Veritas Storage Foundation™ and High Availability Solutions Release Notes

AIX

5.1 Service Pack 1 Rolling Patch 2

November 2011

Storage Foundation and High Availability Solutions Release Notes 5.1 Service Pack 1 Rolling Patch 2

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 RP2

Document version: 5.1SP1RP2.1

Legal Notice

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

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

Contacting Technical Support

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

www.symantec.com/business/support/contact_techsupp_static.jsp

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

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:

doc_feedback@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

Chapter 1

About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- Introduction
- About the installrp and the uninstallrp scripts
- Changes introduced in 5.1 SP1 RP2
- System requirements
- List of products
- Fixed issues
- Known issues
- Software limitations
- Documentation errata
- List of patches
- Downloading the 5.1 SP1 RP2 archive

Introduction

This document provides information about the Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 Rolling Patch 2 release.

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

http://www.symantec.com/docs/TECH75503

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

http://www.symantec.com/docs/TECH74012

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.

This rolling patch applies to the following releases of Storage Foundation and High Availability products:

- Storage Foundation and High Availability Solutions 5.1 SP1
- Storage Foundation and High Availability Solutions 5.1 SP1 RP1
- Storage Foundation and High Availability Solutions 5.1 SP1 PR1

This rolling patch is available as:

- 5.1 SP1 RP2
- 5.1 SP1 PR2 RP2

Given that this rolling patch applies to the previously released 5.1 SP1 platform RP releases, Symantec does not plan on the following releases:

- 5.1 SP1 PR1 RP1

About the installrp and the uninstallrp scripts

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 provides an upgrade script.

See "Supported upgrade paths" on page 92.

Symantec recommends that you use the 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 1-1 The command line options for the product upgrade script

Command Line Option	Function
`[ system1 system2... ]`	Specifies the systems on which to run the upgrade options. If not specified, the command prompts for a system name.
`[ -precheck ]`	Use the `-precheck` option to confirm that systems meet the products' installation requirements before the installation.
`[ -postcheck ]`	Use the `-postcheck` option after an installation or upgrade to help you determine installation-related problems and provide troubleshooting information.
`[ -logpath log_path ]`	Use the `-logpath` option to select a directory other than `/opt/VRTS/install/logs` as the location where the installrp log files, summary file, and response file are saved.
`[ -responsefile response_file ]`	Use the `-responsefile` option to perform automated installations or uninstallations using information stored in a file rather than prompting for information. *response_file* is the full path of the file that contains configuration definitions.
`[ -tmppath tmp_path ]`	Use the `-tmppath` option to select a directory other than `/var/tmp` as the working directory for `installrp`. This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
`[ -hostfile hostfile_path ]`	Use the `-hostfile` option to specify the location of a file containing the system names for installer.
`[ -keyfile ssh_key_file ]`	Use the `-keyfile` option to specify a key file for SSH. When you use this option the `-i` *ssh_key_file* is passed to every SSH invocation.

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

Command Line Option	Function
`[-nim]`	Use to produce a NIM configuration file for installing with NIM. Refer to the product's *Installation Guide* for more information on using NIM.
`[ -patchpath` *`patch_path`* `]`	Use the `-patchpath` 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 `installrp`.

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

Command Line Option	Function
`[ -rsh | -redirect | -listpatches | -makeresponsefile | -pkginfo | -serial | -version ]`	Use the `-rsh` option when rsh and rcp are to be forced for communication though ssh and scp is also setup between the systems. Use the `-redirect` option to display progress details without showing the progress bar. Use the `-listpatches` option to display product patches in the correct installation order. Use the `-makeresponsefile` option to generate a response file without doing an actual installation. The text that displays install, uninstall, start, and stop actions are simulations. These actions are not performed on the system. Use the `-pkginfo` option is used to display the correct installation order of packages and patches. This option is available with or without one of following options: -allpkgs, -minpkgs, and -recpkgs. Use the `-serial` option to perform installation, uninstallation, start and stop operations, typically performed simultaneously on all systems, in a serial fashion. Use the `-version` option to have the installer check and report the installed products and their versions. Identifies the installed and missing fileset and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing fileset and patches where applicable.

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

Command Line Option	Function
`[-upgrade_kernelpkgs | -upgrade_nonkernelpkgs]`	Use the `-upgrade_kernelpkgs` option for the rolling upgrade's upgrade of kernel packages to the latest version Use the `-upgrade_nonkernelpkgs` 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.

The uninstallrp script options

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 provides a new uninstallation script.

See About rolling back Veritas Storage Foundation and High Availability Solutions 5.1 SP1 RP2 for release versions and products that support rolling back.

Symantec recommends that you use the new uninstallation script. The `uninstallrp` script uninstalls all the patches associated with packages installed, and starts the processes. Do not use the `uninstallrp` script for rolling back, because it removes the entire stack.

Table 1-2 The command line options for the product upgrade script

Command Line Option	Function
`[ system1 system2... ]`	Specifies the systems on which to run the upgrade options. If not specified, the command prompts for a system name.
`[ -logpath log_path ]`	Use the `-logpath` option to select a directory other than `/opt/VRTS/install/logs` as the location where the uninstallrp log files, summary file, and response file are saved.

Table 1-2 The command line options for the product upgrade script *(continued)*

Command Line Option	Function
`[ -responsefile` *`response_file`* `]`	Use the `-responsefile` option to perform automated installations or uninstallations using information stored in a file rather than prompting for information. *`response_file`* is the full path of the file that contains configuration definitions.
`[ -tmppath` *`tmp_path`* `]`	Use the `-tmppath` option to select a directory other than `/var/tmp` as the working directory for `uninstallrp`. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
`[ -hostfile` *`hostfile_path`* `]`	Use the `-hostfile` option to specify the location of a file containing the system names for installer.
`[ -keyfile` *`ssh_key_file`* `]`	Use the `-keyfile` option to specify a key file for SSH. When you use this option the `-i` *`ssh_key_file`* is passed to every SSH invocation.

Table 1-2 The command line options for the product upgrade script *(continued)*

Command Line Option	Function
`[ -rsh | -redirect | -makeresponsefile | -serial | -version ]`	Use the `-rsh` option when rsh and rcp are to be forced for communication though ssh and scp is also setup between the systems. Use the `-redirect` option to display progress details without showing the progress bar. Use the `-makeresponsefile` option to generate a response file without doing an actual installation. Text displaying installation, uninstallation, start and stop operations are simulations. These actions are not being performed on the system. Use the `-serial` option to perform installation, uninstallation, start and stop operations, typically performed simultaneously on all systems, in a serial fashion. Use the `-version` option to have the installer check and report the installed products and their versions. Identifies the installed and missing fileset and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing fileset and patches where applicable.

Changes introduced in 5.1 SP1 RP2

This section lists the changes in 5.1 SP1 RP2.

Changes related to Storage Foundation and High Availability

Storage Foundation and High Availability includes the following changes in 5.1 SP1 RP2:

New README files provide detailed information on the included patches

This release includes README_SYMC files that provide detailed information on the patches included in this release. These README_SYMC files provide the following information:

- Patch IDs, incidents fixed through the patch
- Symptom, description, and resolution for the addressed issues

The README_SYMC files are available in the following directory:

- `/patches`

The Rolling Patch Release Notes continue to provide a summary of fixed issues.

Changes related to installing, upgrading and rolling back

The following changes are related to installing, upgrading and rolling back of the product in 5.1 SP1 RP2 release.

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 RP2, only the installrp and uninstallrp scripts can detect 5.1 SP1 RP2 versions.

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, Veritas Dynamic Multi-Pathing, and Veritas Volume Manager.

Changes related to Veritas Volume Manager

Veritas Volume Manager includes the following changes in 5.1 SP1 RP2.

ASM co-existence now enabled by default

The VxVM and ASM co-existence is now enabled by default. VxVM now identifies ASM disks automatically.

support vxcdsconvert utility for EFI disks.

The vxcdsconvert utility is now supported for EFI disks.

Changes related to Veritas Cluster Server

Veritas Cluster Server includes the following changes in 5.1 SP1 RP2:

Combo VCS 5.1SP1 RP2 release for VCS 5.1SP1 or later and VCS 5.1SP1 PR1

This means that you can install the same VCS 5.1SP1 RP2 patches over VCS 5.1SP1 or later and VCS 5.1SP1 PR1.

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 x 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 VCS5.1 SP1 RP2, please ensure that the new `Db2udbTypes.cf` file is used which contains the definition of **IMFRegList** as above.

LVMVG agent supports synchronization of ODM entries

LVMVG agent on AIX now supports synchronization of the ODM entries of all the nodes in a cluster, if the disk configuration of the volume group is changed on a node.

If physical volumes are added, deleted or replaced in a volume group, some or all of the ODM entries on other nodes become stale. This may cause the volume group resource online failure on those nodes. The updated disk information can now be propagated to all nodes in the cluster using `updatepv` action entry point.

You must run this action whenever there are changes in the disk configuration, such as addition, deletion, or replacement of physical volumes of the volume group. Running this action ensures that the other nodes are updated with the new information, which is used when the agent brings the volume group online.

Use the following command to run this action:

```
# hares -action res_name updatepv -sys system_name
```

where *system_name* is the name of the node on which the disks were added to or removed from the volume group.

Running this action sends the updated physical volume information from the local node to all the other nodes. This information is stored in the file `/var/VRTSvcs/log/tmp/resource_name.volume_group_name.pvid` on all the nodes. If this file is present on that node, and if the `SyncODM` attribute is set to 1, then the online entry point uses the PVIDs from the file, exports the volume group, breaks the reservations on all these disks, and uses any one PVID to re-import the volume group. The file is deleted from that node after the volume group is successfully brought online on that node and the ODM is synchronized.

