on AIX, and you might observe some performance degradation only during async cached ODM request. However, synchronous reads will not be affected.

**Workaround:** There is no workaround for this issue.

## Shrinking a file system that is larger than 1 TB takes a long time (2097673)

Shrinking a file system shrink via either the `fsadm` command or `vxresize` command can take a long time to complete in some cases, such as if the shrink size is large and some large extent of a file is overlapping with the area to be shrunk.

**Workaround:** One possible workaround is to use the `vxtunefs` command and `setwrite_pref_io` and `write_nstream` to high values, such that `write_pref_io` multiplied by `write_nstream` is around 8 MB.

## Storage Checkpoints can exceed the quota limit (2102201)

Under some circumstances, Storage Checkpoints can exceed the quota limit set by `fsckptadm setquotalimit` command. This issue can arise if all of the following conditions are met:

- The Storage Checkpoint quota has been enabled.
- The Storage Checkpoint quota is not exceeded.
- A file content modification operation, including removing a file, needs to push some or all blocks of the file to the Storage Checkpoint.
- Number of blocks that need to be pushed to the Storage Checkpoint is enough to exceed Storage Checkpoint quota hard limit.

**Workaround:** There is no workaround for this issue.

## vxfsconvert can only convert file systems that are less than 1 TB (2108929)

The `vxfsconvert` command can only convert file systems that are less than 1 TB. If the file system is greater than 1` TB, the `vxfsconvert` command fails with the "Out of Buffer cache" error.

## Enabling the D\_REFUND parameter on AIX 6.1 causes a hang in some situations on a cluster file system (2166515)

In some situations, enabling the `D_REFUND` parameter on AIX 6.1 causes a hang on a cluster file system. Some example situations include creating a Storage Checkpoint, unmounting a file system after receiving an I/O error, and having a high GLM load.

**Workaround:** Disable the `D_REFUND` parameter.

## Panic due to null pointer de-reference in vx\_unlockmap() (2059611)

A null pointer dereference in the `vx_unlockmap()` call can cause a panic. A fix for this issue will be released in the a future patch.

**Workaround:** There is no workaround for this issue.

## The dynamic vmm buffer allocation feature requires certain AIX APARs to be installed (1849083)

VxFS supports the use of the dynamic vmm buffer allocation (`D_REFUND`) feature, which IBM added to AIX 6.1 TL2 and later releases of AIX. However, IBM fixed some issues in the `D_REFUND` feature through certain APARs, which you must install to use the `D_REFUND` feature with VxFS. The TL of the operating system determines which APAR you must install:

| Operating system | Required APAR                     |
|------------------|-----------------------------------|
| AIX 6.1 TL2      | IZ41494, which is packaged in SP3 |
| AIX 6.1 TL3      | IZ37627                           |
| AIX 6.1 TL4      | IZ38189                           |

## Possible write performance degradation with VxFS local mounts (1837394)

Some applications that allocate large files without explicit preallocation may exhibit reduced performance with the VxFS 5.1 release and later releases compared to the VxFS 5.0 MP3 release due to a change in the default setting for the tunable `max_seqio_extent_size`. One such application is DB2. Hosting DB2 on a single file system extent maximizes the potential for sequential pre-fetch processing. When DB2 detects an application performing sequential reads against database data, DB2 begins to read ahead and pre-stage data in cache using efficient sequential physical I/Os. If a file contains many extents, then pre-fetch processing

is continually interrupted, nullifying the benefits. A larger `max_seqio_extent_size` value reduces the number of extents for DB2 data when adding a data file into a tablespace without explicit preallocation.

The `max_seqio_extent_size` tunable controls the amount of space that VxFS automatically preallocates to files that are allocated by sequential writes. Prior to the 5.0 MP3 release, the default setting for this tunable was 2048 file system blocks. In the 5.0 MP3 release, the default was changed to the number of file system blocks equaling 1 GB. In the 5.1 release, the default value was restored to the original 2048 blocks.

The default value of `max_seqio_extent_size` was increased in 5.0 MP3 to increase the chance that VxFS will allocate the space for large files contiguously, which tends to reduce fragmentation and increase application performance. There are two separate benefits to having a larger `max_seqio_extent_size` value:

- Initial allocation of the file is faster, since VxFS can allocate the file in larger chunks, which is more efficient.
- Later application access to the file is also faster, since accessing less fragmented files is also more efficient.

In the 5.1 release, the default value was changed back to its earlier setting because the larger 5.0 MP3 value can lead to applications experiencing "no space left on device" (ENOSPC) errors if the file system is close to being full and all remaining space is preallocated to files. VxFS attempts to reclaim any unused preallocated space if the space is needed to satisfy other allocation requests, but the current implementation can fail to reclaim such space in some situations.

**Workaround:** If your workload has lower performance with the VxFS 5.1 release and you believe that the above change could be the reason, you can use the `vxtunefs` command to increase this tunable to see if performance improves.

#### To restore the benefits of the higher tunable value

- 1 Increase the tunable back to the 5.0 MP3 value, which is 1 GB divided by the file system block size.  
  
Increasing this tunable also increases the chance that an application may get a spurious ENOSPC error as described above, so change this tunable only for file systems that have plenty of free space.
- 2 Shut down any application that are accessing any large files that were created using the smaller tunable setting.
- 3 Copy those large files to new files, which will be allocated using the higher tunable setting.

- 4 Rename the new files back to the original names.
- 5 Restart any applications were shut down earlier.

### 5.1SP1 sol\_sprac Test LM-stress\_S5 hits an assert of "f:vx\_idelxwri\_off:5a vai vx\_trunc\_tran"(2169326)

When a removable clone mounted in “rw” mode, and if files from the mounted clone are being modified, operations on the mounted clone may fail if the file system is full. Even though the clone will continue to be seen mounted, no operation except an umount will be allowed. Also `fsckptadm list primary mount point` will not show the clone right away.

**Workaround:**It is advised to create non-removable clones to avoid encountering this issue.

## Veritas Cluster Server known issues

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

### NFS resource fails to come online with NFSv4 (2239020)

When the local domain is not set on machine, the NFS resource can fail to come online with NFSv4.

**Workaround:**

If you want to configure NFSv4 for NFS resource, ensure that the local domain is set for the cluster node.

**To set the local domain**

- ◆ Run the following command:

```
# chnfsdom domain_of_the_node
```

### Hang or crash issue in frmalloc recursive lock acquisition

Recursive calls to xmalloc causes hang or crash in frmalloc recursive lock acquisition. This issue is reported on AIX 6.1.

**Workaround:** To resolve the issue, install the following APARs before installing Veritas product:

AIX 6.1 TL4

APAR IZ65498

AIX 6.1 TL4                    APAR IZ64768

For versions earlier than AIX 6.1 TL4, contact IBM for a suitable APAR.

## **ha command does not work when VCS\_DOMAIN or VCS\_DOMAINTYPE is set with remote broker (2272352)**

When VCS\_DOMAIN or VCS\_DOMAINTYPE is set with remote broker, ha command does not work.

### **Workaround:**

- 1 Set VCS\_REMOTE\_BROKER to the remote AB:

```
# export VCS_REMOTE_BROKER=remote_broker
```

- 2 Set VCS\_DOMAIN and VCS\_DOMAINTYPE:

```
# export VCS_DOMAINTYPE=ldap  
# export VCS_DOMAIN=ldap_domain_name
```

- 3 Run halogin:

```
# halogin ldap_user
```

Provide password when prompted.

- 4 Unset VCS\_DOMAIN and VCS\_DOMAINTYPE:

```
# unset VCS_DOMAINTYPE  
# unset VCS_DOMAIN
```

- 5 Run any ha command. The command should run fine if the *ldap\_user* has the correct privileges

## **NFS cluster I/O fails when storage is disabled**

The I/O from the NFS clusters are saved on a shared disk or a shared storage. When the shared disks or shared storage connected to the NFS clusters are disabled, the I/O from the NFS Client fails and an I/O error occurs.

Workaround: If the application exits (fails/stops), restart the application.

## **Issues related to installation**

This section describes the known issues during installation and upgrade.

### While configuring authentication passwords through the Veritas product installer, the double quote character is not accepted (1245237)

The Veritas product installer prompts you to configure authentication passwords when you configure Veritas Cluster Server (VCS) as a secure cluster, or when you configure Symantec Product Authentication Service (AT) in authentication broker (AB) mode. If you use the Veritas product installer to configure authentication passwords, the double quote character (") is not accepted. Even though this special character is accepted by authentication, the installer does not correctly pass the characters through to the nodes.

**Workaround:** There is no workaround for this issue. When entering authentication passwords, do not use the double quote character (").

### Errors observed during partial upgrade of SFHA

While upgrading the VCS packages during an SFHA upgrade from 5.0 MP3 RP2 to 5.1SP1, CPI failed to uninstall the I/O fencing packages (VRTSvxfen, VRTSllt, and VRTSgab). [1779129]

#### Workaround

Before upgrading SFHA from 5.0 MP3 RP2 to 5.1SP1, you must apply the I/O fencing hotfix 5.0MP3RP2HF2.

