

# Veritas Storage Foundation™ and High Availability Solutions Release Notes

HP-UX

5.0.1 Rolling Patch 2



# Veritas Storage Foundation™ and High Availability Solutions Release Notes

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

Product version: 5.0.1 RP2

Document version: 5.0.1RP2.0

## Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas, Veritas Storage Foundation 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.

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

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

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

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

# Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

## Contacting Technical Support

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

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

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

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

- Product release level

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

## Licensing and registration

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

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

## Customer service

Customer service information is available at the following URL:

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

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

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

## Documentation feedback

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:

[sfha\\_docs@symantec.com](mailto:sfha_docs@symantec.com)

## Support agreement resources

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

Asia-Pacific and Japan

[customercare\\_apac@symantec.com](mailto:customercare_apac@symantec.com)

Europe, Middle-East, and Africa

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

North America and Latin America

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

# Contents

Technical Support .....	4	
Chapter 1	Veritas Storage Foundation™ and High Availability Solutions Release Notes .....	9
	Introduction .....	9
	About Veritas Operations Services .....	10
	System requirements .....	11
	What's new in this release .....	12
	New in Veritas Storage Foundation .....	12
	New in Veritas Storage Foundation for Oracle .....	12
	New in SF Oracle RAC .....	12
	HP IVM 4.2 qualification .....	12
	Fixed issues in this release .....	13
	Veritas Storage Foundation fixed issues .....	13
	Veritas Storage Foundation for Oracle fixed issues .....	21
	VCS fixed issues .....	22
	Veritas Storage Foundation Cluster File System fixed issues .....	22
	SF Oracle RAC fixed issues .....	24
	Known issues in this release .....	24
	Veritas Storage Foundation 5.0.1 Rolling Patch 2 known issues .....	24
	Veritas Storage Foundation for Oracle 5.0.1 Rolling Patch 2 known issues .....	26
	VCS 5.0.1 Rolling Patch 2 known issues .....	27
	SF Oracle RAC 5.0.1 Rolling Patch 2 known issues .....	28
	Downloading the patches .....	28
	List of patches .....	28
	Veritas Storage Foundation patches in 5.0.1 Rolling Patch 2 .....	29
	Veritas Storage Foundation for Oracle patches in 5.0.1 Rolling Patch 2 .....	30
	VCS patches in 5.0.1 Rolling Patch 2 .....	31
	Veritas Storage Foundation Cluster File System patches in 5.0.1 Rolling Patch 2 .....	32

Veritas Storage Foundation for Oracle RAC patches in 5.0.1	
Rolling Patch 2 .....	34
Installing the patches .....	37
Installing Veritas Storage Foundation 5.0.1 Rolling Patch 2 .....	37
Installing Veritas Storage Foundation High Availability 5.0.1	
Rolling Patch 2 .....	38
Installing VCS 5.0.1 Rolling Patch 2 .....	38
Installing Veritas Storage Foundation Cluster File System 5.0.1	
Rolling Patch 2 .....	51
Installing SF Oracle RAC 5.0.1 Rolling Patch 2 .....	60
Uninstalling the patches .....	67
Uninstalling Veritas Storage Foundation 5.0.1 Rolling Patch	
2 .....	68
Uninstalling VCS 5.0.1 Rolling Patch 2 .....	69
Uninstalling Veritas Storage Foundation Cluster File System	
5.0.1 Rolling Patch 2 .....	71
Uninstalling SF Oracle RAC 5.0.1 Rolling Patch 2 .....	73
Reference documentation .....	77

## Appendix A

Installing or upgrading to Oracle RAC 11g Release	
2 .....	79
Installing Oracle RAC 11g Release 2 .....	79
Preparing to install Oracle RAC using the SF Oracle RAC installer	
or manually .....	80
Installing Oracle Grid Infrastructure using the Oracle Universal	
Installer .....	92
Configuring LLT links in the GPnP profile .....	95
Installing the Oracle RAC database using the Oracle Universal	
Installer .....	95
Relinking the SF Oracle RAC libraries with Oracle RAC	
manually .....	97
Configuring the CSSD resource manually .....	97
Preventing automatic startup of Oracle Clusterware .....	99
Creating the Oracle RAC database .....	99
Preventing automatic database startup .....	99
Configuring VCS service groups for Oracle RAC .....	100
Upgrading to Oracle RAC 11g Release 2 .....	107
Supported upgrade paths .....	107
Preparing to upgrade Oracle RAC .....	108
Upgrading Oracle RAC binaries .....	108
Migrating the Oracle RAC database .....	109
Performing post-upgrade tasks .....	109



# Veritas Storage Foundation™ and High Availability Solutions Release Notes

This chapter includes the following topics:

- [Introduction](#)
- [System requirements](#)
- [What's new in this release](#)
- [Fixed issues in this release](#)
- [Known issues in this release](#)
- [Downloading the patches](#)
- [List of patches](#)
- [Installing the patches](#)
- [Uninstalling the patches](#)
- [Reference documentation](#)

## Introduction

This document provides information about the products in Veritas Storage Foundation and High Availability Solutions 5.0.1 Rolling Patch 2. Symantec

strongly recommends installing the 5.0.1 Rolling Patch 2 immediately after installing Veritas Storage Foundation and High Availability Solutions 5.0.1 or 5.0.1 Rolling Patch 1.

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

<http://entsupport.symantec.com/docs/331560>

Review this entire document before installing and upgrading your Veritas Storage Foundation and High Availability product.

For further details, depending on the product for which you want to install this Rolling Patch, refer to one of the following release notes:

- *Veritas Storage Foundation Release Notes (Version 5.0.1)*
- *Veritas Cluster Server Release Notes (Version 5.0.1)*
- *Veritas Storage Foundation for Oracle RAC Release Notes (Version 5.0.1)*

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

- Veritas Storage Foundation (SF)
- Veritas Storage Foundation for Oracle (SFORA)
- Veritas Storage Foundation High Availability (SFHA)
- Veritas Storage Foundation Cluster File System (SFCFS)
- Veritas Cluster Server (VCS)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)

## About Veritas Operations Services

Veritas Operations Services (VOS) is a set of Web-based tools and services that lets you proactively manage your Symantec enterprise products. VOS automates and simplifies administration tasks, so you can manage your data center operations more efficiently and get the most out of your Symantec products.

VOS lets you do the following:

- Collect, analyze, and report on server configurations across UNIX or Windows environments. You can use this data to do the following:
  - Assess whether your systems are ready to install or upgrade Symantec enterprise products
  - Tune environmental parameters so you can increase performance, availability, and use

- Analyze your current deployment and identify the Symantec products and licenses you are using
- Upload configuration data to the VOS Web site, so you can share information with coworkers, managers, and Symantec Technical Support
- Compare your configurations to one other or to a standard build, so you can determine if a configuration has "drifted"
- Search for and download the latest product patches
- Get notifications about the latest updates for:
  - Patches
  - Hardware Compatibility Lists (HCLs)
  - Array Support Libraries (ASLs)
  - Array Policy Modules (APMs)
  - VCS Agents
- Determine whether your Symantec enterprise product configurations conform to best practices
- Search and browse the latest product documentation
- Look up error code descriptions and solutions

To access VOS, go to:

<http://vos.symantec.com/>

## System requirements

For information on system requirements, refer to the product documentation for Veritas Storage Foundation and High Availability Solutions 5.0.1.

---

**Note:** This release requires that Version 5.0.1 or 5.0.1 Rolling Patch 1 is installed on your systems.

---

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.