The `SyncODM` attribute must the set to 1 if the updatepv action has been executed for that volume group.

To ensure the high availability of the LVMVG resource, you must run the `updatepv` action immediately after adding, deleting, or replacing the physical volumes in the volume group. When the new disks are added or replaced in the volume group, these disks must be visible and have same PVID on all the cluster nodes.

You must run the updatepv action again for the nodes that were down when updatepv was last run.

Note: The updatepv action does not support the GCO environment.

The full value displays when you use hares -display

For `hares -display`, the maximum number of characters was 20. Any attribute value that was more than 20 characters would be truncated.

Symantec has removed the 20 characters limit. Now the full value displays when you use `hares -display`.

Restrict the max value of FencingWeight to 9999

When Preferred Fencing is enabled, the max weight value that you can assign to the FencingWeight attribute is 9999. `had` adds 1 to every weight that you assign for each node. If the value of FencingWeight is set to 10000, VCS fails to set the node weight to 10001 because the maximum node weight accepted by vxfen driver is 10000.

Below message displays if the value of FencingWeight that you assign is equal to or more than 10000.

```
VCS WARNING V-16-1-50003 FencingWeight should be an \
integer between [0..9999] inclusive
```

Changes to Application Agent

The following are changes about Application Agent for shared disk support

Shared disk support

Symantec enhanced the Application agent to support use of shared disks for StartProgram, StopProgram, CleanProgram and MonitorProgram attributes.

Support for UNIX style return values in MonitorProgram

The Application agent handles the standard UNIX style return values, that is, "0" for success and "1" for failure, and now reports the resource state based on following set of values:

100 or 1 --> OFFLINE

101 to 110 or 0 --> ONLINE

Any other value --> UNKNOWN.

Validating the user that is specified in the User attribute

The Application agent is modified to validate the user name configured in the User attribute. If the user does not exist on the system, the agent reports the state of a configured resource on that system as UNKNOWN.

Changes related to Veritas Dynamic Multi-Pathing

Veritas Dynamic Multi-Pathing includes the following change in 5.1 SP1 RP2:

Performing an operating system upgrade with a single reboot when DMP root support is enabled

When DMP root support is enabled, follow this procedure to upgrade the operating system with a single reboot.

To upgrade the operating system with a single reboot:

1 Change to the methods directory:

```
# cd /usr/lib/methods/
```

2 Unlink the DMP boot device:

```
# unlink vxdmpboot
```

Relink the DMP boot device:

```
# ln -s vxdmpboot.target_os vxdmpboot
```

Where *target_os* is the AIX operating system version to which you are upgrading the host.

3 Enable root support. Do not reboot the system.

4 Upgrade the operating system.

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 7.1 TL0 or later

- 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 TL5 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://www.symantec.com/docs/TECH75503

Database requirements

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

http://www.symantec.com/docs/TECH74389

Note: Veritas Storage Foundation (SF) and Veritas Storage Foundation Cluster File System (SFCFS) do not support running SFDB tools with DB2 and Sybase, but they support running Oracle, DB2, and Sybase on VxFS and VxVM.

http://www.symantec.com/docs/TECH44807

Additional Oracle support for SF Oracle RAC

Table 1-3 Oracle RAC versions that SF Oracle RAC supports

Oracle version	AIX 5.3	AIX 6.1	AIX 7.1
10gR2 10.2 (64-bit)	Yes	Yes	No
11gR1 11.1 (64-bit)	Yes	Yes	No
11gR2 11.2.0.2 (64-bit)	Yes	Yes	Yes

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

http://www.symantec.com/docs/TECH44807

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
  - For a minimum of 256 MB for 1 node and a maximum of 1 GB of swap space for 8 or more nodes

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.

See the README_SYMC.*xxxxx-xx* files in the /patches directory on the installation media for the symptom, description, and resolution of the fixed issue.

Veritas Volume Manager fixed issues

This section describes Veritas Volume Manager fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Volume Manager: Issues fixed in 5.1 SP1 RP2

Table 1-4 describes the incidents that are fixed in Veritas Volume Manager in 5.1 SP1 RP2.

Table 1-4 Veritas Volume Manager 5.1 SP1 RP2 fixed issues

Fixed issues	Description
1791397	VVR:RU thread keeps spinning sending START_UPDATE message repeatedly to the secondary
1675599	Memory leaks in DDL and ASLs
2484685	Race between two vol_subdisk sios while doing 2one2processing which causes one thread to free sio_fsvm_priv before other thread accesses it.
2480600	I/O permanent hung on master node when IO size larger than 512K, and 32+ threads write in parallel.
2440349	DCO volume may grow into any 'site' even when 'alloc=site:xxxx' is specified by a list of 'site' to be limited.
2431470	vxpfto uses DM name when calling vxdisk, but vxdisk will match DA name first and thus cause corruption
2431423	CVR: Panic in vol_mv_commit_check after I/O error on DCM
2428875	I/O on both nodes (wait for the DCM flush started), and crash the slave node, lead to the master reconfiguration hang
2428631	Allow same fence key to be used for all Disk groups
2425722	vxsd move operation failed for disk size greater than or equal to 2 TB

Table 1-4 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2425551	IO hung for 6 mintues when reboot the slave node, if there is I/O on both master and slave.
2424833	Pinnacle while autosync_deport#2 primary logowner hits ted assert nmcom_send_msg_tcp
2421067	Vxconfigd hung in both nodes of primary
2419348	DMP panic: race between dmp reconfig and dmp pass through ioctl
2413904	Multiple issues are seen while performing Dynamic LUN reconfiguration.
2411698	VVR:iohang: On I/O to both master and slave
2410845	Lots of 'reservation conflict' messages seen on 5.1SP1RP1P1 clusters with XIV arrays.
2408771	vxconfigd does not scan and discover all the storage device; some storage devices are skipped.
2407192	Application I/O hangs because of race between CVM reconfiguration and Log-owner change protocol.
2406292	Panic in vol_subdisksio_delete()
2400654	Stale array.info file can cause vxdmpadm commands to hang
2400014	Boot image cannot handle 3 kernel extension versions (AIX 5.3, 6.1 & 7.1) when rootability is enabled
2396293	I/Os loaded, sanboot failed with vxconfigd core dump.
2387993	While testing including/excluding libvxpp.so vxconfigd goes into disabled mode.
2386120	Enhancement request to add diagnostic logging to help triage a CVM master takeover failure situation
2385680	vol_rv_async_childdone+1147
2384473	Ensure vxcdsconvert works with support for greater than 1 TB CDS disks
2383158	VVR: vxio panic in vol_rv_mdship_srv_done+680
2379029	Changing of enclosure name is not working for all devices in enclosure

Table 1-4 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2371685	default tunable parameter volpagemod_max_memsz not updated to 64MB when upgraded 5.1 bits to 5.1SP1RP1 bits
2369786	VVR:A deadloop about NM_ERR_HEADR_IO
2369177	DDL: do_diskio function should be able to handle offset greater than 2TB
2365951	Growto failing with error V-5-1-10128 Unexpected kernel error in configuration update
2364253	VVR: Kernel memory is leaked on VVR secondary while using SO snapshots
2360404	vxmirror operation fails with error "Device has UFS FS on it"
2359814	vxconfigbackup doesn't handle errors well
2357798	CVR:Memory leak due to unfreed vol_ru_update structure
2357507	In presence of large number of NR (Not-Ready) devices, server panics due to NMI triggered and when DMP continuously generates large no of path disable/enable events.
2356744	VxVM script daemons should not allow its duplication instance in itself
2349352	During LUN provisioning in single path IO mode environment a data corruption is observed
2346470	Excluding and including a LUN in a loop triggers a huge memory leak
2337694	TP "vxdisk -o thin list" showing size 0 for over 2TB LUNs
2337353	vxdmpadm include vxvm dmpnodename=*emcpower#* includes all excluded dmpnodes along with the requested one
2334534	In CVM environment, vxconfigd level join is hung when Master returns error "VE_NO_JOINERS" to a joining node and cluster nidmap is changed in new reconfiguration
2322752	Duplicate DA records seen for NR devices upon restart of vxconfigd
2320917	vxconfigd core dump and lost dg config after removing volume and disk on thin reclaim LUN.
2317703	Vxesd/Vxconfigd leaks file descriptors.
2316297	After applying 5.1SP1RP1 error message "Device is in use" appears during boot time

Table 1-4 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2299670	Disk Groups created on EFI LUNs do not auto import at boot time using VxVM version 5.1SP1 and later
2286559	kernel heap corruption detected panic after array controller reboot
2263317	CLONE: Diskgroup import with dgid needs to be clearly documented in manual for the case in which original dg was destroyed and cloned disks are present.
2257678	vxinstall failing due to incorrectly determining boot disk is encapsulated
2255182	Handling misconfiguration of CLARiiON array reporting one failovermode value through one HBA and different from other HBA
2253970	Support per-disk maxiosize for private region I/Os
2253552	Leak in vxsfdefault_parse.y at function vxsf_getdefault (*val)
2249113	vol_ru_recover_primlog_done return the same start address to be read from SRL, if the dummy update is greater than MAX_WRITE
2248730	vxdg import command hangs as vxrecover daemon (spawned by vxdg) doesn't close standard error stream
2242268	panic in voldrl_unlog
2240056	'vxdg move' transaction not completing and backups fail.
2237089	vxrecover might start the recovery of data volumes before the recovery of the associated cache volume is recovered.
2232789	Supporting NetApp Metro Cluster
2228531	cvm master vxconfigd process hung in vol_klog_lock()
2205108	SVS 5.1SP1: vxconfigd clubbing all luns in a single dmpnode
2204752	Multiple VM commands succeed but throw "GPT entries checksum mismatch" error message for hpdisk format.
2200670	vxattachd does not recover disks if disk group is not imported
2197254	While creating volumes on thinrclm disks, the option "logtype=none" does not work with vxassist command.
2196918	Snapshot creation with cachesize fails, as it doesn't take into account diskgroup alignment.