### Issues with keyless licensing reminders after upgrading VRTSvlic

After upgrading from 5.1 to 5.1SP1, some keyless licenses may be left in the system. As a result, you may see periodic reminders being logged if the VOM server is not configured.

This happens if you were using keyless keys before upgrading to 5.1SP1. After the upgrade, you install real keys and run `vxkeyless set NONE`. In this case, the keyless licenses may not be completely removed and you see warning messages being logged after two months (if VOM server is not configured). This does not result in any functionality impact.

To resolve this issue, perform the following steps:

1. Note down the list of products configured on the node for keyless licensing. Run `vxkeyless display` to display the list.
  2. Set the product level to *NONE* with the command:  

```
# vxkeyless set NONE
```
  3. Find and delete the keyless licenses left over in the system. To do this, perform the following steps for every key stored in `/etc/vx/licenses/lic:`
- Verify if the key has `VXKEYLESS` feature Enabled using the following command:

```
# vxlicrep -k <license_key> | grep VXKEYLESS
```

- Delete the key if and only if `VXKEYLESS` feature is Enabled.

---

**Note:** When performing the search, do not include the `.vxlic` extension as part of the search string.

---

4. Restore the previous list of products with the command:

```
# vxkeyless set product1[,product]
```

### **Installer is unable to split a cluster that is registered with one or more CP servers**

Splitting a cluster that uses server-based fencing is currently not supported. [2110148]

You can split a cluster into two and reconfigure on the two clusters using the installer. For example, you can split a cluster `clus1` into `clus1A` and `clus1B`.

However, if you use the installer to reconfigure the , the installer retains the same cluster UUID of `clus1` in both `clus1A` and `clus1B`. If both `clus1A` and `clus1B` use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: None.

### **Issues related to the VCS engine**

#### **Systems with multiple CPUs and copious memory shut-down time may exceed the ShutdownTimeout attribute**

The time taken by the system to go down may exceed the default value of the `ShutdownTimeout` attribute for systems that have a large numbers of CPUs and memory. [1472734 ]

Workaround: Increase the value of the `ShutdownTimeout` attribute based on your configuration.

#### **The `hacf -cmdtocf` command generates a broken `main.cf` file**

The `hacf -cmdtocf` command used with the `-dest` option removes the include statements from the types files. [1728738]

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

## VCS Engine logs messages when it eventually connects to a remote cluster

Description: In a global cluster, if a local cluster fails to connect with a remote cluster in the first attempt but succeeds eventually, then you may see the following warning messages in the engine logs. [2110108]

```
VCS WARNING V-16-1-10510 IpmHandle: pen Bind Failed.  
unable to bind to source address 10.209.125.125. errno = 67
```

Workaround: There is currently no workaround for this issue. This issue has no impact on any functionality.

## Agent framework can reject `hares -action` command

When a probed resource is disabled and later enabled then, the agent framework can reject `hares -action` command till the agent successfully monitors the resource.

## New nodes get added to SystemList and AutoStartList attributes of ClusterService even if AutoAddSystemToCSG is disabled

The AutoAddSystemToCSG attribute determines whether the newly joined or added systems in a cluster become part of the SystemList of the ClusterService service group if the service group is configured. The value 1 (default) indicates that the new systems are added to SystemList of ClusterService.

AutoAddSystemToCSG has an impact only when you execute the `hasys -add` command or when a new node joins the cluster. [2159139]

However, when you use the installer to add a new node to the cluster, the installer modifies the SystemList and AutoStartList attributes irrespective of whether AutoAddSystemToCSG is enabled or disabled. The installer adds the new system to the SystemList and AutoStartList. To add nodes, the installer uses the following commands that are not affected by the value of AutoAddSystemToCSG:

```
# hagrpl -modify ClusterService SystemList -add newnode n  
# hagrpl -modify ClusterService AutoStartList -add newnode
```

## Workaround

The installer will be modified in future to prevent automatic addition of nodes to SystemList and AutoStartList.

As a workaround, use the following commands to remove the nodes from the SystemList and AutoStartList:

```
# hagrpl -modify ClusterService SystemList -delete newnode  
# hagrpl -modify ClusterService AutoStartList -delete newnode
```

## Issues related to the agent framework

### Agent framework cannot handle leading and trailing spaces for the dependent attribute

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

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. [1511211]

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.

### The ArgListValues attribute values for dependent resources may not populate correctly when a target resource is deleted and re-added

For resource attributes, deleting a resource prevents a dependent attribute's value from refreshing in the dependent resource's value.

For example, you have resource (*rD*), which depends on a resource's attribute value (*rT:Attr\_rt*). When you delete the target resource (*rT*), and re-add it (*rT*), the dependent resource (*rD*) does not get the correct value for the attribute (*Attr\_rt*). [1539927]

Workaround: Set the value of the reference attribute (*target\_res\_name*) to an empty string.

```
# hares -modify rD target_res_name ""
```

Where *rD* is the name of the dependent resource, and *target\_res\_name* is the name of the reference attribute that contains the name of the target resource.

Set the value of the reference attribute (*target\_res\_name*) to the name of the target resource (*rT*).

```
# hares -modify rD target_res_name rT
```

### Agent performance and heartbeat issues

Depending on the system capacity and the number of resources configured under VCS, the agent may not get enough CPU cycles to function properly. This can

prevent the agent from producing a heartbeat synchronously with the engine. If you notice poor agent performance and an agent's inability to heartbeat to the engine, check for the following symptoms.

Navigate to `/var/VRTSvcs/diag/agents/` and look for files that resemble:

```
FFDC_AGFWMMain_729_agent_type.log  FFDC_AGFWTimer_729_agent_type.log  core
FFDC_AGFWSvc_729_agent_type.log   agent_typeAgent_stack_729.txt
```

Where *agent\_type* is the type of agent, for example Application or FileOnOff. If you find these files, perform the next step.

Navigate to `/var/VRTSvcs/log/` and check the `engine_*.log` file for messages that resemble:

```
2009/10/06 15:31:58 VCS WARNING V-16-1-10023 Agent agent_type
not sending alive messages since Tue Oct 06 15:29:27 2009
2009/10/06 15:31:58 VCS NOTICE V-16-1-53026 Agent agent_type
ipm connection still valid
2009/10/06 15:31:58 VCS NOTICE V-16-1-53030 Termination request sent to
agent_type agent process with pid 729
```

**Workaround:** If you see that both of the above criteria are true, increase the value of the `AgentReplyTimeout` attribute value. (Up to 300 seconds or as necessary.) [1853285]

## Issues related to global clusters

### Clusters are stuck in INIT and LOST\_CONN states after enabling AT after cluster migration to IPv6

Clusters are stuck in INIT and LOST\_CONN states after enabling Symantec Product Authentication Service (AT) for the first time without secure WAC. [1836428/1847556]

**Workaround:** Reboot the clusters or restart VCS on both the clusters. The command `hastop -all` may not stop WAC. In this case, you have to manually kill WAC.

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

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. [1539646]

**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**

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. [2107386]

**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 LLT**

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

### **LLT port stats sometimes shows recvcnt larger than recvbytes**

With each received packet, LLT increments the following variables:

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

Both these variables are integers. With constant traffic, recvbytes hits and rolls over `MAX_INT` quickly. This can cause the value of recvbytes to be less than the value of recvcnt. [1788315]

This does not impact the LLT functionality.

### **LLT may incorrectly declare port-level connection for nodes in large cluster configurations**

When ports get registered and unregistered frequently on the nodes of the cluster, LLT may declare that a port-level connection exists with another peer node. This occurs in some corner cases even though a port is not even registered on the peer node. [1809827]

## **Issues related to I/O fencing**

This section covers the known issues related to I/O fencing in this release.

### **All nodes in a sub-cluster panic if the node that races for I/O fencing panics**

At the time of a network partition the lowest node in each sub-cluster races for the coordination points on behalf of that sub-cluster. If the lowest node is unable to contact a majority of the coordination points or the lowest node itself unexpectedly panics during the race, then all the nodes in that sub-cluster will panic. [1965954]

### **Coordination Point agent does not provide detailed log message for inaccessible CP servers**

The Coordination Point agent does not log detailed information of the CP servers that are inaccessible. When CP server is not accessible, the agent does not mention the UUID or the virtual IP of the CP server in the engine log. [1907648]

### **Preferred fencing does not work as expected for large clusters in certain cases**

If you have configured system-based or group-based preferred fencing policy, preferred fencing does not work if all the following cases are true:

- The fencing setup uses customized mode with one or more CP servers.
- The application cluster has more than eight nodes.
- The node weight for a single node (say galaxy with node id 0) is more than the sum total of node weights for the rest of the nodes.
- A network fault occurs and the cluster partitions into two with the single node (galaxy) on one part and the rest of the nodes on the other part.

Under such circumstances, for group-based preferred fencing, the single node panics even though more high priority services are online on that node. For system-based preferred fencing, the single node panics even though more weight is assigned to the node. [2161816]

See the *Veritas product Administrator's Guide* for more information on preferred fencing.