For the latest information on supported hardware, visit:

<http://entsupport.symantec.com/docs/283161>

Symantec recommends installing the latest HP-UX patches from HP.

## What's new in this release

This section lists the new features in Veritas Storage Foundation and High Availability Solutions 5.0.1 Rolling Patch 2.

- Changes in Veritas Storage Foundation:  
See “[New in Veritas Storage Foundation](#)” on page 12.
- Changes in Veritas Storage Foundation for Oracle:  
See “[New in Veritas Storage Foundation for Oracle](#)” on page 12.
- Changes in Veritas Storage Foundation for Oracle RAC:  
See “[New in SF Oracle RAC](#)” on page 12.
- HP IVM 4.2 qualification:  
See “[HP IVM 4.2 qualification](#)” on page 12.

### New in Veritas Storage Foundation

This release contains the following new features:

- Quota support  
In Veritas File System (VxFS), quota support is now added for the user `nobody`.
- Read-ahead performance improvement  
In VxFS, the read-ahead performance is now improved.

### New in Veritas Storage Foundation for Oracle

This release contains the following new feature:

- Oracle 11g R2 support on PA-RISC-based systems  
Oracle Database 11g Release 2 is now supported on PA-RISC-based systems.  
Oracle 11g R2 is also supported on IA (Itanium Architecture) based systems.

### New in SF Oracle RAC

SF Oracle RAC now supports Oracle RAC 11g Release 2.

### HP IVM 4.2 qualification

HP Integrity Virtual Machines 4.2 is qualified with Veritas Storage Foundation and High Availability Solutions 5.0.1 Rolling Patch 2. For more information, refer to the document *Veritas Storage Foundation and High Availability Solutions Application Note: Support for HP-UX Integrity Virtual Machines* that is available at: <http://seer.entsupport.symantec.com/docs/332614.htm>.

## Fixed issues in this release

This section describes issues fixed in this release.

- Veritas Storage Foundation:  
See [“Veritas Storage Foundation fixed issues”](#) on page 13.
- Veritas Storage Foundation for Oracle:  
See [“Veritas Storage Foundation for Oracle fixed issues”](#) on page 21.
- Veritas Cluster Server:  
See [“VCS fixed issues”](#) on page 22.
- Veritas Storage Foundation Cluster File System:  
See [“Veritas Storage Foundation Cluster File System fixed issues”](#) on page 22.
- Veritas Storage Foundation for Oracle RAC:  
See [“SF Oracle RAC fixed issues”](#) on page 24.

## Veritas Storage Foundation fixed issues

[Table 1-1](#) lists the Veritas File System issues fixed in this release.

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

Incident	Description
1946128	The <code>fsadm(1M)</code> command <code>reorg</code> operation can take a long time to complete.
1959367	The <code>vxresize(1M)</code> command fails with the following error message:  <pre>UX:vxfs fsadm: ERROR: V-3-20295: device record checksum is invalid UX:vxfs fsadm: ERROR: V-3-20291: could not initialize devices</pre>
2026690	After more than one million files are created, any new operation to create a file using the <code>open (2)</code> system call fails with the following error value:  <pre>error value 22 (EINVAL)</pre>
2047396	On a multi-volume file system when the file system threshold is set, alerts are generated even before the set limit of the file system threshold is hit.

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

Incident	Description
2061973	An assert is hit when the <code>setext(1M)</code> command with the <code>-f growfile</code> option is executed. This occurs only when the file size is not file system block size aligned.
1935631	During a sequential read operation on an HP-UX 11iv3 system, the read-ahead performance is about 50% slower for an odd-size read request (which is greater than the <code>read_pref_io</code> parameter) compared to the read-ahead performance for an even-size page aligned read request.
1945168	When the <code>vxfs_bc_bufhwm(5)</code> tunable does not specify the upper limit of the tunable, and if the tunable is set to a very high value, for example greater than 2 to 3 percent of the file system size, the subsequent system reboot hangs.
2026514	In a Veritas File System (VxFS) environment with a large number of files in the root directory, if any <code>read(1M)</code> , <code>write(1M)</code> , <code>creat(2)</code> , or <code>open(2)</code> operation is performed, while the file unmount operation is in progress, the system may panic.
2026556	A create directory operation appears to be hung.
2026574	<p>The <code>fsadm(1M)</code> command fails to shrink a file system and displays the following errors:</p> <pre>UX:vxfs fsadm: ERROR: V-3-20340: attempt to resize &lt;file system&gt; failed with errno 16 UX:vxfs fsadm: ERROR: V-3-20343: cannot shrink &lt;file system&gt; -blocks are currently in use.</pre>
2026601	<p>Mounting a file system of size more than two terabytes, with the <code>quotas</code> option enabled, fails with the following error message:</p> <pre>UX:vxfs mount: WARNING: V-3-21407: cannot enable quotas</pre>
2026672	<p>While trying to start an Oracle instance, the following errors are displayed:</p> <pre>ORA-01110: data file 6: &lt;file name&gt; ORA-01115: IO error reading block from file 6 (block # 1) ORA-27070: async read/write failed HPUX-ia64 Error: 15: Block device required</pre>

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

Incident	Description
1995634	The <code>mount(1M)</code> command hangs while mounting a VxFS file system with the <code>quota</code> option.
2053272	On a stackable file system, in a race condition, flushing a file to the disk and processing the same file for removal can lead to a system panic.
2048018	<p>The <code>fsck(1M)</code> command dumps core due to the SIGBUS signal. The following error message is displayed:</p> <pre>log replay in progress 13054 Bus Error (coredump) Failed to full fsck cleanly, exiting</pre>
2026623	The <code>inod-ov</code> option of the <code>sar(1M)</code> utility is not available for the VxFS file systems.
2076199	In a scenario where thin arrays are used, when the <code>df(1M)</code> command is executed on Veritas File System (VxFS), the total available size of the file that is returned is greater than the actual size of the file system. Also, the <code>df(1M)</code> command returns a negative value for the <code>%Used</code> field.
2080395	The <code>vxfsstat(1M)</code> command does not reflect the change in the <code>vx_ninode(5)</code> tunable, after the tunable is changed using the <code>kctune(1M)</code> command.
2066039	The <code>mount(1M)</code> operation fails for a file system after a remount of the same file system with a changed <code>logiosize</code> .
2029930	The relocation or the placement of the files in the directories may not occur correctly, after the old policy is unassigned and the new policy is assigned.
2066043	Mounting a file system and then remounting the file system with a different <code>logiosize</code> value, may cause structural damage to the file system and make it unmountable.
2080683	When a hole is created in the file using Data Management Application Programming Interface (DMAPI), the <code>dm_punch_hole()</code> function can leave the file in a corrupted state.
2072119	An <code>umount(1M)</code> operation hangs on a file system that has the read-write clone.

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

Incident	Description
2088083	The <code>getacl(1M)</code> command incorrectly reports the file system information.

[Table 1-2](#) lists the Veritas Volume Manager issues fixed in this release.