Table 1-4 Veritas Volume Manager 5.1 SP1 RP2 fixed issues *(continued)*

Fixed issues	Description
2196480	The disk initialization failed due to wrong number of cylinders reported in devintf_disk_geom_raw() gotten from raw geometry
2194685	vxconfigd daemon core dump during array side switch ports disable and re-enable.
2193429	IO policy not getting preserved when vold is restarted and migration from one devlist to other is taking place.
2190020	Complains dmp_deamon applying 1m continuous memory paging is too large
2179259	DMP SCSI bypass needs to be enhanced to handle I/O greater than 2TB
2165394	CLONE: dg imported by selecting wrong disks. After destroying original dg, when try to import clone devices without useclonedev option with dgname, then it import dg with original disks.
2154287	Improve handling of Not-Ready(NR)devices which are triggering "VxVM vxdmp V-5-3-1062 dmp_restore_node: Unstable path" messages
2152830	In multilevel clone disks environment, regular DG import should be handled properly and in case of DG import failure, it should report correct error message
2144775	Failoverpolicy "local" is not getting preserved after upgrade from 5.1RP1/Sles10Sp2 to 5.1Sp1/Sles10Sp3.
2139179	SSB check invalid when lun copy
2094672	CVR: vxconfigd on master hangs while reconfig is running in cvr stress with 8 users
2033909	In SF-RAC configuration, IO hung after disable secondary path of A/PG array Fujitsu ETERNUS3000

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 1-5 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 coredumps 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
2158438	vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core
2159947	Bump up the dmpslab_minsz to 512 elements

Table 1-5 Veritas Volume Manager 5.1 SP1 RP1 fixed issues *(continued)*

Fixed issues	Description
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 fsvm 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	'vxdmpadm 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
2234844	asm2vxfs conversion fails
2215216	vxkprint does not report TP related values

Veritas File System fixed issues

This section describes Veritas File System fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas File System: Issues fixed in 5.1 SP1 RP2

Table 1-6 describes the incidents that are fixed in Veritas File System in 5.1 SP1 RP2.

Table 1-6 Veritas File System fixed issues

Fixed issues	Description
2529356 (2340953)	cfs.stress.enterprise hit an assert f:vx_iget:1a.
2523084 (2515101)	VxFS crash conflict with svmon
2515569 (2515559)	LM conformance -> aixopen test get panic issues
2508164 (2481984)	file system will hang if customer creates 400 shares
2494464 (2247387)	LM stress.S3 test hit an assert "vx_ino_update:2"
2487976 (2483514)	System panic due to OS upgarde from AIX 5.3 to 6.1
2486597 (2486589)	threads blocked behind vx_ireuse_steal
2482337 (2431674)	panic in vx_common_msgprint() via vx_inactive()
2480949 (2480935)	fsppadm: ERROR: V-3-26626: File Change Log IOTEMP and ACCESSTEMP index creation failure for /vx/fsvm with message Argument list too long
2478237 (2384861)	CFS stress+reconfig test hit assert "f:vx_do_filesnap:1b".

Table 1-6 Veritas File System fixed issues *(continued)*

Fixed issues	Description
2432898 (N/A)	fsvoladm remove failed with "ERROR: V-3-20: 0000:ioctl on /oradata failed: Arg list too long", "UX:vxfs fsvoladm: ERROR: V-3-25572:"
2427281 (2413172)	There is a priority 1 issue reported by AXA Rosenburg for Filestore replication and issue seems related to VxFS
2427269 (2399228)	TRuncate up size updates can be missed
2426065 (2430794)	AXRT51SP1RP2:removing a volume from volume set file system failed by ERROR: V-3-25572
2426039 (2412604)	it does not work when set homedir user softlimit numspace quota after generate data
2425429 (2422574)	Reboot one node and the node can't mount file system , after turn on the homedir quota on
2420060 (2403126)	cfs recovery didn't finished timely in the primary node after one slave left.
2418819 (2283893)	Add functionality of free space defragmentation through fsadm.
2412187 (2401196)	glm panic while doing dynamic reconfiguration of LDPAR
2412181 (2372093)	new fsadm -C hung
2412179 (2387609)	User quota corruption
2412177 (2371710)	user quota information corrupts on 5.1SP1
2412137 (2346730)	Need to find out how much vxglm used at kernel pinned memory.

Table 1-6 Veritas File System fixed issues *(continued)*

Fixed issues	Description
2412029 (2384831)	vxfs panic in iput() from vx_softcnt_flush() ,after filesystem full fsck,and run reboot
2406572 (2146573)	"qdetails" performance downgraded on Aug 16th.
2405590 (2397976)	AIX6.1 SF 5.1SP1 - EXCEPT_DSI panic
2402643 (2399178)	fsck : pass2c needs performance enhancements
2386483 (2374887)	Accessing FS hung. FS marked full fsck after reboot of node.
2386478 (2346730)	Need to find out how much vxglm used at kernel pinned memory.
2373565 (2283315)	cfs-stress_S5 hits assert of "f:vx_reorg_emap:10 via vx_extmap_reorg"
2368738 (2368737)	RCQ processing code should set FULLFSCK flag if it finds a corrupt indirect block.
2360821 (1956458)	fsckpt_fbmap for changed blocks failed with ENXIO due to inode mapped to hole in ILIST of down stream checkpoint
2360819 (2337470)	In the process of shrink fs, the fs out of inodes, fs version is 5.0MP4HF*
2360817 (2332460)	vxedquota slow on some systems
2341007 (2300682)	Question about IOTemp on fsppadm query
2340839 (2316793)	After removing files df command takes 10 seconds to complete

Table 1-6 Veritas File System fixed issues *(continued)*

Fixed issues	Description
2340834 (2302426)	Unaligned Reference Fault in vx_copy_getemap_structs
2340831 (2272072)	Threads stuck in vx_rwsleep_rec_lock_em
2340825 (2290800)	investigation on ilist HOLE
2340817 (2192895)	VxFS 5.0MP3RP4 Panic while set/get acls - possible race condition
2340799 (2059611)	Panic in vx_unlockmap() due to NULL ml_tranp
2340741 (2282201)	vxdump core dumped whilst backing up layout 7 local VxFS file system
2329893 (2316094)	There was discrepancy between vxi_bcache_maxkbyte and vx_bc_bufhwm.
2329887 (2253938)	EAU delegation timeouts
2320044 (2419989)	ncheck -i does not limit output to the specified inodes when using -o device/block/sector
2311490 (2074806)	dm_punch_hole request does not invalidate pages
2296277 (2296107)	Operation not applicable appear on fsppadm query result
2280552 (2246579)	Panic at getblk() when growing a full filesystem with fsadm
2280386 (2061177)	fsadm -de' command erroring with 'bad file number' on filesystem(s) on 5.0MP3RP1

Table 1-6 Veritas File System fixed issues *(continued)*

Fixed issues	Description
2275543 (1475345)	write() system call hangs for over 10 seconds on VxFS 3.5 on 11.23
2257904 (2251223)	df -h after removing files takes 10 seconds
2255786 (2253617)	LM stress aborted due to "run_fsck : Failed to full fsck cleanly".
2249658 (2220300)	vx_sched' is hogging CPU resources.
2247299 (2161379)	repeated hangs in vx_event_wait()
2243063 (1949445)	hang due to large number of files in a directory
2243061 (1296491)	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2169326 (2169324)	5.1SP1 sol_sprac Test LM-stress_S5 hits an assert of "f:vx_idelxwri_off:5a vai vx_trunc_tran"

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 1-7 Veritas File System fixed issues

Fixed issues	Description
1929221	vxrepquota truncating username and groupname to 8 characters is addressed.
2030119	fsppadm 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.

Table 1-7 Veritas File System fixed issues *(continued)*

Fixed issues	Description
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_ilist spin lock with the default values for `vx_idelxwri_timelag`.

Veritas Storage Foundation Cluster File System fixed issues

This section describes Veritas Storage Foundation Cluster File System fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation Cluster File System: Issues fixed in 5.1 SP1 RP2

Table 1-8 describes the incidents that are fixed in Veritas Storage Foundation Cluster File System in 5.1 SP1 RP2.

Table 1-8 Veritas Storage Foundation Cluster File System fixed issues

Fixed issues	Description
2406572 (2146573)	qdetails performance downgraded

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 1-9 Veritas Storage Foundation Cluster File System fixed issues

Fixed issues	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 CFSMount 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 Storage Foundation for Oracle RAC fixed issues

This section describes Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Storage Foundation for Oracle RAC: Issues fixed in 5.1 SP1 RP2

Table 1-10 describes the incidents that are fixed in Veritas Storage Foundation for Oracle RAC in 5.1 SP1 RP2.

Table 1-10 Veritas Storage Foundation for Oracle RAC fixed issues

Fixed issues	Description
2374977	Oracle instance crashed; failure occurred at: vcsipc_dosnd
2390892	memory leak in vcsmm_set_cluster_proto
2429449	The cssd agent explicitly uses hard-coded string "cssd" as resource name.
2374970	CRSResource agent support for 11gR2

Veritas Storage Foundation for Oracle RAC: Issues fixed in 5.1 SP1 RP1

There are no fixed issues in this release.