### **Server-based I/O fencing fails to start after configuration on nodes with different locale settings**

On each (application cluster) node, the vxfen module retrieves and stores the list of the UUIDs of coordination points. When different nodes have different locale settings, the list of UUIDs on one (application) node does not match with that of the other (application) nodes. Hence, I/O fencing does not start after configuration. [2112742]

Workaround: Start I/O fencing after fixing the locale settings to use the same values on all the (application) cluster nodes.

### **Reconfiguring with I/O fencing fails if you use the same CP servers**

When you reconfigure an application cluster that uses server-based I/O fencing (customized fencing mode), the installer does not remove the application cluster information from the CP servers before the reconfiguration. As a result, if you reconfigure the application cluster and choose to configure I/O fencing in customized mode using the same CP servers, then reconfiguration of server-based fencing for the application cluster fails. [2076240]

**Workaround:** Manually remove the application cluster information from the CP servers after you reconfigure but before you reconfigure server-based I/O fencing for the application cluster.

See the *Veritas Cluster Server Administrator's Guide* for instructions to remove the application cluster information from the CP servers.

### **CP server cannot bind to multiple IPs (2085941)**

Coordination point server (CP server) binds only to a single virtual IP and listens on the same. Application clusters cannot access the CP server if it fails to establish connection to this virtual IP. Therefore, if the connection fails because of the subnet in which the virtual IP of the CP server exists, you cannot access the CP server even if there is another subnet through which the client can connect to the CP server over a different IP.

**Resolution:** No known resolution for this issue.

### **I/O fencing driver fails to configure after reboot in a non-SCSI3 fencing environment**

With non-SCSI3 fencing configured, if you reboot cluster nodes using the `reboot -n` command successively without any delay, then the VXFEN driver fails to configure after reboot. [2074279]

**Workaround:** After rebooting on one system, delay the reboot on the next system by 180 seconds.

## **Issues related to Symantec Product Authentication Service with VCS**

This section covers the known issues related to Symantec Product Authentication Service (AT) in this release.

### **The `vcsat` and `cpsat` commands may appear to be hung**

The following commands may appear to be hung when you invoke them from the command shell:

- `/opt/VRTScps/bin/cpsat`
- `/opt/VRTSvc/bin/vcsat`

This issue occurs when the command requires some user interaction. [1841185]

**Workaround:**

- To fix the issue for `vcsat`, run the commands as follows:

```
# export EAT_HOME_DIR=/opt/VRTSvcS  
# /opt/VRTSvcS/bin/vssatvcs command_line_argument  
# unset EAT_HOME_DIR
```

- To fix the issue for cpsat, run the commands as follows:

```
# export EAT_HOME_DIR=/opt/VRTScps  
# /opt/VRTScps/bin/vssatcps command_line_argument  
# unset EAT_HOME_DIR
```

## Veritas Cluster Server agents for Veritas Volume Replicator known issues in

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

### Issues with bunker replay

When ClusterFailoverPolicy is set to Auto and the AppGroup is configured only on some nodes of the primary cluster, global cluster immediately detects any system fault at the primary site and quickly fails over the AppGroup to the remote site. VVR might take longer to detect the fault at the primary site and to complete its configuration changes to reflect the fault.

This causes the RVGPrimary online at the failover site to fail and the following message is displayed:

```
RVGPrimary:RVGPrimary:online:Diskgroup bunkerdgname  
could not be imported on bunker host hostname. Operation  
failed with error 256 and message VxVM  
VVR vradmin ERROR V-5-52-901 NETWORK ERROR: Remote server  
unreachable...
```

```
Timestamp VCS ERROR V-16-2-13066 (hostname) Agent is calling  
clean for resource(RVGPrimary) because the resource  
is not up even after online completed.
```

**Resolution:** To ensure that global clustering successfully initiates a bunker replay, Symantec recommends that you set the value of the OnlineRetryLimit attribute to a non-zero value for RVGPrimary resource when the primary site has a bunker configured.

## Veritas Storage Foundation for Databases (SFDB) tools known issues

The following are known issues in this release of Veritas Storage Foundation products.

### Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0.x to 5.1SP1

The `sfua_rept_migrate` command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 or 5.0MP3 to SFHA or SF for Oracle RAC 5.1SP1. The error message is:

```
/sbin/rc3.d/S*vxdbms3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

When using SFHA or SF for Oracle RAC version 5.0 or 5.0MP3 to SFHA or SF for Oracle RAC 5.1SP1 the `S*vxdbms3` startup script is renamed to `NO_S*vxdbms3`. The `S*vxdbms3` startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the `S*vxdbms3` startup script and gives the above error message.

Workaround:

Before running `sfua_rept_migrate`, rename the startup script `NO_S*vxdbms3` to `S*vxdbms3`.

### Database fails over during Flashsnap operations (1469310)

In an Veritas product environment, if the database fails over during Flashsnap operations such as the `dbed_vmsnap -o resync` command and various error messages appear. This issue occurs because Flashsnap commands do not create a VCS resource for the SNAP disk group. As such, when the database fails over, only the primary disk group is moved to another node.

Workaround:

There is no workaround for this issue.

The error messages depend on the timing of the database failover. To fix the problem, you need to bring the FlashSnap state to `SNAP_READY`. Depending on the failure, you may have to use base VxVM commands to reattach mirrors. After mirrors are attached, you need to wait until the mirrors are in `SNAPDONE` state. Re-validate the snapplan again.

## Reattach command failure in a multiple disk group environment (1840672)

In a multiple disk group environment, if the snapshot operation fails then `dbed_vmsnap` fails to reattach all the volumes. This operation must be performed as root user.

### Workaround

In case the reattach operation fails, use the following steps to reattach the volumes.

To reattach volumes in a multiple disk group environment if the snapshot operation fails

- 1 Join the snapshot disk groups to primary diskgroups. The snapshot disk group name is a concatenation of “SNAPSHOT\_DG\_PREFIX” parameter value in `snapplan` and primary disk group name. Use the following command to join the disk groups:

```
# vxdg join snapshot_disk_group_name  
           primary_disk_group_name
```

- 2 Start all the volumes in primary disk group.

```
# vxvol -g primary_disk_group_name startall
```

- 3 Reattach the snapshot volumes with primary volumes. The snapshot volume names is a concatenation of “SNAPSHOT\_VOL\_PREFIX” parameter value in `snapplan` and primary volume name. Use the following command to reattach the volumes.

```
# vxsnap -g primary_disk_group_name reattach snapshot_volume_name  
source=primary_volume_name
```

Repeat this step for all the volumes.

## Clone command fails if archive entry is spread on multiple lines (1764885)

If you have a `log_archive_dest_1` in single line in the `init.ora` file, then `dbed_vmclonedb` will work but `dbed_vmcloneb` will fail if you put in multiple lines for `log_archive_dest_1`

Workaround:

There is no workaround for this issue.

## VCS agent for Oracle: Health check monitoring is not supported for Oracle database 11g R1 and 11g R2 (1985055)

Health check monitoring is not supported for Oracle database 11g R1 and 11g R2.

Workaround: Set MonitorOption attribute for Oracle resource to 0.

## Veritas Storage Foundation for Oracle RAC known issues

The following are new additional Veritas Storage Foundation for Oracle RAC known issues in this 5.1 SP1 RP1 release.

### Incorrect ownership assigned to the parent directory of ORACLE\_BASE causes Oracle Clusterware/Grid Infrastructure installations to fail

When you use the Veritas product installation program to install Oracle Clusterware/Grid Infrastructure, the ownership of the parent directory of ORACLE\_BASE/GRID\_BASE that is created by the installation program is incorrectly set to root. This causes the Oracle Universal Installer to display errors when it creates the `oraInventory` directory as the `oraInventory` directory must be created on behalf of the `oracle` user (Oracle RAC 10g Release 2/Oracle RAC 11g Release 1) or `grid` user (Oracle RAC 11g Release 2).

#### Workaround:

1. Log into each node in the cluster as the root user.
2. Perform the following operations:
  - If you have not yet installed Oracle Clusterware/Grid Infrastructure, create the directory and set the correct ownership as follows before you invoke the installation program:

```
# mkdir -p oracle_base
# chown user_name:oraInventory_group_name
  oracle_base/..
```

where:

*oracle\_base* is the name of the Oracle base directory.

*user\_name* is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).

*oraInventory\_group\_name* is the name of the `oraInventory` group.

Complete the other preparatory tasks before you invoke the installation program. For instructions, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

- If you faced this issue during the installation of Oracle Clusterware/Grid Infrastructure, open another terminal session, and modify the ownership of the directory on all nodes in the cluster as follows:

```
# chown user_name:oraInventory_group_name  
      oracle_base/..
```

where:

*oracle\_base* is the name of the Oracle base directory.

*user\_name* is the name of the user (For Oracle Clusterware: `oracle`; For Oracle Grid Infrastructure: `grid`).

*oraInventory\_group\_name* is the name of the oraInventory group.

Return to the former session and proceed with the installation.

## Adding a node to a cluster fails with errors

When you add a node to an existing Veritas product cluster using the product installer, the shared volumes fail to mount on the new node.

The following message is displayed:

```
Mount vol on /testmnt for new_node ..... Failed
```

This causes CVM to fault on the new node and the new node fails to join the cluster. [2242561]

**Workaround:** Perform the steps in the following procedure:

1. Log into the new node as the root user.
2. Stop all `vxfsckd` processes.

Obtain the process IDs (PID) of all `vxfsckd` processes:

```
# ps -ef | grep vxfsckd
```

Kill the process IDs:

```
# kill -9 PID
```

3. Unregister VxFS from GAB port membership (f):

```
# fsclustadm cfsdeinit
```

4. Stop VCS on the new node:

```
# hastop -local
```

5. Start VCS on the new node:

```
# hastart
```

## Software limitations

This section covers the software limitations of this release.

### Veritas Storage Foundation software limitations

The following are software limitations in the 5.1 SP1 RP1 release of Veritas Storage Foundation.

#### Upgrades on alternate disk supported only from version 5.1

Veritas product supports upgrade on an alternate disk only from version 5.1 to version 5.1 SP1 RP1. If you are running earlier versions of Veritas product, perform a full or phased upgrade to version 5.1 and then upgrade to version 5.1 SP1 RP1 using an alternate disk.

### Veritas Volume Manager software limitations

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

#### Limitation with device renaming on AIX 6.1TL6

If you rename an operating system (OS) path with the `rendev` command on AIX 6.1TL6, the operation might remove the paths from DMP control. DMP cannot discover these paths.

#### DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment, set the following DMP tunables:

Table 1-8

| Parameter name                    | Definition               | New value    | Default value |
|-----------------------------------|--------------------------|--------------|---------------|
| <code>dmp_restore_internal</code> | DMP restore daemon cycle | 60 seconds.  | 300 seconds.  |
| <code>dmp_path_age</code>         | DMP path aging tunable   | 120 seconds. | 300 seconds.  |

The change is persistent across reboots.

### To change the tunable parameters

- 1 To change the tunable parameters, run the following commands:

```
# vxddmpadm settune dmp_restore_internal=60
# vxddmpadm settune dmp_path_age=120
```

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

```
# vxddmpadm gettune dmp_restore_internal
# vxddmpadm gettune dmp_path_age
```

## DMP support in AIX virtualization environment (2038475)

A single enclosure cannot have both NPIV and vSCSI LUNs. Each enclosure can have either vSCSI or NPIV LUNs. DMP does not support a mixed configuration.

## Veritas Volume Replicator software limitations

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

### Replication in a shared environment

Currently, replication support is limited to 4-node cluster applications.

### IPv6 software limitations

VVR does not support the following Internet Protocol configurations:

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

- VVR does not support Edge and NAT-PT routers that facilitate IPv4 and IPv6 address translation.

## VVR support for replicating across Storage Foundation versions

VVR supports replication between Storage Foundation 5.1SP1 RP1 and the prior major releases of Storage Foundation (5.0 MP3 and 5.1). Replication between versions is supported for disk group versions 140, 150, and 160 only. Both the Primary and Secondary hosts must be using a supported disk group version.

## Cannot relayout data volumes in an RVG from concat to striped-mirror (2162537)

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

### Workaround:

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

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

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

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

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

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

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

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

## Veritas Cluster Server software limitations

The following are software limitations in this release of Veritas Cluster Server.

### Limitations related to installing and upgrading VCS

#### Upgrades on alternate disk supported only from version 5.1

Veritas product supports upgrade on an alternate disk only from version 5.1 to version 5.1 SP1. If you are running earlier versions of Veritas product, perform a full or phased upgrade to version 5.1 and then upgrade to version 5.1 SP1 using an alternate disk.

### Limitations related to the VCS engine

#### VCS deletes user-defined VCS objects that use the HostMonitor object names

If you had defined the following objects in the main.cf file using the reserved words for the HostMonitor daemon, then VCS deletes these objects when the VCS engine starts. [1293092]

- Any group that you defined as VCSHmg along with all its resources.
- Any resource type that you defined as HostMonitor along with all the resources of such resource type.
- Any resource that you defined as VCSHm.

### Limitations related to the bundled agents

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

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

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

Symantec recommends creating users locally. To reflect local users, configure: `/etc/netsvc.conf`

### Limitations of the DiskGroup agent

Volumes in disk group are started automatically if the Veritas Volume Manager default value of `AutoStartVolumes` at system level will be set to `ON` irrespective of the value of the `StartVolumes` attribute defined inside the VCS. Set `AutoStartVolumes` to `OFF` at system level if you do not want to start the volumes as part of import disk group.

### Volume agent clean may forcibly stop volume resources

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

### False concurrency violation when using PidFiles to monitor application resources

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

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

### Volumes in a disk group start automatically irrespective of the value of the StartVolumes attribute in VCS

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

#### 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.

### Limitations related to IMF

- IMF registration on Linux for “bind” file system type is not supported.
- In case of SLES11 SP1:
  - IMF should not be enabled for the resources where the `BlockDevice` can get mounted on multiple `MountPoints`.

- If FSType attribute value is nfs, then IMF registration for “nfs” file system type is not supported.

## Limitations related to the VCS database agents

### DB2 RestartLimit value

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. [1231311]

### Limitation with intentional offline functionality of VCS agent for Oracle

The Oracle resource never faults after an intentional offline.

Intentional offline functionality of VCS agent for Oracle requires you to enable health check monitoring. The agent uses Oracle's Health Check API to find the state of the database. If the API returns a graceful shutdown for the database, then the agent marks the resource state as INTENTIONAL OFFLINE. Later if the Oracle agent's online function does not succeed, the agent does not mark the resource as FAULTED. The state remains as INTENTIONAL OFFLINE because the agent receives the database state from the API as graceful shutdown during each monitor cycle. [1805719]

## 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.

## Systems in a cluster must have same system locale setting

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

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

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

If you attempt to shut down the node where the fire drill service group or the disk group is online, the node goes to a LEAVING state. The VCS engine attempts to take all the service groups offline on that node and rejects all action entry point requests. Therefore, the DiskGroupSnap agent cannot invoke the action to reattach the fire drill site to the target disk group. The agent logs a message that the node is in a leaving state and then removes the lock file. The agent's monitor function declares that the resource is offline. After the node restarts, the disk group site still remains detached. [1272012]

Workaround:

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

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

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

## Limitations with DiskGroupSnap agent

The DiskGroupSnap agent has the following limitations:

- The DiskGroupSnap agent does not support layered volumes. [1368385]
- 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: [1391445]
  - 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.

## Virtualizing shared storage using VIO servers and client partitions

AIX 5.3, with proper patches to the operating system and client partitions, is capable of running multiple virtualized partitions within a single frame. You can split the CPU, memory, and certain adapters (networking and storage), into smaller virtual units that the partitions can then use.

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

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

The minimum patch level for using VIO servers with VCS is: Fix Pack 7.1.2.0.0.

### Supported storage

Refer to the IBM data sheet:

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

### Disk Restrictions

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

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

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

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

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

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

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

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

Get the serial number of the disk of interest:

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

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

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

List the virtual SCSI adapters.

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

---

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

---

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

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

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

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

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

Next create a virtual device to map `hdisk20` to `vhost1`.