**Table 1-2** Veritas Volume Manager fixed issues

Incident	Description
2059239	Veritas Volume Manager (VxVM) corrupts the address space of Veritas File system (VxFS). This causes the system to panic in the VxFS function.
2059300	In a Cluster Volume Manager (CVM) environment, the <code>split</code> operation of the <code>vx dg(1M)</code> command fails with the following error message:  Internal configuration daemon error
2060306	The <code>addmir</code> operation of the <code>vx snap(1M)</code> command dumps core.
2071596	When the exclude array operation is performed with the EMC Clariion PowerPath, the system panics.
2020060	For a scenario where the disks are of size larger than 2 terabytes, the <code>vx disksetup(1M)</code> command fails with the following error message:  Disk is too small for supplied parameters
1670279	I/O hangs on a volume when the volume mirror is added to the volume for snapshot purpose using <code>vx snap -g &lt;dg&gt; addmir mirvol=&lt;vol&gt;</code> .
1670291	In an EMC Symmetrix Remote Data Facility (SRDF) environment, there is a failover from Server1 (which contains the R1 devices) to Server2 (which contains the R2 devices). After the failover, any further updates of the Unique Disk Identification (UDID) on the devices in the disk group on Server 2, corrupt the disk group configuration copies. The following error message is displayed when the disk group is deported on Server2 and imported back on Server1 after a failback:  VxVM vx dg ERROR V-5-1-10978 Disk group <disk group name>: import failed



**Table 1-2** Veritas Volume Manager fixed issues (*continued*)

Incident	Description
1925842	In a CVM environment, the master node panics during a node join process or when a shared disk group is imported.
1937726	The user must not be able to set the <code>vol_maxspecialio</code> tunable value higher than the I/O size supported by the operating system.
2049163	<p>When the <code>vxprint(1M)</code> command reads the configuration information from the standard input, it fails with the following error message:</p> <pre>VxVM vxprint ERROR V-5-1-326 unrecognized value for variable, context: activation=shared-write</pre>
2052521	<p>The fix for JAGaf60923 introduced the problem importing the disk groups with <code>vol_maxspecialio</code> and/or <code>vol_maxio</code> tunable value being smaller than 1024 or if an underlying disk driver does not support I/O size as large as 1MB. The typical error observed during the disk group import is:</p> <pre>VxVM vxconfigd ERROR V-5-1-569 Disk group rootdg, Disk [disk]: Cannot auto-import group: Disk group has no valid configuration copies</pre>
2052524	When the volume recovery operation I/O statistics are checked, the volume recovery operation I/O occurs in chunks of 64 KB, despite of the volume recovery operation I/O size being set to 1 MB.
2052527	<p>In a CVM environment, the system panics with the message:</p> <pre>Can't sleep in interrupt context</pre>
2059155	The <code>vxbrk_rootmir</code> script does not complete executing and stops responding.
1670281	<p>In the Campus Cluster environment, the “reattach” operation of the site may fail with the following error:</p> <pre>VxVM vxdg WARNING V-5-1-13835 Disk List is empty for site &lt;site name&gt;</pre>
2059119	The primary host of Veritas Volume Replicator (VVR) panics during the automatic synchronization or the DCM (Data Change Map) replay.

**Table 1-2** Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2059125	<p>The operation to join a disk group fails with the following error message:</p> <pre>VxVM vxdg ERROR V-5-1-4597 vxdg join &lt;dg1&gt; &lt;dg2&gt; failed join failed : Configuration too large for configuration copies</pre>
2059170	<p>On a VxVM volume, when a snapshot is taken or when a snapshot is reattached, or when any other snapshot operations are performed, the system panics.</p>
2060303	<p>In the campus cluster environment, the sites or plexes may not be reattached automatically when the faulted disks are reconnected.</p>
2071515	<p>In the campus cluster environment, all the VxVM sites may enter a DETACHED state in the disk group.</p>
2039118	<p>Memory may get exhausted due to the allocations made by the vxio driver. The memory leak occurs when smartsync is enabled on volumes under Veritas Volume Replicator (VVR).</p>
2056317	<p>Depending on the workload on the system, the read or write I/O performance degrades after VxVM is upgraded from version 5.0 to version 5.0.1.</p>
1670283	<p>In a campus cluster environment, the disk group import operation fails with the following error message:</p> <pre>VxVM vxdg ERROR V-5-1-XXXXX Disk group &lt;dgname&gt;: import failed: Required lock not held in transaction</pre>

**Table 1-2** Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2059169	<p>In a Cluster Volume Replicator (CVR) environment, the <code>vradmind</code> daemon dumps core. There are three different core dumps that occur. The three different scenarios under which the command dumps core are as follows:</p> <p>In the first scenario, the <code>vradmind</code> daemon dumps core when the log owner node is switched frequently.</p> <p>In the second scenario, when the user runs multiple <code>vrstat(1M)</code> commands in parallel, the <code>vxadmind</code> daemon dumps core when the <code>vrstat(1M)</code> command exits.</p> <p>In the third scenario, the operation to create and delete Replicator Volume Group (RVG) repeatedly along with multiple <code>vrstat -n 2</code> commands causes the <code>vradmind</code> daemon to dump core.</p>
2060312	<p>In a CVM environment, the master node loses access to all the disks in the shared disk group and the disk group <code>fail_policy</code> is set to <code>leave</code>. As a result, the CVM master node panics and leaves the cluster, and a new master node takes over. For the site-consistent disk groups, the new master node detaches one or more plexes and one of the campus cluster sites randomly.</p>
2071517	<p>In a campus cluster environment, when a site that has the CVM master node loses the site storage, the I/O hangs on all the nodes.</p>
2071582	<p>In a Veritas Cluster Server (VCS) environment with Cluster Volume Manger (CVM), the <code>vxclustadm -m vcs -t gab startnode</code> command results in a segmentation fault. The <code>vxclustadm(1M)</code> command dumps core.</p>
1913450	<p>When the Dynamic Multipathing (DMP) I/O policy is changed, the <code>vxconfigd(1M)</code> daemon hangs. This also causes the other VxVM commands to hang.</p>
1956960	<p>The volumes that are created on the thin provisioned Logical Unit Numbers (LUNs) are not reclaimed using the <code>fsadm -R</code> or the <code>vxdisk reclaim</code> commands.</p>

**Table 1-2** Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2059303	<p>I/O errors are observed on the DMP node and the disks are marked as failing. In the DMP events log /etc/vx/dmpevents.log file, the following messages are seen:</p> <pre>I/O error occured on Path &lt;c#t#d#&gt; belonging to Dmpnode &lt;dmp_nodename&gt; I/O analysis done on Path &lt;c#t#d#&gt; belonging to Dmpnode &lt;dmp_nodename&gt;</pre>
2013736	<p>When the coexistence of Veritas Volume Manager (VxVM) 5.0.1 with EMC PowerPath (PP) 5.1.2 is tested, EMC PP 5.1.2 encounters a panic. The disks are claimed by the libvxCLARion.sl Array Support Library (ASL) instead of being claimed by the libvxpp.sl ASL.</p>
2075323	<p>The vxconfigd(1M) daemon may dump core when the JBOD definition is added or removed.</p>
2071587	<p>When a LUN is replaced , the disk goes into the error state during the dynamic reconfiguration operation .</p>
2046497	<p>In a CVM environment, the VxVM commands may take a long time to complete. The following messages are displayed in the system log on the master node when the issue occurs:</p> <pre>&lt;date &amp; time&gt; &lt;system name&gt; syslog: advresp_master: GAB returned EAGAIN, retrying &lt;date &amp; time&gt; &lt;system name&gt; syslog: advresp_master: Retry Successful</pre>
2059166	<p>In a 4-node CVR setup, I/O hangs.</p>
2059168	<p>When a break- off style snapshot of the original volume (which has a large number of volumes) is used and the snapshot volume is moved to another group for off-host processing and again the snapshot volume is moved to the original disk group and refreshed, I/O hangs. This causes VCS to panic. I/O is stuck in volsioq_restart.</p>
2059167	<p>Followed by a system crash, the I/O to read or write to a volume that has associated space optimized snapshots may hang. The system may also panic and unexpected read or write errors may get displayed. The following error message may get displayed:</p> <pre>VxVM vxio V-5-0-1266 Volume snap1 block 324: Uncorrectable read error</pre>

**Table 1-2** Veritas Volume Manager fixed issues (*continued*)

Incident	Description
2092315	In a campus cluster environment, the disabled volumes start automatically.