Veritas Cluster Server fixed issues

This section describes Veritas Cluster Server fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP2

Table 1-11 describes the incidents that are fixed in Veritas Cluster Server in 5.1 SP1 RP2.

Table 1-11 Veritas Cluster Server 5.1 SP1 RP2 fixed issues

Fixed Issues	Description
2518609	Clean EP of WPAR does not stop the WPAR
2516926	Changes to Application agent to support physical to virtual failover.
2516856	Changes to Mount agent to support physical to virtual failover.
2512840	Changes to Oracle, Netlsnr and ASMInst agents to support physical to virtual failover.
2511385	Sybase online script should honor RECOVERY state of the database.
2508637	system got crashed when uninstall GAB package.
2483044	Changes to VCS engine to skip state update requests when resource is already in that state

Table 1-11 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2481086	LLT: Improve LLT-over-UDP performance
2477372	LLT: reduce "lltd" CPU consumption by reducing the wakeup calls
2477305	Changes to WPAR agent to support physical to virtual failover
2477296	Application service group did not fail over on node panic
2477280	Application resource is not failover when system reboot after Concurrency Violation
2476901	Changes to IP agent to support physical to virtual failover.
2439895	LLT: lltconfig reports its own cluster node as part of duplicate cluster
2439772	WAC resource offline failed after network interruption
2438261	Failed to perform online migration from scsi raw to scsi dmp policy.
2435596	NFS resource failed to come online with NFSv4 on AIX, because of local domain not set on machine.
2434782	ContainerInfo should be allowed to be set for group that is already online.
2426663	On OCPR from customized mode to scsi3 mode, vxfend does not terminate
2426572	Changes to VCS engine to reports a persistent resource as FAULTED when a system is added to group using hagrp -modify command.
2423990	Changes to Application agent to handle non-existent user
2416842	Changes to VCS engine to gracefully handle the case when the file descriptor limit is exhausted.
2411860	Agent entry points time out due to non-responsive NFS mount
2411653	GAB: Add check for MAX message size in GAB
2407755	Changes to the agent framework to avoid memory allocation between fork and exec system calls.
2406748	Changes to AMF module to prevent registration of already online process for offline monitor with AMF.
2405391	LLT: The arp ack packet should include the nodename of the node
2403851	AMF status is showing Module loaded but not configured.

Table 1-11 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2403782	Sybase agent should use perl file I/O for password file specified in SApswd attribute with "VCSSY:" key.
2403633	ContainerInfo attribute should be allowed to be updated even when Group is not completely offline
2400485	Once vxfenconfig -c with mode A has returned EFAULT ("1036 Unable to configure..."), all subsequent runs of vxfenconfig -c with mode B fail with error EBADMSG ("1050 Mismatched modes...")
2400330	whyonlining does not behave as advertised in VCS 5.1SP1
2399898	hagrp -switch of child group fails if 2 or more parent groups online on alternate node
2399658	System panicked while executing the installrp to update RP1.
2398807	VCS should set a soft limit for file descriptors in /opt/VRTSvcs/bin/vcsenv
2394176	vxfenswap process hangs, "ps -ef" shows "vxfenconfig -o modify" on one node but not on other. "vxfenswap -a cancel" kills the stuck operation.
2386326	cannot configure fencing, vxfenadm prints same Serial Number for all LUNs which have more than 96 bytes of SCSI Inquiry data in page 0x83
2382583	CP Agent doesn't show coordination point information in engine log when CP server is not accessible.
2382582	Vxfen tunable resetting when node comes up after panic.
2382575	Cannot modify VxFEN tunable parameters
2382559	Online Migration fails with the message I/O fencing does not appear to be configured on node.
2382493	Parent service group does not failover in case of online local firm dependency with child service group.
2382460	Configuring fencing is successful with 3 disks even when single_cp=1 and formatting of warning messages required in vxfend_A.logo
2382452	Syntax errors while unconfiguring CP server using configure_cps.pl scripto.
2382335	vxfentsthdw fails to choose the same fencing disk on two nodes.

Table 1-11 Veritas Cluster Server 5.1 SP1 RP2 fixed issues *(continued)*

Fixed Issues	Description
2380922	The default route gets deleted when two IPMultiNICB resources are configured for a given network and one resource is brought offline.
2377788	IPMultiNICB agent dumps core when configured for IPv6
2367721	The owner.vfd virtual fire drill check of Oracle agent should only match the uid and gid from the id command output.
2366201	Allow Fencing to start when a majority of the coordination points are available.
2354932	hacli -cmd' triggers had coredump
2330980	When a node is added or deleted from the Group's SystemList, notifications about resources should be sent to agents running only on the added or deleted systems.
2330045	RemoteGroup resource does not go offline when network fails
2330041	VCS group dependencies do not online parallel parent group after upgrading SF 5.0MP3 RP2 to SF5.1SP1.
2318334	Oracle agent should set $Oracle_home/lib library to be first in LD_LIBRARY_PATH
2301731	Panic in amf_lock() due to bad mutex during system shutdown.
2296172	Failover Policy does not take into consideration AutoFailover = 2 and SystemZones when the nodes within a SystemZone are brought down or rebooted.
2276622	Cannot configure SCSI-3 fencing using RamSan DMP devices.
2271882	MonitorMethod attribute does not reflect IMF value without setting Listener attribute on Netlsnr Resource
2253349	When netmask changed outside VCS, VCS should show a warning message
2393939	Enhanced Apache agent version parsing to accommodate IBM HTTP server 7.0.

Veritas Cluster Server: Issues fixed in 5.1 SP1 RP1

This section describes the incidents that are fixed in Veritas Cluster Server in 5.1SP1RP1 release.

Table 1-12 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 1-12 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.

Storage Foundation Manager fixed issues

This section describes Storage Foundation Manager fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Storage Foundation Manager: Issues fixed in 5.1 SP1 RP2

There are no Storage Foundation Manager fixed issues in this release.

Storage Foundation Manager: Issues fixed in 5.1 SP1 RP1

There are no Storage Foundation Manager fixed issues in this release.

Veritas Storage Foundation for Databases (SFDB) tools fixed issues

This section describes Veritas Storage Foundation for Databases (SFDB) tools fixed issues in 5.1 SP1 RP2 and 5.1 SP1 RP1.

Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP2

Table 1-13 describes the incidents that are fixed in Storage Foundation for Databases (SFDB) tools in 5.1 SP1 RP2.

Table 1-13 Veritas Storage Foundation for Databases (SFDB) tools fixed issues

Fixed issues	Description
2429359	dbed_update does not work on AIX 7.1 Power 7
2509867	vxdbed looping doing read/write to IDLE sockets

Storage Foundation for Databases (SFDB) tools: Issues fixed in 5.1 SP1 RP1

There are no SFDB fixed issues in 5.1 SP1 RP1.

Veritas Enterprise Administrator fixed issues

This section describes Veritas Enterprise Administrator fixed issues in 5.1 SP1 RP2.

Veritas Enterprise Administrator: Issues fixed in 5.1 SP1 RP2

Table 1-14 describes the incidents that are fixed in Veritas Enterprise Administrator fixed issues in 5.1 SP1 RP2.

Table 1-14 Veritas Enterprise Administrator fixed issues

Fixed issues	Description
2394915	VEA service(vxsvc) running on port 2148 crashes and dumps core

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.

Incorrect version listed after upgrading (2121881)

When you upgrade from Veritas product 5.1 SP1 to Veritas product 5.1 SP1 RP2, the previous version is incorrectly listed as 5.1.001.000

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

The installer displays the space that all the product fileset and patches needs. During migration some fileset are already installed and during migration some fileset 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.

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...
```

installrp fails to install 5.1SP1RP2 when the root user shell is set to csh (2523643)

VCS installation fails if super user (root) logged-in is using C shell (csh). Currently the installer does not support c-shell (/usr/bin/csh).

Workaround: Change your super-user (root) shell to shell (/usr/bin/sh) and retry the 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 enviroment, but it has not been tested or verified.

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

After upgrading SF stack from 5.0 MP3 SP1 RP2 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_rac11dg1: No such disk
group SFORA
vxsnapadm ERROR V-81-5623 Could not get CVM information for
SNAP_rac11dg1.
SFORA dbed_vmclonedb ERROR V-81-5578 Import SNAP_rac11dg1 failed.
```

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

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

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.

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 vxdctl 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 5.1SP1RP2 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.

Co-existence check might fail for CDS disks

In Veritas Volume Manager (VxVM) 5.1 SP1, 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.

Changes in enclosure attributes are not persistent after an upgrade to VxVM 5.1 SP1 RP2 (2082414)

The Veritas Volume Manager (VxVM) 5.1 SP1 RP2 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 RP2, 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 RP2. Manually reconfigure the enclosure attributes to resolve the issue.

Table 1-15 shows the Hitachi arrays that have new array names.