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

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

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

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

## Cluster Manager (Java console) limitations

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

### Use the VCS 5.1 Java Console to manage clusters

Cluster Manager (Java Console) from versions lower than VCS 5.1 cannot be used to manage VCS 5.1SP1 clusters. Symantec recommends using the latest version of Cluster Manager.

See the *Veritas Cluster Server Installation Guide* for instructions on upgrading Cluster Manager.

### Run Java Console on a non-cluster system

Symantec recommends not running Cluster Manager (Java Console) for an extended period on a node in the cluster. The Solaris version of the Java Virtual Machine has a memory leak that can gradually consume the host system's swap space. This leak does not occur on Windows systems.

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

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

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

### VCS Simulator does not support I/O fencing

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

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

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

### Limitations related to LLT

This section covers LLT-related software limitations.

#### LLT does not start automatically after system reboot

After you reboot the systems, if you had not completed the terminal setting procedure, LLT does not start automatically and does not log any error messages. You can manually start LLT using the `/etc/init.d/llt.rc` command. [2058752]

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

#### Workaround: To resolve the LLT startup issue

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

### Limitations related to I/O fencing

This section covers I/O fencing-related software limitations.

### **Limitation with RDAC driver and FASTT array for coordinator disks that use raw disks**

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

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

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

### **Stopping systems in clusters with I/O fencing configured**

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect, or “split brain.” See the *Veritas Cluster Server Administrator's Guide* for a description of the problems a failed interconnect can create and the protection I/O fencing provides.

I/O fencing uses SCSI-3 PR keys to implement data protection. Keys are placed on I/O fencing coordinator points and on data disks. The VCS administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordinator 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 coordinator 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.

### **Cannot modify VxFEN tunable parameters**

You cannot change the VxFEN tunable parameters due to a software limitation. [1863916]

## **Veritas Storage Foundation for Databases tools software limitations**

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

## Oracle Data Guard in an Oracle RAC environment

Database snapshots and Database Checkpoints are not supported in a Data Guard and Oracle RAC environment.

### Upgrading if using Oracle 11.1.0.6

If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 5.1SP1 RP1: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to SP1.

## List of patches

This section lists the patches for 5.1 SP1 RP1.

---

**Note:** You can also view the following list using the `installrp` command, type:  
`./installrp -listpatches`

---

**Table 1-9** Patches for AIX

| BFF file       | Size in bytes | Patches    | Version     |
|----------------|---------------|------------|-------------|
| VRTSamf.bff    | 3737600       | VRTSamf    | 5.1.101.0   |
| VRTScavf.bff   | 256000        | VRTScavf   | 5.1.101.0   |
| VRTScps.bff    | 48076800      | VRTScps    | 5.1.101.0   |
| VRTSglm.bff    | 768000        | VRTSglm    | 5.1.101.0   |
| VRTSllt.bff    | 2867200       | VRTSllt    | 5.1.101.0   |
| VRTSodm.bff    | 716800        | VRTSodm    | 5.1.101.0   |
| VRTSsfmh.bff   | 39270400      | VRTSsfmh   | 3.1.429.401 |
| VRTSvc.bff     | 318976000     | VRTSvc     | 5.1.101.0   |
| VRTSvcsg.bff   | 19353600      | VRTSvcsg   | 5.1.101.0   |
| VRTSvcsea.bff  | 6297600       | VRTSvcsea  | 5.1.101.0   |
| VRTSvcxfen.bff | 3430400       | VRTSvcxfen | 5.1.101.0   |
| VRTSvcxfs.bff  | 26931200      | VRTSvcxfs  | 5.1.101.0   |
| VRTSvcxvm.bff  | 80793600      | VRTSvcxvm  | 5.1.101.0   |

## Downloading the 5.1 SP1 RP1 archive

The patches that are included in the 5.1 SP1 RP1 release are available for download from the Symantec website. After downloading the 5.1 SP1 RP1 rolling patch, use `gunzip` and `tar` commands to uncompress and extract it.

For the 5.1 SP1 RP1 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/334998>

## About the `installrp` script

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1 provides a new upgrade script. To upgrade from Veritas Storage Foundation and High Availability Solutions version 5.1 SP1 or later, Symantec recommends that you use the new upgrade script. The `installrp` script allows you to upgrade all the patches associated with the packages installed, after which you can reboot to start all of the processes.

## The `installrp` script options

**Table 1-10** The command line options for the product upgrade script

| Command Line Option                          | Function   |
|--|--|
| [ <i>system1 system2...</i> ]                | Specifies the systems on which to run the upgrade options. If not specified, the command prompts for a system name.  |
| [ <code>-precheck</code> ]                   | Use the <code>-precheck</code> option to confirm that systems meet the products' installation requirements before the installation.  |
| [ <code>-logpath log_path</code> ]           | Use the <code>-logpath</code> option to select a directory other than <code>/opt/VRTS/install/logs</code> as the location where the <code>installrp</code> log files, summary file, and response file are saved.   |
| [ <code>-responsefile response_file</code> ] | Use the <code>-responsefile</code> option to perform automated installations or uninstallations using information stored in a file rather than prompting for information. <i>response_file</i> is the full path of the file that contains configuration definitions. |

**Table 1-10** The command line options for the product upgrade script (*continued*)

| Command Line Option                | Function  |
|------------------------------------|---|
| [ -makeresponsefile ]              | Use the <code>-makeresponsefile</code> option to generate a response file without doing an actual installation. The text that displays <code>install</code> , <code>uninstall</code> , <code>start</code> , and <code>stop</code> actions are simulations. These actions are not performed on the system. |
| [ -tmppath <i>tmp_path</i> ]       | Use the <code>-tmppath</code> option to select a directory other than <code>/var/tmp</code> as the working directory for <code>installrp</code> . This destination is where initial logging is performed and where filesets are copied on remote systems before installation.                             |
| [ -hostfile <i>hostfile_path</i> ] | Use the <code>-hostfile</code> option to specify the location of a file containing the system names for installer.  |
| [ -keyfile <i>ssh_key_file</i> ]   | Use the <code>-keyfile</code> option to specify a key file for SSH. When you use this option the <code>-i <i>ssh_key_file</i></code> is passed to every SSH invocation.   |
| [-nim]                             | Use to produce a NIM configuration file for installing with NIM.<br><br>Refer to the product's <i>Installation Guide</i> for more information on using NIM.   |
| [ -patchpath <i>patch_path</i> ]   | Use the <code>-patchpath</code> option to define the complete path of a directory that is available to all install systems (usually NFS mounted) that contains all patches to be installed by <code>installrp</code> .  |

**Table 1-10** The command line options for the product upgrade script (*continued*)

| Command Line Option   | Function  |
|---|---|
| <pre>[ -rsh   -redirect   -listpatches   -pkginfo   -serial   -upgrade_kernelpkgs   -upgrade_nonkernelpkgs   -version ]</pre> | <p>Use the <code>-rsh</code> option when <code>rsh</code> and <code>rcp</code> are to be forced for communication though <code>ssh</code> and <code>scp</code> is also setup between the systems.</p> <p>Use the <code>-redirect</code> option to display progress details without showing the progress bar.</p> <p>Use the <code>-listpatches</code> option to display product patches in the correct installation order.</p> <p>Use the <code>-pkginfo</code> option is used to display the correct installation order of packages and patches. This option is available with or without one of following options: <code>-allpkgs</code>, <code>-minpkgs</code>, and <code>-recpkgs</code>.</p> <p>Use the <code>-serial</code> option to perform installation, uninstallation, start and stop operations, typically performed simultaneously on all systems, in a serial fashion.</p> <p>Use the <code>-upgrade_kernelpkgs</code> option for the rolling upgrade's upgrade of kernel packages to the latest version</p> <p>Use the <code>-upgrade_nonkernelpkgs</code> option for the rolling upgrade's upgrade of non-kernel packages. In this phase, VCS packages and other agent packages are upgraded to the latest versions. Product kernel drivers are upgraded to the latest protocol version.</p> <p>Use the <code>-version</code> option to have the installer check and report the installed products and their versions. Identifies the installed and missing and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing and patches where applicable.</p> |



# Installing the products for the first time

This chapter includes the following topics:

- [Installing the Veritas software using the script-based installer](#)
- [Installing Veritas software using the Web-based installer](#)

## Installing the Veritas software using the script-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and install 5.1 SP1 RP1. Review the 5.1 SP1 Installation Guide and Release Notes for your product.

### To install the Veritas software for the first time

- 1 Mount the 5.1 SP1 product disc and navigate to the folder that contains the installation program to install 5.1 SP1. Choose one of the following to start the installation:
  - For Veritas Storage Foundation:

```
# ./installsf node1 node2 ... nodeN
```
  - For Veritas Storage Foundation High Availability:

```
# ./installsfha node1 node2 ... nodeN
```
  - For Veritas Storage Foundation Cluster File System and Veritas Storage Foundation Cluster File System High Availability:

```
# ./installsfcfs node1 node2 ... nodeN
```

- For Veritas Storage Foundation for Oracle RAC:

```
# ./installsfrac node1 node2 ... nodeN
```

- For Veritas Cluster Server:

```
# ./installvcs node1 node2 ... nodeN
```

- For Dynamic Multi-Pathing:

```
# ./installdmp node1 node2 ... nodeN
```

- 2 Review the installation prerequisites for upgrading to 5.1 SP1 RP1.

See “[Prerequisites for upgrading to 5.1 SP1 RP1](#)” on page 81.

- 3 Copy the patch archive downloaded from patch central to temporary location, untar the archive and browse to the directory containing the `installrp` script.

- If the 5.1 SP1 product is installed, run the `installrp` script to install 5.1 SP1 RP1.

```
# ./installrp node1 node2 ... nodeN
```

If you did not configure the product after the 5.1 SP1 installation, the installer prompts you to configure it. If you do not want to configure the product now, answer **n** to the prompt. To configure the product in the future, run the product installation script from the 5.1 SP1 installation media with the `-configure` option.

## Installing Veritas software using the Web-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and then to install 5.1 SP1 RP1 using the Web-based installer. For detailed instructions on how to install 5.1 SP1 using the Web-based installer, follow the procedures in the 5.1 SP1 Installation Guide and Release Notes for your products.

---

**Note:** Installing SF Oracle RAC using the Web-based installer is not supported in this release.

---

## Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

### To start the Web-based installer

- 1 Start the Veritas XPortal Server process `xprt1wid`, on the installation server:

```
# ./webinstaller start
```

The webinstaller script displays a URL.

- 2 Start the Web browser on the system from which you want to perform the installation.
- 3 Navigate to the URL displayed from step 1.
- 4 The browser may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

- 5 When prompted, enter `root` and root's password of the installation server.

## Obtaining a security exception on Mozilla Firefox

You may need to get a security exception on Mozilla Firefox.

### To obtain a security exception

- 1 Click **Or you can add an exception** link.
- 2 Click **Add Exception** button.
- 3 Click **Get Certificate** button.
- 4 Uncheck **Permanently Store this exception checkbox (recommended)**.
- 5 Click **Confirm Security Exception** button.
- 6 Enter `root` in User Name field and root password of the web server in the Password field.

## Installing products with the Veritas Web-based installer

This section describes installing Veritas product with the Veritas Web-based installer.

### To install Veritas product

- 1 The 5.1SP1 version of the Veritas product must be installed before upgrading to 5.1 SP1 RP1.

- 2 On the **Select a task and product** page, select **Install SP1 RP1** from the **Task** drop-down list, and click **Next**
- 3 Stop all applications accessing the filesystem. Unmount all mounted file systems before installation.
- 4 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.
- 5 After the validation completes successfully, click **Next** to install 5.1 SP1 RP1 patches on the selected system.
- 6 Select the checkbox to specify whether you want to send your installation information to Symantec.

Would you like to send the information about this installation to Symantec to help improve installation in the future?

Click **Finish**.

## Upgrading Veritas product with the Veritas Web-based installer

This section describes upgrading Veritas product with the Veritas Web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

### To upgrade Veritas product

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.  
See [“Starting the Veritas Web-based installer”](#) on page 79.
- 3 Stop all applications accessing the file system. Unmount all mounted filesystems before installation.
- 4 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Validate**.
- 5 Click **Next** to complete the upgrade.

After the upgrade completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

# Upgrading to 5.1 SP1 RP1

This chapter includes the following topics:

- [Prerequisites for upgrading to 5.1 SP1 RP1](#)
- [Supported upgrade paths](#)
- [Upgrading from 5.1 SP1 to 5.1 SP1 RP1](#)
- [Verifying software versions](#)

## Prerequisites for upgrading to 5.1 SP1 RP1

The following list describes prerequisites for upgrading to the 5.1 SP1 RP1 release:

- For any product in the Veritas Storage Foundation stack, you must have the 5.1 SP1 release installed before you can upgrade that product to the 5.1 SP1 RP1 release.
- Each system must have sufficient free space to accommodate patches.
- Stop all applications accessing the vxfs filesystem. Unmount all mounted vxfs file systems before upgrading.
- The full list of prerequisites can be obtained by running `./installrp -precheck`

## Supported upgrade paths

You can upgrade to this release of Veritas product from version 5.1 SP1 or later.

## Upgrading from 5.1 SP1 to 5.1 SP1 RP1

This section describes how to upgrade from 5.1 SP1 to 5.1 SP1 RP1 on a cluster or a standalone system.

- [Performing a full upgrade to 5.1 SP1 RP1 on a cluster](#)  
Use the procedures to perform a full upgrade to 5.1 SP1 RP1 on a cluster that has Veritas Cluster Server (VCS), Veritas Storage Foundation and High Availability Solutions (SFHA), Veritas Storage Foundation Cluster File System (SFCFS), or Veritas Storage Foundation for Oracle RAC (SFRAC) installed and configured.
- [Upgrading to 5.1 SP1 RP1 on a standalone system](#)  
Use the procedure to upgrade to 5.1 SP1 RP1 on a system that has SF installed.
- [Performing a rolling upgrade using the script-based installer](#)  
Use the procedure to upgrade your Veritas product with a rolling upgrade.

### Performing a full upgrade to 5.1 SP1 RP1 on a cluster

Performing a full upgrade on a cluster requires stopping cluster failover functionality during the entire procedure. However, if you use Veritas Storage Foundation Cluster File System (SFCFS) and Cluster Volume Manager (CVM), the SFCFS and CVM services remain available.

The following are the stages of performing a full upgrade on a cluster:

- Freeze service group operations and stop Veritas Cluster Server (VCS) on the cluster.
- Take the nodes offline and install the software patches.
- Bring the nodes online to restart cluster failover services.

Depending on your cluster's configuration, select one of the following procedures to upgrade to 5.1 SP1 RP1:

- [Performing a full upgrade to 5.1 SP1 RP1 for Veritas Cluster Server](#)
- [Performing a full upgrade to 5.1 SP1 RP1 on an SFHA cluster](#)
- [Performing a full upgrade to 5.1 SP1 RP1 on an SFCFS cluster](#)
- [Performing a full upgrade to 5.1 SP1 RP1 on an SF Oracle RAC cluster](#)

### Performing a full upgrade to 5.1 SP1 RP1 for Veritas Cluster Server

The following procedure describes performing a full upgrade on a Veritas Cluster Server (VCS) cluster.

---

**Note:** You need to make sure that IPv4RouteOptions attribute is configured for MultiNICA resources, otherwise network connection may be interrupted.

---

### To upgrade VCS

- 1 Log in as superuser.
- 2 Upgrade the Operating System and reboot the systems if required.  
See [“System requirements”](#) on page 15.
- 3 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 SP1 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check:

```
# ./installrp -precheck node1 node2 ... nodeN
```

See [“About the installrp script”](#) on page 73.

- 4 Resolve any issues that the precheck finds.
- 5 Start the upgrade:

```
# ./installrp node1node2 ... nodeN
```

See [“About the installrp script”](#) on page 73.

- 6 After the upgrade, review the log files for any issues.

## Performing a full upgrade to 5.1 SP1 RP1 on an SFHA cluster

The following procedure describes performing a full upgrade on a SFHA and VCS cluster.

### To perform a full upgrade to 5.1 SP1 RP1 on an SFHA cluster

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.
- 3 On each node, enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

- 4 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

- 5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvg stop` command to stop each RVG individually:

```
# vxrvg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

---

**Caution:** To avoid data corruption, do not proceed until all RLINKs are up-to-date.

---

- 6 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

- 7 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

Verify that no volumes remain open:

```
# vxprint -Aht -e v_open
```

- 8 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```

- 9 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 SP1 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check.

```
# ./installrp -precheck [-rsh] node1node2 ... nodeN
```

The program proceeds in a noninteractive mode to examine the systems for licenses, filesets, disk space, system-to-system communications, etc.

- 10 Review the output as the program displays the results of the check and saves the results of the check in a log file.
- 11 Make sure all your nodes meet the prerequisites, and all the issues reported by above pre-check process have been resolved.

- 12 Start the upgrade.

```
# ./installrp [-rsh] node1node2 ... nodeN
```

Review the output.

- 13 Restart all the volumes. Enter the following command for each disk group:

```
# vxvol -g diskgroup startall
```

- 14 If you stopped any RVGs in step 8, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

- 15 Remount all VxFS file systems on all nodes in the selected group:

```
# mount /filesystem
```

- 16 Remount all Storage Checkpoints on all nodes in the selected group:

```
# mount /checkpoint_name
```

## Performing a full upgrade to 5.1 SP1 RP1 on an SFCFS cluster

The following procedure describes performing a full upgrade on an SFCFS cluster.

### To perform a full upgrade to 5.1 SP1 RP1 on an SFCFS cluster

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.

- 3 On each node, enter the following command to check if any Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

If any Storage Checkpoints are mounted, on each node in the cluster unmount all Storage Checkpoints.

```
# umount /checkpoint_name
```

- 4 On each node, enter the following command to check if any VxFS file systems are mounted:

```
# mount | grep vxfs
```

- If any VxFS file systems are present, on each node in the cluster unmount all of the VxFS file systems:

```
# umount /filesystem
```

- 5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the `vxrvg stop` command to stop each RVG individually:

```
# vxrvg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

---

**Caution:** To avoid data corruption, do not proceed until all RLINKs are up-to-date.

---

- 6 Stop activity to all VxVM volumes.

For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

- 7 On each node, stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

Verify that no volumes remain open:

```
# vxprint -Aht -e v_open
```

- 8 If required, apply the OS kernel patches.  
See IBM's documentation for the procedures.
- 9 On each node, check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```

- 10 From the directory that contains the extracted and untarred 5.1 SP1 RP1 rolling patch binaries, change to the directory that contains the `installrp` script.