For details about the issues fixed in the rollup patches, go to the following Web page on VOS:

<https://vos.symantec.com/patch/detail/3452>

## Veritas Storage Foundation for Oracle fixed issues

[Table 1-3](#) lists the Veritas Storage Foundation for Oracle issues fixed in this release.

**Table 1-3** Veritas Storage Foundation for Oracle fixed issues

Incident	Description
2063215	<p>The DBED commands do not work with Oracle Database 11g Release 2 on the PA-RISC-based systems. For example, the <code>dbed_update(1M)</code> command fails and generates the following error messages:</p> <pre> /usr/lib/dld.sl: Bad magic number for shared library: /orabinvol/oracle/lib/libclntsh.sl /usr/lib/dld.sl: Exec format error /opt/VRTS/bin/dbed_update[767]: 27590 Abort(coredump) SFORA dbed_update ERROR V-81-8828 An error occurred while updating the repository for &lt;database&gt; </pre>
2081834	<p>On Oracle Database 11g R2, the <code>dbed_ckptrollback -F datafile</code> command fails with the following error:</p> <pre> SFORA rb.file ERROR V-81-3038 Error occurred while querying Oracle Database. ORA-01135: file 1015 accessed for DML/query is offline ORA-01110: data file 1015: '/snap_data11r2/FLAS11r2/MYTBS.dbf' </pre>

[Table 1-4](#) lists the Veritas Storage Foundation Common Utilities for Databases issue fixed in this release.

**Table 1-4** Veritas Storage Foundation Common Utilities for Databases fixed issue

Incident	Description
2088331	<p>The <code>dbdst_file_move(1M)</code> command fails and displays the following error:</p> <pre>SFORA dbdst_file_move ERROR V-81-6244 Could not execute tsdb_oraquery.</pre>

For details about the issues fixed in the rollup patches, go to the following Web page on VOS:

<https://vos.symantec.com/patch/detail/3452>

## VCS fixed issues

[Table 1-5](#) lists the issues fixed in 5.0.1 Rolling Patch 2.

**Table 1-5** VCS 5.0.1 Rolling Patch 2 fixed issues

Incident	Description
2052116	<p>The <code>hastart</code> command fails to start VCS on one of the nodes with the following error message in the syslog :</p> <pre>GabHandle::open failed errno = 16</pre>
2052118	<p>Cascaded failing of node with lowest LLT-node IDs may cause GAB panic during sequence space recovery.</p>
2052121	<p>Joining of a node with lowest node id at the time an existing cluster node is moving out of the cluster, may cause another node to panic with the following panic string :</p> <pre>Port f halting system due to network failure</pre>

## Veritas Storage Foundation Cluster File System fixed issues

[Table 1-6](#) lists the Veritas Storage Foundation Cluster File System issues fixed in this release.

**Table 1-6** Veritas Storage Foundation Cluster File System fixed issues

Incident	Description
1935629	In a Cluster File System (CFS) environment, the <code>df(1M)</code> command and the <code>rm(1M)</code> command get hung.
1935637	In a Cluster File System (CFS) environment, file system operations can get into an endless loop after state map corruptions are reported.
1959375	The <code>rm(1M)</code> operation to delete a large number of files is about seven times slower on the Cluster File System (CFS) environment as compared to the same operation on the local mount.
2081510	In a Cluster File System (CFS) environment, the startup operation of CFS hangs.
1984249	<p>While trying to modify the attributes of the CFSMount resources using the <code>hares(1M)</code> command, the following message gets logged in the Veritas Cluster Server (VCS) engine log:</p> <pre>hares cmd (write) failed</pre> <p>However, the command is executed successfully.</p>
1989995	<p>A CFS file system cannot be mounted on a secondary cluster and the following error message gets logged in the VCS engine log:</p> <pre>(Date &amp; Time) VCS ERROR V-16-20007-1037 (s117072) CVMVolDg:cvmvoldg1:monitor:Recheck of status check failed (Date &amp; Time) VCS ERROR V-16-20007-1052 (s117071) CVMVolDg:cvmvoldg1:monitor:Volume (Vol_Num) is in an RVG. Can not monitor volumes which are in an RVG</pre>
2039842	<p>The following information gets logged in the VCS engine log:</p> <pre>(Date &amp; Time) VCS INFO V-16-2-13716 (Node_Name) Resource (cvmvoldg2) : Output of the completed operation (monitor) ===== /opt/VRTSvcs/bin/CVMVolDg/monitor[33]: -: not found =====</pre> <p>However, this issue does not affect any functionality.</p>

For details about the issues fixed in the rollup patches, go to the following Web page on VOS:

<https://vos.symantec.com/patch/detail/3452>

## SF Oracle RAC fixed issues

[Table 1-7](#) lists the issues fixed in this release.

**Table 1-7** SF Oracle RAC fixed issues

Incident	Description
1877596	Low Latency Transport Multiplexer (LMX) may cause the system to panic with the following message:  <code>kernel heap corruption detected</code>
1931590	When the system starts, the following error message may be displayed on the console and in the <code>/var/adm/syslog/syslog.log</code> file:  <code>mmp1_reconfig_ioctl: dev_ioctl failed</code>
1931604	The PrivNIC agent logs error messages in the <code>/var/VRTSvcs/log/engine_A.log</code> file if non existent interfaces are specified while configuring the PrivNIC resource.

## Known issues in this release

This section describes the known issues in this release.

- Veritas Storage Foundation:  
See [“Veritas Storage Foundation 5.0.1 Rolling Patch 2 known issues”](#) on page 24.
- Veritas Storage Foundation for Oracle:  
See [“Veritas Storage Foundation for Oracle 5.0.1 Rolling Patch 2 known issues”](#) on page 26.
- Veritas Storage Foundation for Oracle RAC:  
See [“SF Oracle RAC 5.0.1 Rolling Patch 2 known issues”](#) on page 28.

## Veritas Storage Foundation 5.0.1 Rolling Patch 2 known issues

The Veritas Storage Foundation known issues in the 5.0.1 release are listed in *Veritas Storage Foundation Release Notes (Version 5.0.1)*.

The Veritas Storage Foundation known issues in the 5.0.1 Rolling Patch 1 release are listed in *Veritas Storage Foundation and High Availability Solutions Release Notes (Version 5.0.1 RP1)*.



The Veritas Storage Foundation known issues in this release are as follows:

- File Change Log tunable setting for proper functioning of Dynamic Storage Tiering (DST) applications:

If the active placement policy of a given file system uses I/O or access temperatures, after the policy becomes active by being assigned, you must tune the file system's `fcl_malloc` tunable with the following command:

```
# vxtunefs -o fcl_maxalloc=0 mount_point
```

However, if any applications other than DST use FCL, this setting may conflict with those applications.