Table 1-15 Hitachi arrays with new array names

Previous name	New name
TagmaStore-USP	Hitachi_USP
TagmaStore-NSC	Hitachi_NSC

Table 1-15 Hitachi arrays with new array names *(continued)*

Previous name	New name
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 RP2. 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 specifed (in seconds) by the tunable dmp_restore_interval.

```
# vxdmpadm 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.

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 set`write_pref_io` and `write_nstream` to high values, such that `write_pref_io` multiplied by`write_nstream` is around 8 MB.

Online migration can miss copying some files and directories if you issue the umount command during a migration (2097960)

When the `umount` command is issued on a file system, the namespace momentarily disappers even before the `umount` command executes.This is true even for a busy `umount`. Because of this, if you run the `umount` command on a file system while an online migration is in progress, the background copy process can miss copying some files and directories. Files and directories that are being copied in the window when the namespace disappears are skipped. The background copy will not be able to detect this because calls such as `opendir`() fail with the ENOENT error at that time. This results in giving a false migration completion.

Workaround: Do not unmount a file system that is being online migrated.

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 fials 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.

Possible write performance degradation with VxFS local mounts

Some applications that allocate large files without explicit preallocation may exhibit reduced performance with the VxFS 5.1 release 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 data 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.

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.

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 data 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_sizevalue 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.

NFS issues with VxFS checkpoint (1974020)

NFS clients mounting VxFS checkpoints that are NFS-exported by SFCFS cluster nodes using a Virtual IP may receive the following error message upon Virtual IP failover:

```
Stale NFS file handle
```

This is a result of major numbers of VxFS checkpoints not necessarily being the same on all SFCFS cluster nodes.

There is no workaround at this time.

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
   ```

Veritas Cluster Server known issues

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

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

Parent service groups fail to restart after a child service group that has recovered from a fault restarts (2330038)

A child service group that has recovered from a fault will not be able to restart its faulted parent service group, if the parent service group's fault is detected before the child service group's fault.

Workaround:

Set the child service group's `OnlineClearParent` attribute to 1. When the child service group recovers from a fault and comes online, VCS clears the fault of the parent service group. This allows the VCS to bring the parent service group online.

hacmd -display G or A or O or E or S or C dumps core (2415578)

VCS ha commands do not work when object names are single character long.

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.

had stopped before vxfen startup (2429272)

VCS stops with error `VCS CRITICAL V-16-1-10031 VxFEN driver not configured. VCS Stopping. Manually restart VCS after configuring fencing`.

If UseFense is set to SCSI3, upon starting VCS checks if the VxFen driver is configured. If VxFen driver is configured to use CPS or if there is a pre-existing split brain, the driver takes a long time to complete configuration. VCS periodically queries the driver until the driver is configured or exits after 90 seconds.

Workaround

VCS must be restarted manually after the VxFen driver is configured.

If LinkTestRatio is set to 1 Group is going to be in faulted state (2492608)

When using the MultiNICB resource with IPv6 protocol, the value of the `LinkTestRatio` attribute must be 0. The `MultiNICB resource` shows unexpected behavior when `LinkTestRatio` is set to some other value.

Application Agent does not handle a case when user is root, envfile is set and shell is csh. (2490299)

The Application Agent uses the `system` command to execute the Start/Stop/Monitor/Clean Programs for root user. This executes Start/Stop/Monitor/Clean Programs in sh shell, due to which there is an error when root user has csh shell and `EnvFile` is written accordingly.

Workaround:

Do not set csh as shell for root user. Use sh as shell for root instead.

When the WPAR is down, IMF registration does not succeed for an application resource running inside of WPAR (2529278)

If an application agent runs inside of WPAR, when WPAR is down, it is not able to retrieve user information. It fails to check the user verification and it is not able to register the resource with AMF.

Workaround: There is no workaround for this issue. Only IMF monitoring does not function for that application resource. Traditional monitoring functions without any problem.

Pre-requisites for the hawparsetup.pl script (2523171)

- If a service group already exists, the `hawparsetup.pl` script does not check whether the service group is completely OFFLINE or not. If the service group where the WPAR resource needs to be added already exists and is not OFFLINE, the `hawparsetup.pl` script does not modify the `ContainerInfo` attribute for the system.
  Workaround: If a service group already exists, make sure that it is completely OFFLINE before running the `hawparsetup.pl` script.
- The `hawparsetup.pl` script does not check whether the key `Enabled` in attribute `ContainerInfo` has been set or not. If the `ContainerInfo` attribute is already set for a service group and the key `Enabled` is set to some value other than 1, running the `hawparsetup.pl` script will over-write the value for key *Enabled* to 1.
  Workaround: After running the `hawparsetup.pl` script, manually set the value of key `Enabled` in attribute `ContainerInfo` to the desired value.

Forcefully un-configuring AMF does not change the monitor method of agent to TRADITIONAL (2564376)

If you forcefully unconfigure AMF when a script based agent is registered with AMF, the monitor method of the agent is not changed to `TRADITIONAL`. It remains `IMF`.

Workaround: Restarting the agent will resolve the issue.

Forcefully un-configuring AMF causes the engine log to be flooded with error messages (2535690)

If you forcefully unconfigure AMF when a script based agent is registered with AMF, the `getnotification` thread continuously polls and displays error messages in the engine log.

Workaround: Restarting the agent will resolve the issue.

OracleTypes.cf needs updates for WPAR support (2163956)

To support the Oracle and the Netlsnr agents in WPAR environment, you must modify the OracleTypes.cf file to include the following attributes in both Netlsnr and Oracle type:

```
static int ContainerOpts{}
= { RunInContainer=1, PassCInfo=0 }
```

Workaround: To modify the OracleTypes.cf for WPAR support

1. Stop VCS on all nodes.

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

2. Edit the `OracleTypes.cf` file on one of the nodes. Add the following line to Netlsnr and Oracle type:

   ```
   static int ContainerOpts{}
   = { RunInContainer=1, PassCInfo=0 }
   ```

3. Start VCS on the node with modified configuration:

   ```
   # hastart
   ```

4. After VCS has started on the node, start VCS on other cluster nodes:

   ```
   # hastart
   ```

NFS resource goes offline on its own and errors out when restarted (2490415)

If multiple agent processes are running because an agent process is restarted multiple times by _had, then only one of the agent process survives and other agent processes go offline on its own. Even though the agent process is running, _had does not recognize it and hence does not perform any resource operations.

Workaround: Kill the agent process to recover from this situation. Refer to the engine log for further actions (if required) to restart the agent.

HAD dumps core when hagrp -clear is executed on a group in OFFLINE|FAULTED state and a resource in the fault path is waiting to go online (2536404)

This issue occurs if you have a resource dependency, such as r1 -> r2 -> r3. While resources r2 and r3 are online and you initiate bringing resource r1 online, before the OnlineTimeout occurs, resources r2 and r3 suffer a fault. Resource r2 faults first, and then r3 faults. After the fault of both resources is detected, the group is becomes in an OFFLINE|FAULTED state and resource r1 is stuck waiting to become online. If you execute the `hagrp -clear` command to clear the fault, then HAD dumps core on all nodes due to assertion failure.

Workaround: Flush the pending online operation using the `hagrp -clear` command before clearing the fault.

Errors observed while rollback of VRTSvxfen patch (2556918)

The error message **"/usr/lib/methods/vxfenext -stop -dvxfend failed."** displays when you remove the fencing package (VRTSvxfen) or reject the patch. You will also see a pseudo device file `/dev/vxfend` left in the system after the removal of the package.

Workaround: After the removal of VRTSvxfen package or rejection of patch, manually delete the `/dev/vxfend` node.

The hares -display command fails if the resource is part of a global service group (2358600)

The `hares -display` command incorrectly processes the response received from the had process. Due to the processing error, `hares -display` does not show the resource details.

Workaround: Use the -localclus or -clus option with `hares -display`.

Excessive delay messages in one-node configurations of Storage Foundation High Availability (SFHA) (2486415)

The GAB error, "`Excessive delay between successive calls to GAB heartbeat`" appear in the engine log while running a single node or standalone VCS cluster where GAB is disabled.

GAB heartbeat log messages are logged as informational when there is delay between heartbeats (HAD being stuck). When HAD runs in -onenode, GAB does not need to be enabled. When HAD is running in -onenode, for self-checking purposes, HAD simulates heartbeat with an internal component of HAD itself. These log messages are logged because of delay in simulated heartbeats.

Workaround: Log messages are for informational purpose only. When HAD is running in -onenode, no action is needed on excessive delay between heartbeats.

Agent does not retry resource registration with IMF to the value specified in RegisterRetryLimit key of IMF attributes (2411882)

Agent maintains internal count for each resource to track the number of times a resource registration is attempted with IMF. This count gets incremented if any attempts to register a resource with IMF fails. At present, any failure in un-registration of resource with IMF, also increments this internal count. This means that if resource un-registration fails, next time agent may not try as many resource registrations with IMF as specified in RegisterRetryLimit. This issue is also observed for multiple resources.

Workaround: Increase the value of RegisterRetryLimit if un-registration and registration of one or more resources is failing frequently. After increasing the RegisterRetryLimit, if resource registration continues to fail, report this problem to Symantec.

Oracle agent incorrectly reports the global resource as online when the resource inside the local zone is online and the Sid's are same (2561563)

Oracle agent incorrectly reports the resource configured for Oracle instance running in global container as online, if the resource configured for Oracle instance running in local container also has same value for Sid attribute and the instance in local container is online.

The above issue is also applicable for ASMInst and Netlsnr agents.

For Netlsnr agent the above issue appears when the Home and listener attributes of the resources running in global and local container are same.

The issue does not appear for Oracle and ASMInst agents when multiple local containers have resources configured with the same value of Sid attribute.

The issue does not appear for Netlsnr agent when multiple local containers have resources configured with the same value of Home and listener attributes.

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:dg1: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 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, ues 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 snapshop_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 snapshop_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.

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 RP2 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.

Failure to set the MTU (Maximum Transmission Unit) size in LLT over UDP environments causes issues with the PrivNIC/MultiPrivNIC agents (2557144)

If the MTU size field is not set explicitly when you configure the PrivNIC/MultiPrivNIC agents in an LLT over UDP environment, the agents may fail to plumb the private IP addresses during their operations or may configure incorrect MTU size on the LLT interfaces.