```
# ./installrp node1 node2
```

where `node1` and `node2` are nodes which are to be upgraded.

- 11 After all the nodes in the cluster are upgraded, the processes restart. If the `installrp` script finds issues, it may require you to reboot the nodes.
- 12 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node.
- 13 Bring the CVM service group online on each node:

```
# hagr -online cvm -sys nodename
```

- 14 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

- 15 If you stopped any RVGs in step 5, restart each RVG:

```
# vxrv -g diskgroup start rvg_name
```

16 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

17 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

## Performing a full upgrade to 5.1 SP1 RP1 on an SF Oracle RAC cluster

The following procedure describes performing a full upgrade on an SF for Oracle RAC cluster.

To upgrade to 5.1 SP1 RP1 on a SF Oracle RAC cluster

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your **PATH** so that you can execute all product commands.
- 3 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw  
# hagrpl -modify oracle_group AutoStart 0  
# haconf -dump -makero
```

4 Stop Oracle database on the cluster:

- If the Oracle RAC instance is managed by VCS:

```
# hagrpl -offline oracle_group -sys galaxy  
# hagrpl -offline oracle_group -sys nebula
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on all nodes in the cluster and shut down the instances:

```
$ srvctl modify database -d db-name -y manual
```

- 5 Stop all applications on the cluster that are not configured under VCS. Use native application commands to stop the application.
- 6 Unmount the VxFS and CFS file systems that are not managed by VCS.
  - Ensure that no processes are running that make use of mounted shared file system or shared volumes. To verify that no processes use the VxFS or CFS mount point, enter the following commands:

```
# mount | grep vxfs
# fuser -cu /mount_point
# umount /mount_point
```

Unmount the VxFS or CFS file system:

```
# umount /mount_point
```

- 7 Stop all VxVM and CVM volumes for each diskgroup that are not managed by VCS on the cluster:

```
# vxvol -g disk_group stopall
```

Verify that no volumes remain open:

```
# vxprint -Aht -e v_open
```

- 8 From the directory that contains the extracted and untarred 5.1 SP1 RP1 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

```
# ./installrp node1 node2 ...
```

- 9 Restart the nodes:

```
# shutdown -r now
```

- 10 Relink the SF Oracle RAC libraries with Oracle.

Refer to *Veritas Storage Foundation for Oracle RAC 5.1 Installation and Configuration Guide* for more information.

- 11 Enter the following command on each node to unfreeze HA service group operations:

```
# haconf -makerw
# hagr -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- 12 Start VCS on each of the nodes:

- For parallel groups:

```
# hagr -online group_name -sys nodename
```

- For failover groups:

```
# hagrpl -online group_name -any
```

- 13 If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on all nodes in the cluster and start the instances:

```
$ srvctl modify database -d db-name -y AUTOMATIC
```

- 14 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 15 Start all applications on the cluster that are not configured under VCS. Use native application commands to start the application.

## Upgrading to 5.1 SP1 RP1 on a standalone system

You can use this procedure to upgrade on a standalone system that runs SF.

### To upgrade to 5.1 SP1 RP1 on a standalone system

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 3 If required, apply the OS kernel patches.  
See IBM's documentation for the procedures.
- 4 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

- 5 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

- 6 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

---

**Caution:** To avoid data corruption, do not proceed until all RLINKs are up-to-date.

---

**7** Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

**8** Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

Verify that no volumes remain open:

```
# vxprint -Aht -e v_open
```

**9** Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

**10** Copy the patch archive downloaded from the patch central to temporary location and untar the archive and browse to the directory containing the installrp installer script. Enter the `installrp` script:

```
# ./installrp nodename
```

**11** If necessary, reinstate any missing mount points in the `/etc/filesystems` file.

**12** Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

**13** If you stopped any RVGs in step 6, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

**14** Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem  
# mount /checkpoint_name
```

**15** Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

## Performing a rolling upgrade using the script-based installer

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

- [About rolling upgrades](#)
- [Prerequisites for a rolling upgrades](#)
- [Performing a rolling upgrade on kernel packages: phase 1](#)
- [Performing a rolling upgrade on non-kernel packages: phase 2](#)

### About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime. Rolling upgrades are not compatible with phased upgrades. Do not perform "mixed" rolling upgrades with phased upgrades.

Rolling upgrades take two discrete phases. In the first, you upgrade the kernel packages with exception of VCS packages and agent packages. In the second, you upgrade the non-kernel packages, which are VCS packages and agents packages.

You can perform a rolling upgrade for the following products:

- Veritas Cluster Server
- Storage Foundation and High Availability
- Storage Foundation Cluster File System
- Storage Foundation Cluster File System and High Availability
- Storage Foundation for Oracle RAC

You can perform a rolling upgrade from 5.1 SP1 or later.

## Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrades.
- Split up your clusters into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you are logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.

**Limitation:** During VCS and agents upgrade, you must bring down the application High Availability (HA) for several minutes. This does not affect the application running on the cluster. You can restore the application's high availability after VCS and the agent packages are upgraded.

## Performing a rolling upgrade on kernel packages: phase 1

Note that in the following instructions a sub-cluster can represent one or more nodes in a full cluster, but is represented by nodeA.

### To perform the rolling upgrade on kernel packages: phase 1

- 1 Stop all applications that access volumes.
- 2 Unmount any file systems on the nodes that you plan to upgrade.  
You only need to unmount locally mounted file systems. The installer unmounts file systems that SFCFS has mounted.
- 3 On the first sub-cluster, start the installer for the rolling upgrade with the `-upgrade_kernelpkgs` option.  

```
./installrp -upgrade_kernelpkgs nodeA
```
- 4 The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.
- 5 The installer loads new kernel modules.
- 6 The installer starts all the relevant processes and brings all the service groups online.
- 7 Before you proceed to phase 2, complete step 2 to step 6 on the second subcluster.

## Performing a rolling upgrade on non-kernel packages: phase 2

In this phase installer installs all non-kernel filesets on all the nodes in cluster and restarts VCS cluster.

### To perform the rolling upgrade on non-kernel packages: phase 2

- 1 Start the installer for the rolling upgrade with the `-upgrade_nonkernelpkgs` option. Specify all the nodes in the cluster:

```
./installrp -upgrade_nonkernelpkgs nodeA nodeB nodeC...
```

- 2 The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1.
- 3 The installer upgrades non-kernel filesets.
- 4 The installer loads the new kernel modules. It then starts all relevant processes and brings all the service groups online.
- 5 Verify the cluster's status:

```
# hastatus -sum
```

## Verifying software versions

To list the Veritas filesets installed on your system, enter the following command:

```
# ls1pp -l VRTS\*
```

# Rolling back Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [About rolling back Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1](#)
- [Prerequisites for rolling back](#)
- [Rolling back using the `uninstallrp` script](#)
- [Rolling back manually](#)

## About rolling back Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP1

This section describes how to roll back either by using the `uninstallrp` script or manually.

### Prerequisites for rolling back

Perform the following prerequisites before you roll back Veritas Storage Foundation and High Availability Solutions.

- See [“Committing files prior to rolling back”](#) on page 96.

## Committing files prior to rolling back

Before you perform a roll back, make sure that you have committed the fileset up to the 5.1 SP1 or 5.1 SP1 P1/P2 patch level.

### To commit files prior to roll back

- 1 Run the `ls1pp -l` command to determine the levels of the current fileset.

```
# ls1pp -l
```

- 2 Commit the filesets.

```
# installp -c fileset_name
```

For example:

```
installp -c VRTSvxfs 5.1.100.0
```

## Rolling back using the `uninstallrp` script

Use the following procedure to roll back from any Veritas product to 5.1 SP1 using the `uninstallrp` script.

### To roll back

- 1 Browse to the directory that contains the `uninstallrp` script.
- 2 Stop all VxVM volumes. For each disk group enter:

```
# vxvol -g diskgroup stopall
```

Verify that no volumes remain open.

```
# vxprint -Aht -e v_open
```

- 3 Unmount all the Storage Checkpoints and the file systems.

```
# umount /checkpoint_name
```

```
# umount /filesystem
```

Verify that you unmounted the the Storage Checkpoints and the file systems.

```
# mount | grep vxfs
```

- 4 Run the `uninstallrp` script to rollback patches, type:

```
# ./uninstallrp
```

- 5 The `uninstallrp` script checks whether the patches are at 5.1 SP1 committed level, and 5.1 SP1 RP1 applied level. If this is not the case, error messages showing the list of packages and commit levels will be shown.
- 6 The `uninstallrp` script removes 5.1 SP1 RP1 patches. After patch rollback completes, modules are loaded and processes are restarted. `uninstallrp` will also report any warning happened during uninstallation.

## Rolling back manually

Use one the following procedures to roll back to 5.1 SP1 manually.

- [Rolling back Storage Foundation or Storage Foundation and High Availability manually](#)
- [Rolling back Storage Foundation Cluster File System manually](#)
- [Rolling back Storage Foundation for Oracle RAC manually](#)
- [Rolling back Veritas Cluster Server manually](#)
- [Rolling back Dynamic Multi-Pathing manually](#)

---

**Note:** You must reboot systems that you roll back manually at the end of the roll back procedure.

---

## Rolling back Storage Foundation or Storage Foundation and High Availability manually

Use the following procedure to roll back to 5.1 SP1 manually.

### To roll back SF or SFHA

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 3 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
# umount /filesystem
```
- 4 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

- 5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:
  - Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
  - Use the `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

---

**Caution:** To avoid data corruption, do not proceed until all RLINKs are up-to-date.

---

- 6 Stop activity to all VxVM volumes.
- 7 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

To verify that no volumes remain open, enter the following command:

```
# vxprint -Aht -e v_open
```

- 8 Stop VCS and its modules manually.

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

- 9 Stop I/O fencing:

```
# /etc/init.d/vxfen.rc stop
```

- 10 Stop GAB:

```
# /etc/init.d/gab.rc stop
```

- 11 Stop LLT:

```
# /etc/init.d/llt.rc stop
```

12 Unmount `/dev/odm`:

```
# umount /dev/odm
```

13 Unload the ODM module:

```
# genkex | grep odm  
# vxkextadm vxodm unload
```

14 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

15 Remove the Storage Foundation or Storage Foundation and High Availability 5.1 SP1 RP1 patches.

- Create a file that contains all the 5.1 SP1 RP1 patches. In this example, it is called `/reject.list`:

```
# /reject.list
```

- Reject each patch from the patch list file, for example:

```
# installp -rBfX /reject.list
```

16 Reboot the systems. On each system, run the following command.

```
# /usr/sbin/shutdown -r
```

## Rolling back Storage Foundation Cluster File System manually

Use the following procedure to roll back to 5.1 SP1 manually.