- Accessing the files through Oracle Disk Manager (ODM) may not work (1842547):  
Accessing the files through ODM in the serially-exclusive mode does not work. ODM access is only available for files that are not simultaneously accessed from multiple nodes.

- File system checkpoint mount unlock operation using the `fsadm(1M)` command fails (2087260):

A file system checkpoint mount unlock operation using the `fsadm(1M)` command fails with the following error message:

```
fsadm -F vxfs -o mntunlock=<mnt lock id> <mntpt>  
UX:vxfs fsadm: ERROR: V-3-20000: file system not mount locked
```

Workaround:

Unmount the file system using the `vxumount(1M)` command with the `mntunlock` option.

- Unmount operation on a file system checkpoint mounted as read-only fails (2091311):

A `umount(1M)` operation on any file system checkpoint mounted as read-only fails with the following message:

```
umount: cannot unmount <device path> : Device busy  
umount: return error 1
```

The mount tab file for the read-only file systems is not updated with the mount lock ID. As a result, all the subsequent unmount operations fail in unmounting the file system.

Workaround:

Remount the file system as read-write. You can now find the mount lock ID in the mount options. Then, unmount the file system using the `vxumount(1M)` command with the `mntunlock` option.

- The VRTSperl patch takes more than 10 minutes to install on an HP Integrity system node:  
On an HP Integrity system node, installing the VRTSperl patch takes more than 10 minutes and requires that VCS is offline during this period. The installation time may vary based on the configuration of the machine on which the VRTSperl patch is being installed.
- Number of paths for disks not showing in the Storage Foundation Manager (SFM) dashboard:  
In SFM 2.1 Rolling Patch 1, if the disk naming scheme is changed to 'new' OS based naming, the SFM dashboard does not show the number of paths for the respective disks.  
Workaround:  
To display the number of paths correctly, change the disk naming scheme to enclosure based naming or 'default' OS based naming by using the following commands:  
Change the disk naming scheme to enclosure based naming: **vxddladm set namingscheme=ebn**  
Change the disk naming scheme to 'default' OS based naming: **vxddladm set namingscheme=osn mode=default**
- Some disks of a site are not detached when the site is detached (2092869):  
In a campus cluster setup, when a slave loses all storage (complete storage failure), some of the disks might fail to get detached. Also, the disks might remain in the FAILING state. This occurs because the transaction to detach a disk would fail when the site has already been detached.  
This issue does not affect any functionality because the plexes that belong to the disks are already detached from the kernel.  
Workaround:  
When the disks come back up, use the following command to reattach the site:  

```
vxdg reattachsite <site>
```

  
To clear the FAILING flag on the disks, use the following command:  

```
vxedit -g <dg> set failing=(off/false) <disk>
```

## Veritas Storage Foundation for Oracle 5.0.1 Rolling Patch 2 known issues

The Veritas Storage Foundation for Oracle known issues in this release are as follows:

- Errors in the swremove log of the VRTSdbms3 patch (PHCO\_40672):

Users may observe errors in the swremove log of the VRTSdbms3 patch, if the `base_resize` parameter is set to a non-default value and the patch is removed. This occurs because after the removal of this patch, the issue fixed by this patch resurfaces.

- The `dbed_analyzer` utility fails on the IA based systems:  
The `dbed_analyzer` utility fails with the following error on the IA based systems:

```
SFORA dbed_analyzer ERROR V-81-5900 Operation not supported.
```

**Workaround:**

There is no workaround for this issue. The `dbed_analyzer` utility is deprecated and it is no longer supported.

- `dbed_clonedb -o restartdb` fails after database group switches to a second node with spfile (1721965):  
In Veritas Storage Foundation for Oracle High Availability, the `restartdb` command for a checkpoint clone database can be run only on the node on which you have created the clone.

## VCS 5.0.1 Rolling Patch 2 known issues

The VCS issues in this release are as follows:

- The Oracle agent with 11g Release 2 does not support Health check monitoring using the `MonitorOption` attribute. If the database is 11g Release 2, the `MonitorOption` attribute for the Oracle agent should be set to 0.  
The Oracle agent with 11g Release 2 database does not support the Intentional Offline feature. [1975007]
- The `ASMInst` agent does not support pfile or spfile for the ASM Instance on the ASM diskgroups in 11g Release 2. Symantec recommends that you store the file on the local file system. [1975010]
- If you try to enable debug logs for the DB2 agent, the logs are not written to the `engine_A`. log file. [1954752]  
Workaround: Download and install the GNU Awk software from the GNU Web site. Then, create a soft link to the default awk binary on the cluster nodes as follows:

```
# ln -s /usr/local/bin/gawk /bin/awk
```
- The `VRTSperl` patch takes more than 10 minutes to install on an HP Integrity system node:

On an HP Integrity system node, installing the VRTSperl patch takes more than 10 minutes and requires that VCS is offline during this period. The installation time may vary based on the configuration of the machine on which the VRTSperl patch is being installed.

## SF Oracle RAC 5.0.1 Rolling Patch 2 known issues

There are no known issues in SF Oracle RAC 5.0.1 Rolling Patch 2. For known issues in version 5.0.1, see the *Veritas Storage Foundation for Oracle RAC Release Notes (5.0.1)*.

## Downloading the patches

The patches included in Veritas Storage Foundation and High Availability Solutions 5.0.1 Rolling Patch 2 are available for download from the Symantec website. After downloading the file, use gunzip and tar to uncompress and extract.

For the Rolling Patch 2 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/331560>

## List of patches

This section lists the patches included in this release.

- Veritas Storage Foundation:  
See “[Veritas Storage Foundation patches in 5.0.1 Rolling Patch 2](#)” on page 29.
- Veritas Storage Foundation for Oracle:  
See “[Veritas Storage Foundation for Oracle patches in 5.0.1 Rolling Patch 2](#)” on page 30.
- Veritas Cluster Server:  
See “[VCS patches in 5.0.1 Rolling Patch 2](#)” on page 31.
- Veritas Storage Foundation Cluster File System:  
See “[Veritas Storage Foundation Cluster File System patches in 5.0.1 Rolling Patch 2](#)” on page 32.
- Veritas Storage Foundation for Oracle RAC  
See “[Veritas Storage Foundation for Oracle RAC patches in 5.0.1 Rolling Patch 2](#)” on page 34.

## Veritas Storage Foundation patches in 5.0.1 Rolling Patch 2

[Table 1-8](#) lists the Veritas Storage Foundation patches included in this release.

**Table 1-8** Veritas Storage Foundation patches

Patch	Version	Description
PHCO_41194	1.0	VRTS 5.0.1 GARP2 VRTSvxfs Command Patch (Veritas File System )
PHKL_41074	1.0	VRTS 5.0.1 GARP2 VRTSvxfs Kernel Patch (Veritas File System )
PHCO_41192	1.0	VRTS 5.0.1 GARP2 VRTSvxvm Command Patch (Veritas Volume Manager)
PHKL_41193	1.0	VRTS 5.0.1 GARP2 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PHCO_40771	1.0	VRTS 5.0.1 GARP1 VRTSfspro Command Patch (Veritas File System Management Services Provider) (Rollup patch)
PHCO_40770	1.0	VRTS 5.0.1 GARP1 VRTSfsman Command Patch (Veritas File System Manual Pages ) (Rollup patch)
PHCO_40520	1.0	VRTS 5.0.1 GARP1 VRTSweb Patch (Symantec Web Server) (Rollup patch)
PHCO_39783	1.0	VRTS 5.0.1 GARP1 VRTSsfmh Command Patch (Veritas Storage Foundation Manager) (Rollup patch)
PHCO_40673	1.0	VRTS 5.0.1 GARP1 VRTSvmpro Command Patch (Veritas Volume Manager Management Services Provider) (Rollup patch)
PHCO_40674	1.0	VRTS 5.0.1 GARP1 VRTSddlpr Command Patch (Veritas Device Discovery Layer services Provider) (Rollup patch)

**Table 1-8** Veritas Storage Foundation patches (*continued*)

Patch	Version	Description
PHCO_40782	5.8.8.10	VRTS 5.0.1 GARP1 VRTSperl Command Patch (Perl Redistribution )  (Rollup patch)

## Veritas Storage Foundation for Oracle patches in 5.0.1 Rolling Patch 2

[Table 1-9](#) lists the Veritas Storage Foundation for Oracle patches included in this release.