The agents use the `lltstat -l` command to retrieve MTU size information for LLT interfaces. In an LLT over UDP environment, the command retrieves 8192 as the MTU size. When the PrivNIC/MultiPrivNIC agents use this size information to plumb the IP addresses, the operation may fail causing the agents to fault. However, even if the plumbing operation succeeds, the incorrect MTU configuration may still cause issues in the cluster later.

Workaround:

To update the PrivNIC/MultiPrivNIC resource configuration in an LLT over UDP environment

1 Retrieve the MTU size of the network interfaces configured under PrivNIC/MultiPrivNIC agents:

For AIX: `# lsattr -El en1`

For HPUX: `# lanadmin -m 1`

For Linux: `# ifconfig eth1`

For Solaris: `# ifconfig ce0`

2 Set the MTU attribute for the PrivNIC/MultiPrivNIC resource:

```
# haconf -makerw
```

Run the following command for all the network interfaces configured under PrivNIC/MultiPrivNIC agents:

```
# hares -modify resource_name MTU -add interface_name mtu_size
```

Where:

resource_name is the name of the PrivNIC/MultiPrivNIC resource

interface_name is the name of the network interface for which the MTU size is set

mtu_size is the MTU size retrieved in step 1.

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

Veritas Enterprise Administrator known issues

The following is a Veritas Enterprise Administrator known issue in this 5.1 SP1 RP2 release.

After uninstalling 5.1SP1RP2 patch on AIX 7, the file "/etc/vx/isis/Registry.conf" shows "Version" as 3.4.235.0 instead of 3.4.290.0 (2557174)

Workaround: This issue is safe to ignore.

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 RP2 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 SP1 or 5.1 SP1 PR1 to version 5.1 SP1 RP2. If you are running earlier versions of Veritas product, perform an upgrade to version 5.1 SP1 or 5.1 SP1 PR1 and then upgrade to version 5.1 SP1 RP2 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 1-16

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

The change is persistent across reboots.

To change the tunable parameters

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

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

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

   ```
   # vxdmpadm gettune dmp_restore_internal
   # vxdmpadm 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.

PowerPath 5.5 managing EMC storage (2433061, 2478307)

Disks may go into error state in the following two scenarios:

- Changing failover mode for EMC Symmetrix or Clariiion LUNs from array side
- Unmanaging disks from PowerPath view and initiating VxVM discovery (vxdctl enable/vxdisk scandisks)

Workaround: If disks go into error state, you can try using `vxddladm -c assign` *`names`* to clear the names, and use the default OSN or EBN names.

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 5.1 SP1 RP2 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 Storage Foundation for Databases tools software limitations

The following are software limitations in this release of Veritas Storage Foundation for Databases.

Oracle Data Guard in an Oracle RAC environment

Database snapshots and Database Checkpoints are not supported in an environment where there are both Data Guard and Oracle RAC. But separately, either Data Guard or Oracle RAC supports Database snapshots and Database Checkpoints.

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 5.1 SP1 RP2: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to 5.1 SP1 RP2.

Documentation errata

The following section provides documentation updates.

Veritas Cluster Server Administrator's Guide (2444653)

In the "VCS environment variables" section, the definition of the variable `VCS_GAB_RMTIMEOUT` should be "Timeout in milliseconds for HAD to register with GAB."

The value of VCS_GAB_RMTIMEOUT is specified in milliseconds. The minimum value is 200000 milliseconds or 200 seconds. If VCS_GAB_RMTIMEOUT is less than the minimum value then VCS overrides and sets it to 200000 milliseconds.

Veritas Installation Guides (2521411)

The ignorepatchreqs option is no longer valid and does not work.

List of patches

This section lists the patches for 5.1 SP1 RP2.

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

Table 1-17 Patches for AIX

BFF file	Size in bytes	Patches	Version
VRTSamf.bff	3737600	VRTSamf	05.01.0112.0000
VRTScavf.bff	307200	VRTScavf	05.01.0112.0000
VRTScps.bff	48076800	VRTScps	05.01.0112.0000
VRTSdbac.bff	8345600	VRTSdbac	05.01.0112.0000
VRTSdbed.bff	39168000	VRTSdbed	05.01.0112.0000
VRTSgab.bff	6195200	VRTSgab	05.01.0112.0000
VRTSglm.bff	768000	VRTSglm	05.01.0112.0000
VRTSgms.bff	307200	VRTSgms	05.01.0112.0000
VRTSllt.bff	3481600	VRTSllt	05.01.0112.0000
VRTSob.bff	58572800	VRTSob	03.04.0312.0000
VRTSodm.bff	921600	VRTSodm	05.01.0112.0000
VRTSsfmh.bff	39270400	VRTSsfmh	03.01.0429.0401
VRTSvcs.bff	319436800	VRTSvcs	05.01.0112.0000
VRTSvcsag.bff	19456000	VRTSvcsag	05.01.0112.0000
VRTSvcsea.bff	6297600	VRTSvcsea	05.01.0112.0000

Table 1-17 Patches for AIX *(continued)*

BFF file	Size in bytes	Patches	Version
VRTSvxfen.bff	4198400	VRTSvxfen	05.01.0112.0000
VRTSvxfs.bff	34611200	VRTSvxfs	05.01.0112.0000
VRTSvxvm.bff	266598400	VRTSvxvm	05.01.0112.0000

Downloading the 5.1 SP1 RP2 archive

The patches that are included in the 5.1 SP1 RP2 release are available for download from the Symantec website. After downloading the 5.1 SP1 RP2 rolling patch, use gunzip and tar commands to uncompress and extract it.

For the 5.1 SP1 RP2 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

http://www.symantec.com/docs/TECH75503

Chapter 2

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 RP2. Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

See the 5.1 SP1 *Installation Guide* and *Release Notes* or 5.1SP1 PR1 *Installation Guide* for your product for more information.

See "Upgrading to 5.1 SP1 RP2" on page 92.

To install the Veritas software for the first time

1. Download Storage Foundation and High Availability Solutions 5.1 SP1 from http://fileConnect.symantec.com.
2. Extract the tar ball into a directory called /tmp/sfha51sp1.
3. Check http://sort.symantec.com/patches to see if there are any patches available for the 5.1SP1 Installer. Download applicable P-patches and extract them to the /tmp directory.

4 Change the directory to /tmp/sfha51sp1:

```
# cd  /tmp/sfha51sp1
```

5 Run the installer to install SFHA 5.1SP1. See the Installation Guide for instructions on installing the 5.1 SP1 version of this product.

```
# ./installer -require complete_path_to_SP1_installer_patch
```

6 Download SFHA 5.1 SP1 RP2 from http://sort.symantec.com/patches.

7 Extract it to a directory called /tmp/sfha51sp1rp2.

8 Check http://sort.symantec.com/patches to see if there are patches available for the 5.1SP1RP2 installer. Download applicable P-patches and extract them to the /tmp directory.

9 Change the directory to /tmp/sfha51sp1rp2:

```
# cd  /tmp/ sfha51sp1rp2
```

10 Invoke the `installrp` script to install 5.1SP1 RP2:

```
# installrp -require complete_path_to_SP1RP2_installer_patch
```

11 If you did not configure the product after the 5.1 SP1 installation, the installer prompts you to configure the product during RP installation. If you do not want to configure the product now, answer **n** when prompted. To configure the product in the future, run the product installation script from the 5.1 SP1 installation media or from /opt/VRTS/install directory 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 RP2 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.

See “Upgrading to 5.1 SP1 RP2” on page 92.

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 `xprtlwid`, 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 5.1 SP1 RP2 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 or VCS 5.1 SP1 PR1 version of the Veritas product must be installed before upgrading to 5.1 SP1 RP2.

2 On the **Select a task and product** page, select **Install SP1 RP2** 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 RP2 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.

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 RP2

This chapter includes the following topics:

Prerequisites for upgrading to 5.1 SP1 RP2

The following list describes prerequisites for upgrading to the 5.1 SP1 RP2 release:

- For any product in the Veritas Storage Foundation stack, you must have the 5.1 SP1 (or later) or 5.1SP1 PR1 installed before you can upgrade that product to the 5.1 SP1 RP2 release.
- Each system must have sufficient free space to accommodate patches.
- The full list of prerequisites can be obtained by running `./installrp -precheck`
- Make sure to download the latest patches for the installer.
  See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 91.

Downloading required software to upgrade to 5.1 SP1 RP2

This section describes how to download the latest patches for the installer.

To download required software to upgrade to 5.1 SP1 RP2

1. Download SFHA 5.1 SP1 RP2 from http://sort.symantec.com/patches.
2. Extract it to a directory, say /tmp/sfha51sp1rp2.
3. Check http://sort.symantec.com/patches to see if there are patches available for the 5.1 SP1 RP2 installer. Download applicable P-patches and extract them to the /tmp directory.
4. When you run the `installrp` script, use the `-require` option and specify the location where you downloaded the patches.

Supported upgrade paths

You can upgrade to this release of Veritas product from version 5.1 SP1 (or later) or VCS 5.1SP1 PR1.

Upgrading to 5.1 SP1 RP2

This section describes how to upgrade from 5.1 SP1 (or later) to 5.1 SP1 RP2 on a cluster or a standalone system.

- Performing a full upgrade to 5.1 SP1 RP2 on a cluster
  Use the procedures to perform a full upgrade to 5.1 SP1 RP2 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 RP2 on a standalone system
  Use the procedure to upgrade to 5.1 SP1 RP2 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.

See “Installing the Veritas software using the script-based installer” on page 87.

Performing a full upgrade to 5.1 SP1 RP2 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.

Depending on your cluster’s configuration, select one of the following procedures to upgrade to 5.1 SP1 RP2:

- Performing a full upgrade to 5.1 SP1 RP2 on a Veritas Cluster Server
- Performing a full upgrade to 5.1 SP1 RP2 on an SFHA cluster
- Performing a full upgrade to 5.1 SP1 RP2 on an SFCFS cluster
- Performing a full upgrade to 5.1 SP1 RP2 on an SF Oracle RAC cluster
  See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 91.

Performing a full upgrade to 5.1 SP1 RP2 on a 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. Make sure you have downloaded the latest software required for the upgrade.

   See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 91.

2. Log in as superuser.

   Note: Upgrade the Operating System and reboot the systems if required.

3. Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 SP1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check:

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

4. Resolve any issues that the precheck finds.

5. Start the upgrade:

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

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

Performing a full upgrade to 5.1 SP1 RP2 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 RP2 on an SFHA cluster

1. Make sure you have downloaded the latest software required for the upgrade.

   See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 91.

2. Log in as superuser.

3. Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.

4. On each node, 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 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.

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 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 5.1 SP1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the pre-upgrade check.