To roll back SFCFS or SFCFS HA manually

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.

- 3 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

- 4 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

- 5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

---

**Caution:** To avoid data corruption, do not proceed until all RLINKs are up-to-date.

---

- 6 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

- 7 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

To verify that no volumes remain open, enter the following command:

```
# vxprint -Aht -e v_open
```

- 8** Stop VCS along with all the resources. Then, stop the remaining resources manually:

```
# /etc/rc.d/rc2.d/vcs stop
# /etc/rc.d/rc2.d/S99vcs stop
```

- 9** Unmount `/dev/odm`:

```
# umount /dev/odm
```

- 10** Unload the ODM module:

```
# genkex | grep odm
# vxkextadm vxodm unload
```

- 11** Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```

- 12** Remove the Storage Foundation Cluster File System 5.1 SP1 RP1 patches.

- Create a file that contains all the 5.1 SP1 RP1 patches. In this example, it is called `/reject.list`:

```
# /reject.list
```

- Reject each patch from the patch list file, for example:

```
# installp -rBfX /reject.list
```

- 13** Reboot the systems. On each system, run the following command.

```
# /usr/sbin/shutdown -r
```

## Rolling back Storage Foundation for Oracle RAC manually

Use the following procedure to roll back to 5.1 SP1 manually.

### To roll back SF for Oracle RAC manually

- 1** Stop Oracle and CRS on each node of the cluster.

- If CRS is controlled by VCS, log in as superuser on each system in the cluster and enter the following command:

```
# hastop -all
```

- If CRS is not controlled by VCS, enter the following command on each node of the cluster to stop CRS:

```
# crsctl stop crs
```

Unmount all VxFS file system used by a database or application and enter the following command to each node of the cluster:

```
# hastop -local
```

- 2 Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped. In the `gabconfig -a` command output, the VCS engine or high availability daemon (HAD) port `h` is not displayed. This indicates that VCS has been stopped.

```
# /sbin/gabconfig -a
```

Sample output:

```
GAB Port Memberships
=====
Port a gen 5c3d0b membership 01
Port b gen 5c3d10 membership 01
Port d gen 5c3d0c membership 01
Port o gen 5c3d0f membership 01
```

- 3 Remove the Storage Foundation for Oracle RAC 5.1 SP1 RP1 patches.

- Create a file that contains all the 5.1 SP1 RP1 patches. In this example, it is called `/reject.list`:

```
# /reject.list
```

You can use the following list as the reject list for Storage Foundation for Oracle components:

```
VRTSvxvm VRTSsfmh VRTSvxfv VRTSllt VRTSvxfen VRTSamf VRTSvcs
VRTScps VRTSvcsag VRTSvcssea VRTSg1m VRTScavf VRTSodm
```

- Reject each patch from the patch list file, for example:

```
# installp -rBfX /reject.list
```

- 4 Reboot the systems. On each system, run the following command.

```
# /usr/sbin/shutdown -r
```

## Rolling back Veritas Cluster Server manually

Use the following procedure to roll back VCS 5.1 SP1 RP1 to VCS 5.1 SP1 on your cluster manually. To uninstall VCS, see the *Veritas Cluster Server Installation Guide*.

---

**Note:** Use this procedure only when rolling back VCS. Do not roll back VCS when it is part of other products that rely on VCS, for example Storage Foundation Clustered File System or Storage Foundation for Oracle RAC.

---

### To roll back Veritas Cluster Server manually

- 1 Verify that all of the VCS 5.1 SP1 RP1 patches are in the `APPLIED` state.

The following is a list of the VCS 5.1 SP1 RP1 patches:

```
VRTSvcsea    5.1.101.0
VRTSvcstag   5.1.101.0
VRTScps      5.1.101.0
VRTSvcsc     5.1.101.0
VRTSamf      5.1.101.0
VRTSvxfen    5.1.101.0
VRTS11t      5.1.101.0
```

- 2 List the service groups in your cluster and their status. On any node, type:

```
# hagrps -state
```

- 3 Take the ClusterService service group offline if it is running. On any node, type:

```
# hagrps -offline -force ClusterService -any
```

- 4 Make the VCS configuration writable. On any node, type:

```
# haconf -makerw
```

- 5 Freeze all service groups except the ClusterService service group. On any node, type:  

```
# hagrps -freeze groupname -persistent
```
- 6 Save the configuration (`main.cf`) file with the groups frozen. On any node, type:  

```
# haconf -dump -makero
```
- 7 Make a backup copy of the current `main.cf` file and all the `types.cf` configuration files. For example, on one node in the cluster, type:  

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/types.cf.save
```
- 8 Shut down VCS. On any node, type:  

```
# /opt/VRTSvcs/bin/hastop -all -force
```
- 9 Shut down CmdServer. On each node, type:  

```
# /opt/VRTSvcs/bin/CmdServer -stop
```

10 Verify that VCS has shut down. On any node, type:

```
# /sbin/gabconfig -a
```

The output resembles:

GAB Port Memberships

Port a gen 23dc0001 membership 01

The output shows no membership for port h.

On each node, run the command:

```
# ps -ef | egrep "had|hashadow|CmdServer"
```

Terminate any instances of had, hashadow, or CmdServer that still run after 60 seconds.

Stop fencing, GAB, and LLT. On each node, type:

```
# /etc/init.d/vxfen.rc stop
```

```
# /etc/init.d/gab.rc stop
```

```
# /etc/init.d/llt.rc stop
```

11 Perform the patch removal. On each node, type:

```
# installp -r patch_name
```

Review the summaries at the end of each run and confirm that all of the intended patches removed successfully.

12 Remove the Storage Foundation Cluster File System 5.1 SP1 RP1 patches.

- Create a file that contains all the 5.1 SP1 RP1 patches. In this example, it is called `/reject.list`:

```
# /reject.list
```

- Reject each patch from the patch list file, for example:

```
# installp -rBfX /reject.list
```

13 After VCS has started, perform the following steps:

- Verify all resources have been probed. On any node, type:

```
# hastatus -summary
```

- Unfreeze all service groups. On any node, type:

```
# haconf -makerw  
# hagrpl -unfreeze Group Name -persistent  
# haconf -dump -makero
```

- 14 Bring the ClusterService service group online, if necessary. On any node where system is the node name, type:

```
# hagrpl -online ClusterService -sys system
```

- 15 Reboot the systems. On each system, run the following command.

```
# /usr/sbin/shutdown -r
```

## Rolling back Dynamic Multi-Pathing manually

Use the following procedure to roll back to 5.1 SP1 manually.

### To roll back DMP manually

- 1 Stop activity to all VxVM volumes.
- 2 Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

To verify that no volumes remain open, enter the following command:

```
# vxprint -Aht -e v_open
```

- 3 Perform the following commands to determine whether root support or DMP native support is enabled.

```
■ # vxdmpadm gettune dmp_native_support
```

If the command returns an "on" value, DMP native support is enabled on the system. If the command returns any other value, DMP native support is disabled.

```
■ # vxdmpadm native list vgroupname=rootvg
```

If the output is a list of hdisks, root support is enabled on this system. If the command returns any other value, root support is disabled.

- Once you have determined if root support or DMP native support is enabled, go to step 4

- Once you have determined that root support and DMP native support is not enabled, go to step 5
- 4 If root support or DMP native support is enabled:
- You must disable DMP native support.  
Run the following command to disable DMP native support and to disable root support:  

```
# vxdmpadm settune dmp_native_support=off
```
  - If only root support is enabled, run the following command to disable root support:  

```
# vxdmpadm native disable vgname=rootvg
```
  - Reboot the system:  

```
# shutdown -r now
```
  - Before backing out patch, stop the VEA server's vxsvc process:  

```
# /opt/VRTSob/bin/vxsvcctrl stop
```
  - Create a file that contains all the 5.1 SP1 RP1 patches. In this example, it is called `/reject.list`:  

```
# /reject.list
```
  - Reject each patch from the patch list file, for example:  

```
# installp -rBFX /reject.list
```
  - Reboot the system:  

```
# shutdown -r now
```
  - Enable DMP native support, this also enables root support:  

```
# vxdmpadm settune dmp_native_support=on
```
  - Reboot the system:  

```
# reboot
```
  - Verify DMP native or root support is enabled:

```
# vxdmpadm gettune dmp_native_support
```

**5** If root support or DMP native support is not enabled:

- Before you back out the patch, kill the VEA Server's vxsvc process:

```
# /opt/VRTSob/bin/vxsvcctrl stop
```

- To reject the patch if it is in `APPLIED` state

```
# installp -r patch_name
```

- Reboot the system:

```
# shutdown -r now
```