**Table 1-9** Veritas Storage Foundation for Oracle patches

Patch	Version	Description
PHCO_41221	1.0	VRTS 5.0.1 GARP2 VRTSdbed Command Patch (Veritas Storage Foundation for Oracle )
PHCO_41321	1.0	VRTS 5.0.1 GARP2 VRTSdbcom Command Patch (Veritas Storage Foundation Common Utilities for Databases )
PHCO_41194	1.0	VRTS 5.0.1 GARP2 VRTSvxfs Command Patch (Veritas File System )
PHKL_41074	1.0	VRTS 5.0.1 GARP2 VRTSvxfs Kernel Patch (Veritas File System )
PHCO_41192	1.0	VRTS 5.0.1 GARP2 VRTSvxvm Command Patch (Veritas Volume Manager)
PHKL_41193	1.0	VRTS 5.0.1 GARP2 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PHCO_40771	1.0	VRTS 5.0.1 GARP1 VRTSfspro Command Patch (Veritas File System Management Services Provider)  (Rollup patch)
PHCO_40770	1.0	VRTS 5.0.1 GARP1 VRTSfsman Command Patch (Veritas File System Manual Pages )  (Rollup patch)

**Table 1-9** Veritas Storage Foundation for Oracle patches (*continued*)

Patch	Version	Description
PHCO_40520	1.0	VRTS 5.0.1 GARP1 VRTSweb Patch (Symantec Web Server) (Rollup patch)
PHCO_39783	1.0	VRTS 5.0.1 GARP1 VRTSsfmh Command Patch (Veritas Storage Foundation Manager) (Rollup patch)
PHCO_40672	1.0	VRTS 5.0.1 GARP1 VRTSdbms3 Command Patch (Veritas Shared DBMS) (Rollup patch)
PHCO_40673	1.0	VRTS 5.0.1 GARP1 VRTSvmpro Command Patch (Veritas Volume Manager Management Services Provider) (Rollup patch)
PHCO_40674	1.0	VRTS 5.0.1 GARP1 VRTSddlpr Command Patch (Veritas Device Discovery Layer services Provider ) (Rollup patch)
PHCO_40782	5.8.8.10	VRTS 5.0.1 GARP1 VRTSperl Command Patch (Perl Redistribution ) (Rollup patch)

## VCS patches in 5.0.1 Rolling Patch 2

[Table 1-10](#) lists the VCS patches in this release.

**Table 1-10** VCS patches

Patch	Version	Description
PHNE_41196	1.0	VRTS 5.0.1 GARP2 VRTSlt Kernel Patch (Veritas Low Latency Transport)
PHNE_41197	1.0	VRTS 5.0.1 GARP2 VRTSgab Kernel Patch (Veritas Group Membership and Atomic Broadcast)

**Table 1-10** VCS patches (*continued*)

Patch	Version	Description
PHCO_40690	1.0	VRTS 5.0.1 GARP1 VRTSvxfen Command Patch (Veritas I/O Fencing) (Rollup patch)
PHKL_40443	1.0	VRTS 5.0.1 GARP1 VRTSvxfen Kernel Patch (Veritas I/O Fencing) (Rollup patch)
PVCO_03905	1.0	VRTS 5.0.1 GARP1 VRTSvcs Command Patch (Veritas Cluster Server) (Rollup patch)
PVCO_03906	1.0	VRTS 5.0.1 GARP1 VRTSvcsag Command Patch (Veritas Cluster Server Bundled Agents) (Rollup patch)
PVCO_03907	1.0	VRTS 5.0.1 GARP1 VRTSvcsor Command Patch (Veritas High Availability Agent for Oracle) (Rollup patch)
PHCO_40520	1.0	VRTS 5.0.1GARP1VRTSweb Patch (SymantecWeb Server) (Rollup patch)
PHCO_40782	5.8.8.10	VRTS 5.0.1 GARP1 VRTSperl Command Patch (Perl 5.8.8 Redistribution) (Rollup patch)

## Veritas Storage Foundation Cluster File System patches in 5.0.1 Rolling Patch 2

[Table 1-11](#) lists the Veritas Storage Foundation Cluster File System patches included in this release.

**Table 1-11** Veritas Storage Foundation Cluster File System patches

Patch	Version	Description
PVCO_03916	1.0	VRTS 5.0.1 GARP2 VRTScavf Command Patch (Veritas Cluster Server Agents for Cluster File System)



**Table 1-11** Veritas Storage Foundation Cluster File System patches (*continued*)

Patch	Version	Description
PHNE_41196	1.0	VRTS 5.0.1 GARP2 VRTSllt Kernel Patch (Veritas Low Latency Transport )
PHNE_41197	1.0	VRTS 5.0.1 GARP2 VRTSgab Kernel Patch (Veritas Group Membership and Atomic Broadcast )
PHCO_41194	1.0	VRTS 5.0.1 GARP2 VRTSvxfs Command Patch (Veritas File System )
PHKL_41074	1.0	VRTS 5.0.1 GARP2 VRTSvxfs Kernel Patch (Veritas File System )
PHCO_41192	1.0	VRTS 5.0.1 GARP2 VRTSvxvm Command Patch (Veritas Volume Manager)
PHKL_41193	1.0	VRTS 5.0.1 GARP2 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PHKL_40772	1.0	VRTS 5.0.1 GARP1 VRTSglm Kernel Patch (Veritas Group Lock Manager) (Rollup patch)
PHCO_40771	1.0	VRTS 5.0.1 GARP1 VRTSfspro Command Patch (Veritas File System Management Services Provider) (Rollup patch)
PHCO_40770	1.0	VRTS 5.0.1 GARP1 VRTSfsman Command Patch (Veritas File System Manual Pages ) (Rollup patch)
PHKL_40443	1.0	VRTS 5.0.1 GARP1 VRTSvxfen Kernel Patch (Veritas I/O Fencing) (Rollup patch)
PHCO_40690	1.0	VRTS 5.0.1 GARP1 VRTSvxfen Command Patch (Veritas I/O Fencing) (Rollup patch)
PVCO_03905	1.0	VRTS 5.0.1 GARP1 VRTSvcs Command Patch (Veritas Cluster Server) (Rollup patch)

**Table 1-11** Veritas Storage Foundation Cluster File System patches (*continued*)

Patch	Version	Description
PVCO_03906	1.0	VRTS 5.0.1 GARP1 VRTSvcsag Command Patch (Veritas Cluster Server Bundled Agents) (Rollup patch)
PHCO_40520	1.0	VRTS 5.0.1 GARP1 VRTSweb Patch (Symantec Web Server) (Rollup patch)
PHCO_39783	1.0	VRTS 5.0.1 GARP1 VRTSsfmh Command Patch (Veritas Storage Foundation Manager) (Rollup patch)
PHCO_40673	1.0	VRTS 5.0.1 GARP1 VRTSvmpro Command Patch (Veritas Volume Manager Management Services Provider) (Rollup patch)
PHCO_40674	1.0	VRTS 5.0.1 GARP1 VRTSddlpr Command Patch (Veritas Device Discovery Layer services Provider ) (Rollup patch)
PHCO_40782	5.8.8.10	VRTS 5.0.1 GARP1 VRTSperl Command Patch (Perl Redistribution ) (Rollup patch)

## Veritas Storage Foundation for Oracle RAC patches in 5.0.1 Rolling Patch 2

[Table 1-12](#) lists the Veritas Storage Foundation for Oracle RAC patches included in this release.