```
# ./installrp -precheck [-rsh] node1 node2 ... nodeN
```

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

11 Review the output as the program displays the results of the check and saves the results of the check in a log file.

12 Make sure all your nodes meet the prerequisites, and all the issues reported by above pre-check process have been resolved.

13 Start the upgrade.

```
# ./installrp [-rsh] node1 node2 ... nodeN
```

Review the output.

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

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

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

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

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

```
# mount /filesystem
```

17 Remount all Storage Checkpoints on all nodes in the selected group:

```
# mount /checkpoint_name
```

Performing a full upgrade to 5.1 SP1 RP2 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 RP2 on an SFCFS cluster

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 91.

2 Log in as superuser.

3 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.

4 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
```

5 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
  ```

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 `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.

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 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
```

9 If required, apply the OS kernel patches.

See IBM's documentation for the procedures.

10 On each node, 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
```

11 From the directory that contains the extracted and untarred 5.1 SP1 RP2 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.

12 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.

13 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node.

14 Bring the CVM service group online on each node:

```
# hagrp -online cvm -sys nodename
```

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

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

16 If you stopped any RVGs in step 6 , restart each RVG:

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

17 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

18 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

Performing a full upgrade to 5.1 SP1 RP2 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 RP2 on a SF Oracle RAC cluster

1 Make sure you have downloaded the latest software required for the upgrade.

See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 91.

2 Log in as superuser.

3 Verify that `/opt/VRTS/bin` is in your **PATH** so that you can execute all product commands.

4 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

And set the the AutoStart attribute of Oracle Agent to 0:

```
# hagrp -modify oracle_group AutoStart 0
# haconf -dump -makero
```

5 If the Oracle DB is not managed by VCS, prevent auto startup of Oracle DB:

```
# srvctl modify database -d db_name -y manual
```

6 Stop Oracle database on the cluster:

- If the Oracle RAC instance is managed by VCS:

```
# hagrp -offline oracle_group -sys galaxy
# hagrp -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 stop database -d db_name
```

7 Stop all applications on the cluster that are not configured under VCS. Use native application commands to stop the application.

8 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
```

- Unmount the VxFS or CFS file system:

```
# umount /mount_point
```

9 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
```

10 Stop VCS.

```
# hastop -all
```

11 From the directory that contains the extracted and untarred 5.1 SP1 RP2 rolling patch binaries, change to the directory that contains the installrp script. Start the upgrade.

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

12 Manually mount the VxFS and CFS file systems that are not managed by VCS.

13 Start all applications on the cluster that are not configured under VCS. Use native application commands to start the application.

14 Relink the SF Oracle RAC libraries with Oracle.

Refer to *Veritas Storage Foundation for Oracle RAC 5.1 SP1 or later Installation and Configuration Guide* for more information.

15 Start Oracle Group on All nodes.

```
# hagrp -online oracle_group -any
```

16 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 start database -d db_name
```

17
- If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

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

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

  ```
  $ srvctl modify database -d db_name -y AUTOMATIC
  ```

Upgrading to 5.1 SP1 RP2 on a standalone system

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

To upgrade to 5.1 SP1 RP2 on a standalone system

1. Make sure you have downloaded the latest software required for the upgrade.

   See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 91.

2. Log in as superuser.

3. Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.

4. If required, apply the OS kernel patches.

   See IBM's documentation for the procedures.

5. Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

   ```
   # mount | grep vxfs
   ```

6. Unmount all Storage Checkpoints and file systems:

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

7. 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.

8 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.

9 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
```

10 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
```

11 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
```

12 If necessary, reinstate any missing mount points in the `/etc/filesystems` file.

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

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

14 If you stopped any RVGs in step 6, restart each RVG:

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

15 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem
# mount /checkpoint_name
```

16 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl 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 for VCS, SFHA, SFCFS, and SFCFSHA: phase 1
- Performing a rolling upgrade on non-kernel packages for VCS, SFHA, SFCFS, and SFCFSHA: phase 2
- Performing a rolling upgrade on kernel packages for SF Oracle RAC: phase 1
- Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: 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 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.
- Make sure you have downloaded the latest software required for the upgrade. See "Downloading required software to upgrade to 5.1 SP1 RP2 " on page 91.

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 for VCS, SFHA, SFCFS, and SFCFSHA: 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: 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 Veritas File System 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 for VCS, SFHA, SFCFS, and SFCFSHA: 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 nodeD
```

2 The installer checks system communications, fileset 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
```

6 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

Performing a rolling upgrade on kernel packages for SF Oracle RAC: phase 1

Note that in the following instructions that a subcluster can represent one or more nodes in a full cluster, but is represented by nodeA, nodeB as subcluster1 and nodeC, nodeD as subcluster2.

To perform the rolling upgrade on kernel: phase 1

1. Log in as superuser to one of the nodes in the cluster.

2. Back up the following configuration files on your system: main.cf, types.cf, CVMTypes.cf, CFSTypes.cf, OracleTypes.cf, OracleASMTypes.cf,PrivNIC.cf, MultiPrivNIC.cf, /etc/llttab, /etc/llthosts./etc/gabtab, /etc/vxfentab, /etc/vxfendg, /etc/vxfenmode.

   For example:

   ```
   # cp /etc/VRTSvcs/conf/config/main.cf \
         /etc/VRTSvcs/conf/config/main.cf.save
   # cp /etc/VRTSvcs/conf/config/types.cf \
         /etc/VRTSvcs/conf/config/types.cf.save
   # cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
         /etc/VRTSvcs/conf/config/OracleTypes.cf.save
   # cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
         /var/tmp/PrivNIC.cf.save
   # cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
         /var/tmp/MultiPrivNIC.cf.save
   ```

3. If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

   ```
   # haconf -makerw
   # hagrp -modify oracle_group AutoStart 0
   ```

   If the Oracle database is not managed by VCS, change the management policy for the database to manual. Execute the command with oracle database user credentials.

   ```
   $ srvctl modify database -d db_name -y manual
   ```

4. If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to SF Oracle RAC 6.0.

5. - If the applications are not under VCS control, stop the applications that use VxFS or VxVM disk groups on each node of subcluster, whether local or CFS. Use native application commands to stop the application.
   - If the database instances are not managed by VCS, stop the Oracle RAC database resources on each node of subcluster, run the following from one node. Execute the command with oracle database user credentials.

```
$ srvctl stop instance -d db_name -i instance_name
```

6 Unmount all the VxFS file system which is not under VCS control on each node of subcluster.

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

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -c /mount_point
```

7 On subcluster, stop all VxVM and CVM volumes (for each disk group) that are not managed by VCS:

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

Verify that no volumes remain open under the diskgroups which are not managed by VCS:

```
# vxprint -g disk_group -ht -e v_open
```

Note: Installer will automatically stop all the applications, database instances, filesystems and volumes which are under VCS control on nodes, while using upgrade_kernelpkgs option.

8 On the sub-cluster, start the installer for the rolling upgrade with the -upgrade_kernelpkgs option.

```
# ./installrp -upgrade_kernelpkgs nodeA nodeB
```

9 The installer checks system communications, fileset versions, product versions, and completes prechecks. It will stop/failover the applications, database instances, filesystems which are under VCS control. It then upgrades applicable product kernel.

10 Reboot the upgraded sub-cluster.

Note: The Oracle service group at this point will be offline as the AutoStart attribute is set to 0 to prevent the service group from starting automatically.

11 Relink the SF Oracle RAC libraries with Oracle on upgraded subcluster by choosing the option Relink Oracle Database Binary from the program menu.

12 Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

  ```
  # hagrp -online oracle_group -sys nodeA
  # hagrp -online oracle_group -sys nodeb
  ```

- If VCS does not manage the Oracle database:

  ```
  $ srvctl start instance -d db_name-I instance_name
  ```

13 Start all applications that are not managed by VCS. Use native application commands to start the applications.

14 Before you proceed to phase 2, complete step 4 to 13 on the remaining subcluster.

15 Perform one of the following steps:

- If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

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

- If the VCS does not manage the Oracle database, change the management policy for the database to automatic:

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

16 Migrate the SFDB repository database.

Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: 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 nodeD
   ```

2. The installer checks system communications, fileset versions, product versions, and completes prechecks. It verifies completion of phase 1.
3. Installer will start HA daemon (had) on all nodes, HA will be available once HA daemon is up.
4. Verify the cluster's status:

   ```
   # hastatus -sum
   ```

5. If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

Verifying software versions

To list the Veritas filesets installed on your system, enter the following command:

```
# lslpp -L VRTS\*
```

The output version for 5.1 SP1 RP2 is `5.1.112.0`.

Chapter 4

Rolling back and removing 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 RP2

This section describes how to roll back either by using the `uninstallrp` script or manually.