**Table 1-12** Veritas Storage Foundation for Oracle RAC patches

Patch	Version	Description
PVKL_03904	1.0	VRTS 5.0.1 GARP1 VRTSdbac Kernel Patch (Veritas Oracle Real Application Clusters Support Package)

**Table 1-12** Veritas Storage Foundation for Oracle RAC patches (*continued*)

Patch	Version	Description
PVCO_03916	1.0	VRTS 5.0.1 GARP2 VRTScavf Command Patch (Veritas Cluster Server Agents for Cluster File System)
PHCO_40672	1.0	VRTS 5.0.1 GARP1 Symantec Shared DBMS Command Patch
PHNE_41196	1.0	VRTS 5.0.1 GARP2 VRTSllt Kernel Patch (Veritas Low Latency Transport )
PHNE_41197	1.0	VRTS 5.0.1 GARP2 VRTSgab Kernel Patch (Veritas Group Membership and Atomic Broadcast )
PHCO_41194	1.0	VRTS 5.0.1 GARP2 VRTSvxfs Command Patch (Veritas File System )
PHKL_41074	1.0	VRTS 5.0.1 GARP2 VRTSvxfs Kernel Patch (Veritas File System )
PHCO_41192	1.0	VRTS 5.0.1 GARP2 VRTSvxvm Command Patch (Veritas Volume Manager)
PHKL_41193	1.0	VRTS 5.0.1 GARP2 VRTSvxvm Kernel Patch (Veritas Volume Manager)
PHCO_41221	1.0	VRTS5.0.1GARP2VRTSdbedCommandPatch (Veritas Storage Foundation for Oracle)
PHCO_41321	1.0	VRTS 5.0.1 GARP2 VRTSdbcom Command Patch (Veritas Storage Foundation Common Utilities for Databases)
PHKL_40772	1.0	VRTS 5.0.1 GARP1 VRTSglm Kernel Patch (Veritas Group Lock Manager) (Rollup patch)
PHCO_40771	1.0	VRTS 5.0.1 GARP1 VRTSfspro Command Patch (Veritas File System Management Services Provider) (Rollup patch)
PHCO_40770	1.0	VRTS 5.0.1 GARP1 VRTSfsman Command Patch (Veritas File System Manual Pages ) (Rollup patch)

**Table 1-12** Veritas Storage Foundation for Oracle RAC patches *(continued)*

Patch	Version	Description
PHKL_40443	1.0	VRTS 5.0.1 GARP1 VRTSvxfen Kernel Patch (Veritas I/O Fencing) (Rollup patch)
PHCO_40690	1.0	VRTS 5.0.1 GARP1 VRTSvxfen Command Patch (Veritas I/O Fencing) (Rollup patch)
PVCO_03905	1.0	VRTS 5.0.1 GARP1 VRTSvcs Command Patch (Veritas Cluster Server) (Rollup patch)
PVCO_03906	1.0	VRTS 5.0.1 GARP1 VRTSvcsag Command Patch (Veritas Cluster Server Bundled Agents) (Rollup patch)
PHCO_40520	1.0	VRTS 5.0.1 GARP1 VRTSweb Patch (Symantec Web Server) (Rollup patch)
PHCO_39783	1.0	VRTS 5.0.1 GARP1 VRTSsfmh Command Patch (Veritas Storage Foundation Manager) (Rollup patch)
PHCO_40673	1.0	VRTS 5.0.1 GARP1 VRTSvmpro Command Patch (Veritas Volume Manager Management Services Provider) (Rollup patch)
PHCO_40674	1.0	VRTS 5.0.1 GARP1 VRTSddlpr Command Patch (Veritas Device Discovery Layer services Provider ) (Rollup patch)
PHCO_40782	5.8.8.10	VRTS 5.0.1 GARP1 VRTSperl Command Patch (Perl Redistribution ) (Rollup patch)

# Installing the patches

You can install Storage Foundation and High Availability Solutions 5.0.1 Rolling Patch 2 by using the `installrp` script.

Depending on the product you want to install, perform the steps in one of the following sections:

- Veritas Storage Foundation:  
See [“Installing Veritas Storage Foundation 5.0.1 Rolling Patch 2”](#) on page 37.
- Veritas Storage Foundation High Availability:  
See [“Installing Veritas Storage Foundation High Availability 5.0.1 Rolling Patch 2”](#) on page 38.
- Veritas Cluster Server:  
See [“Installing VCS 5.0.1 Rolling Patch 2”](#) on page 38.
- Veritas Storage Foundation Cluster File System:  
See [“Installing Veritas Storage Foundation Cluster File System 5.0.1 Rolling Patch 2”](#) on page 51.
- Veritas Storage Foundation for Oracle RAC  
See [“Installing SF Oracle RAC 5.0.1 Rolling Patch 2”](#) on page 60.

## Installing Veritas Storage Foundation 5.0.1 Rolling Patch 2

The following sample procedure is based on installing Veritas Storage Foundation 5.0.1 Rolling Patch 2 on a single system.

---

**Note:** The same procedure can also be used to install Veritas Storage Foundation for Oracle 5.0.1 Rolling Patch 2.

---

### To install Veritas Storage Foundation 5.0.1 Rolling Patch 2

- 1 Mount the software disc.
- 2 Move to the top-level directory on the disc.

- 3 From this directory, type the following command to install on the local system only. Also use this command to install on remote systems using the secure shell (ssh) utilities:

```
# ./installrp
```

The sample installation assumes that you are using ssh. However, if you use the remote shell utilities to install on remote systems, additionally specify the `-rsh` option:

```
# ./installrp -rsh
```

- 4 At the prompt, enter the system names on which the Rolling Patch 2 is to be installed.

```
Enter the system names separated by spaces on which to  
install 5.0.1 RP2: host1
```

- 5 After the list of patches is displayed, press Return to continue with installation.
- 6 After installation is complete, reboot the system by using the following command:

```
# shutdown -r now
```

If any patches fail to install, manually install them to complete the installation.

- 7 Verify the installation of the patches by using the following commands:

```
# swlist | grep -i garp
```

```
# swverify [patch names]
```

## Installing Veritas Storage Foundation High Availability 5.0.1 Rolling Patch 2

You can add high availability functionality to Veritas Storage Foundation High Availability (SFHA) by installing Veritas Cluster Server (VCS). For information on installing VCS 5.0.1 Rolling Patch 2:

See [“Installing VCS 5.0.1 Rolling Patch 2”](#) on page 38.