Rolling back using the uninstallrp script

Use the following procedure to roll back from any Veritas product to the previous version using the `uninstallrp` script.

To roll back

1. Browse to the directory that contains the uninstallrp script.
2. Stop all the processes and services accessing the file systems. 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 (or later) or 5.1 SP1 PR1 commited level, and 5.1 SP1 RP2 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 RP2 patches. After patch rollback completes, modules are loaded and processes are restarted. `uninstallrp` will also report any warning happened during uninstallation.

- For other products:

1. Run the `uninstallrp` command, type:

   ```
   # ./uninstallrp system_list
   ```

2. If you performed a roll back on a system that has an encapsualted boot disk, you must reboot the system. After reboot, you may need to run `hagrp -list Frozen=1` to get the frozen SG list . Then run `hagrp -unfreeze <group> -persistent` to unfreeze all the frozen SGs manually.

Rolling back manually

Use one of 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.1

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 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.

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 on each node:

   ```
   # /etc/rc.d/rc2.d/S97vxfen stop
   ```

10. Stop GAB:

    ```
    # /etc/rc.d/rc2.d/ S92gab stop
    ```

11. Stop LLT:

    ```
    # /etc/rc.d/rc2.d/S70llt stop
    ```

12. Unmount `/dev/odm`:

    ```
    # umount /etc/rc.d/rc2.d/S99odm
    ```

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 Use `uninstallrp` to roll back to 5.1SP1.

```
# uninstallrp node1 node2 ... nodeN
```

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 `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
```

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/S99vcs stop
```

9 Unmount `/dev/odm`:

```
# umount /dev/odm
```

10 Unload the ODM module:

```
# genkex | grep odm
# vxkextadm vxodm unload
```

11 Stop I/O fencing on each node:

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

12 Stop GAB:

```
# /etc/rc.d/rc2.d/S92gab stop
```

13 Stop LLT:

```
# /etc/rc.d/rc2.d/S70llt stop
```

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 Cluster File System 5.1 SP1 RP2 patches.

- Create a file that contains all the 5.1 SP1 RP2 patches. In this example, it is called `/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 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 Oracle Clusterware is controlled by VCS, log in as superuser on each system in the cluster and enter the following command:

  ```
  # hastop -all
  ```

- If Oracle Clusterware is not controlled by VCS, enter the following command on each node of the cluster to stop Oracle Clusterware:

  ```
  # 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 Bring down the rest of the stack:

Stop vcsmm:

```
# /etc/rc.d/rc2.d/S98vcsmm stop
```

Stop lmx:

```
# /etc/rc.d/rc2.d/S71lmx stop
# /usr/lib/methods/lmxext -stop
```

Stop odm:

```
# /etc/rc.d/rc2.d/S99odm stop
```

Stop vxgms:

```
# /etc/methods/gmskextadm unload
```

Stop vxglm:

```
# /etc/methods/glmkextadm unload
```

Stop vxfen:

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

Stop gab:

```
# /sbin/gabconfig -U
# /etc/methods/gabkext -stop
```

Stop llt:

```
# /sbin/lltconfig -U
```

4 Remove the Storage Foundation for Oracle RAC 5.1 SP1 RP2 patches.

- Create a file that contains all the 5.1 SP1 RP2 patches. In this example, it is called `/reject.list`:

  You can use the following list as the reject list for Storage Foundation for Oracle components:

  ```
  VRTSamf VRTScavf VRTScps VRTSdbac VRTSdbed VRTSgab VRTSglm VRTSllt
  VRTSodm  VRTSvcs VRTSvcsag VRTSvcsea VRTSvxfen VRTSvxfs VRTSvxvm
  ```

- Reject each patch from the patch list file, for example:

```
# installp -rBfX /reject.list
```

5 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 RP2 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 5.1 SP1 RP2:

1 Verify that all of the VCS 5.1 SP1 RP2 patches are in the APPLIED state. Create a text file called filesets.to.reject that contains the name and version of each fileset, one per line, exactly as shown below.

```
VRTSamf                   5.1.112.0
VRTScps                   5.1.112.0
VRTSgab                   5.1.112.0
VRTSllt                   5.1.112.0
VRTSvcs                   5.1.112.0
VRTSvcsag                 5.1.112.0
VRTSvcsea                 5.1.112.0
VRTSvxfen                 5.1.112.0
```

2 On each node, make a local copy of filesets.to.reject and then type:

```
# nohdr='^Z$'
  # while read pkg ver; do
  lslpp -l  $pkg | egrep -v "$nohdr"
  nohdr='^  Fileset +Level  State '
  done  < filesets.to.reject
```

Note: Any updates that are in COMMITTED state cannot be rejected (undone). You must remove each one and then re-install it.

3 List the service groups in your cluster and their status. On any node, type:

```
# hagrp -state
```

4 Take the ClusterService service group offline if it is running. On any node, type:

```
# hagrp -offline -force ClusterService -any
```

5 Make the VCS configuration writable. On any node, type:

```
# haconf -makerw
```

6 Freeze all service groups except the ClusterService service group. On any node, type:

```
# hagrp -list | sort -u +0b -1 | \
   while read grp sys ; do
       hagrp -freeze $grp -persistent
   done
```

You can safely ignore the warning about the failure to freeze the ClusterService group.

7 Save the configuration (main.cf) file with the groups frozen. On any node, type:

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

8 Make a backup copy of the current main.cf and all 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
```

9 Shut down VCS. On any node, type:

```
# /opt/VRTSvcs/bin/hastop -all -force
```

10 Shut down CmdServer. On each node, type:

```
# /opt/VRTSvcs/bin/CmdServer -stop
```

11 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
  ```

  Output for membership for port h does not appear.

- 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.

12 Stop AMF, fencing, GAB, and LLT.

```
# /etc/rc.d/rc2.d/S93amf stop
# /etc/rc.d/rc2.d/S97vxfen stop
# /etc/methods/vxfenext -stop
# /etc/rc.d/rc2.d/S92gab stop
# /etc/methods/gabkext -stop
# /etc/rc.d/rc2.d/S70llt stop
```

13 Preview the patch removal selection and validity tests. On each node, type:

```
# installp -pr -gXv -f filesets.to.reject
```

Confirm that the patches to be removed are exactly the same as those listed in the filesets.to.reject file that you created in step 1.

14 Perform the patch removal. On each node, type:

```
# installp -r -gXv -f filesets.to.reject
```

Review the summaries at the end of each run and confirm that all of the intended patches removed successfully.

15 Reboot all nodes in the cluster.

16 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
   # hagrp -list | sort -u +0b -1 | \
    while read grp sys ; do
     hagrp -unfreeze $grp -persistent
        done
   # haconf -dump -makero
```

You can safely ignore the warning about the failure to unfreeze the ClusterService group.

17 Bring the ClusterService service group online, if necessary. On any node, type:

```
# hagrp -online ClusterService -sys system
```

where system is the node name.

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 vgname=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 RP2 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
     ```

Removing the Veritas product

Use one the following procedures to remove the Veritas product.

Removing 5.1 SP1 RP2 on Veritas Storage Foundation or Veritas Storage Foundation Cluster File System

You can use the following procedure to uninstall 5.1 SP1 RP2 on SF or SFCFS.

To uninstall 5.1 SP1 RP2 on Veritas Storage Foundation or Veritas Storage Foundation Cluster File System

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 `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
```

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 VXFEN:

```
# /etc/rc.d/rc2.d/S97vxfen stop
```

10 Unload the ODM module:

```
# genkex | grep odm
# vxkextadm vxodm unload
```

11 Stop GAB:

```
# /etc/rc.d/rc2.d/S92gab stop
```

12 Check if the LLT is running:

```
# lltconfig
# lltstat -nvv
```

If the LLT is running, stop it:

```
# /etc/rc.d/rc2.d/S70llt stop
```

13 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
```

14 To shut down and remove the installed Veritas packages, use the appropriate command in the /opt/VRTS/install directory. For example, to uninstall the Storage Foundation or Veritas Storage Foundation Cluster File System, enter the following commands:

```
# cd /opt/VRTS/install
# ./uninstallsfcfs [-rsh]
```

You can use this command to remove the packages from one or more systems. For other products, substitute the appropriate script for `uninstallsf` such as `uninstallsfcfs` for the Storage Foundation Cluster File System software. The `-rsh` option is required if you are using the remote shell (RSH) rather than the secure shell (SSH) to uninstall the software simultaneously on several systems.

Note: Provided that the remote shell (RSH) or secure shell (SSH) has been configured correctly, this command can be run on a single node of the cluster to install the software on all the nodes of the sub-cluster.

15 After uninstalling the Veritas software, refer to the appropriate product's 5.1 SP1 Installation Guide document to reinstall the 5.1 SP1 software.

Removing 5.1 SP1 RP2 on Veritas Storage Foundation for Oracle RAC

You can use the following procedure to uninstall the 5.1 SP1 RP2 on Storage Foundation for Oracle RAC systems.

Note: This procedure will remove the complete SF for Oracle RAC stack from all nodes.

To uninstall the 5.1 SP1 RP2 on Veritas Storage Foundation for Oracle RAC

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 Uninstall Storage Foundation for Oracle RAC.

```
# cd /opt/VRTS/install
# ./uninstallsfrac MyNode1 MyNode2
```

See the *Veritas Storage Foundation for Oracle RAC 5.1 SP1 Installation and Configuration Guide* for more information.

4 After uninstalling the packages, refer to the *Veritas Storage Foundation for Oracle RAC 5.1 SP1 Installation and Configuration Guide* to reinstall the 5.1 SP1 software.