## Installing VCS 5.0.1 Rolling Patch 2

You can perform the installation in one of the following ways:

- Full installation  
See “[VCS 5.0.1 Rolling Patch 2 full installation](#)” on page 39.
- Phased installation  
See “[VCS 5.0.1 Rolling Patch 2 phased installation](#)” on page 42.

For single-instance Oracle 11g Release 2 databases with ASM, set the Oracle Cluster Synchronization Services tunable.

## VCS 5.0.1 Rolling Patch 2 full installation

The following sample procedure is based on installing VCS 5.0.1 Rolling Patch 2 on the cluster nodes.

If you are applying this patch to a Veritas suite of products, such as SFRAC that includes VCS as a component, refer to the patch installation procedures for that product.

### To install VCS 5.0.1 Rolling Patch 2

- 1 Backup the llttab, llthosts, gabtab, and main.cf files.

```
# cp /etc/llttab /etc/llttab.bkp
# cp /etc/llthosts /etc/llthosts.bkp
# cp /etc/gabtab /etc/gabtab.bkp
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.bkp
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/config/types.cf.bkp
# cp /etc/VRTSvcs/conf/types.cf \
/etc/VRTSvcs/conf/types.cf.bkp
```

- 2 If the ClusterService group is configured, take it offline. You can run the `# hagr - state` command to verify whether the service group is configured or not.

```
# hagr -offline -force ClusterService -sys system
```

---

**Note:** The ClusterService group cannot be frozen.

---

- 3 Freeze all the service groups persistently, on all nodes.

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

where `service_group` is the name of the service group.

- 4 Stop the cluster on all nodes.

From any node, run one of the following command.

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

- 5 Verify that the cluster is stopped on all nodes by running the `ha` command.

```
# hasys -state
```

On all nodes, make sure that both the `had` and the `hashadow` processes are stopped.

- 6 Mount the software disc and move to the top-level directory on the disc.

- 7 From the top-level directory, type the following command to install on the local system only. Also, use this command to install on remote systems using the secure shell (ssh) utilities:

```
# ./installrp
```

The sample installation assumes you are using `ssh`. However, if you use the remote shell utilities to install on remote systems, additionally specify the `-rsh` option:

```
# ./installrp -rsh
```

---

**Note:** If you install using `installrp`, `installrp` stops VCS and other VCS modules (VXFEN/GAB/LLT). Therefore, before you upgrade to VCS 5.0.1 RP2, you should offline all the groups and stop VCS.

---



- 8 At the prompt to enter the system names to install the Rolling Patch 1 and press Return.

```
Enter the system names separated by spaces on which to
install 5.0.1 RP2: node1 node2 node3 node4
```

- 9 After the list of patches is displayed, press Return to continue the installation.
- 10 After installation is complete, reboot the system. If any patches fail to install, manually install them to complete the installation.

```
Rolling Patch Installation Completed.
```

```
It is strongly recommended to reboot the following systems. If any
package or patch failed to install, then manually install them to
complete the installation. node1 node2 node3 node4
```

```
Execute '/usr/sbin/shutdown -r now' to properly restart your
systems.
```

---

**Note:** If you change the configuration using Java GUI or command line the types.cf file is modified. If you have modified the types.cf file, you need to apply the same changes to the new types.cf file.

---

- 11 After installation is complete, you can verify that the patches are installed by running the following command on any node:

```
# swlist | grep GARP2
```

After the patches are installed the following information displays:

PHNE_41196	1.0	VRTS 5.0.1 GARP2 VRTSllt Kernel Patch (Veritas Low Latency Transport)
PHNE_41197	1.0	VRTS 5.0.1 GARP2 VRTSgab Kernel Patch (Veritas Group Membership and Atomic Broadcast)

12 Unfreeze all the groups.

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

13 Bring the service group online.

```
# hagrps -online service_group -sys system  
  
# hagrps -online ClusterService -sys system
```

where `system` is the system name.

## VCS 5.0.1 Rolling Patch 2 phased installation

This section explains the procedure to perform a phased upgrade of VCS on four nodes with four service groups.

---

**Note:** In this scenario, the service groups cannot stay online while you upgrade the second subcluster.

---

---

**Warning:** Do not add, remove, or change resources or service groups on any nodes during the upgrade as these changes are likely to be lost after the upgrade.

---

### Establishing running service groups

Perform the following steps for establishing running service groups.

## To establish running service groups

- 1 Establish the nodes on which your service groups are online.

```
# hagr -state
```

The output resembles the following:

```
#Group Attribute System Value
sg1 State node01 |ONLINE|
sg1 State node02 |ONLINE|
sg1 State node03 |ONLINE|
sg1 State node04 |ONLINE|
sg2 State node01 |ONLINE|
sg2 State node02 |ONLINE|
sg2 State node03 |ONLINE|
sg2 State node04 |ONLINE|
sg3 State node01 |ONLINE|
sg3 State node02 |OFFLINE|
sg3 State node03 |OFFLINE|
sg3 State node04 |OFFLINE|
sg4 State node01 |OFFLINE|
sg4 State node02 |ONLINE|
sg4 State node03 |OFFLINE|
sg4 State node04 |OFFLINE|
```

- 2 Switch the failover service groups (sg3 and sg4) from the first subcluster (node01 and node02) to the nodes on the second subcluster (node03 and node04).

```
# hagr -offline sg1 -sys node01
```

```
# hagr -offline sg2 -sys node01
```

```
# hagr -offline sg1 -sys node02
```

```
# hagr -offline sg2 -sys node02
```

```
# hagr -switch sg3 -to node03
```

```
# hagr -switch sg4 -to node04
```

- 3 Make the VCS configuration writable on the second subcluster.

```
# haconf -makerw
```

4 Freeze the service groups.

```
# hagr -freeze sg1 -persistent
# hagr -freeze sg2 -persistent
# hagr -freeze sg3 -persistent
# hagr -freeze sg4 -persistent
```

5 Dump the configuration and make it read-only.

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

6 Verify that your service groups are offline on the targeted nodes.

```
# hagr -state
```

The output resembles the following:

```
#Group Attribute System Value
sg1 State node01 |OFFLINE|
sg1 State node02 |OFFLINE|
sg1 State node03 |ONLINE|
sg1 State node04 |ONLINE|
sg2 State node01 |OFFLINE|
sg2 State node02 |OFFLINE|
sg2 State node03 |ONLINE|
sg2 State node04 |ONLINE|
sg3 State node01 |OFFLINE|
sg3 State node02 |OFFLINE|
sg3 State node03 |ONLINE|
sg3 State node04 |OFFLINE|
sg4 State node01 |OFFLINE|
sg4 State node02 |OFFLINE|
sg4 State node03 |OFFLINE|
sg4 State node04 |ONLINE|
```

---

**Note:** During the next procedure, do not perform any configuration tasks. Do not start any modules.

---

- 7 Backup llttab, llthosts, gabtab, types.cf, and main.cf files on the first subcluster.

```
# cp /etc/llttab /etc/llttab.bkp  
# cp /etc/llthosts /etc/llthosts.bkp  
# cp /etc/gabtab /etc/gabtab.bkp  
# cp /etc/VRTSvcs/conf/config/main.cf \  
/etc/VRTSvcs/conf/config/main.cf.bkp  
# cp /etc/VRTSvcs/conf/config/types.cf \  
/etc/VRTSvcs/conf/config/types.cf.bkp  
# cp /etc/VRTSvcs/conf/types.cf \  
/etc/VRTSvcs/conf/types.cf.bkp
```

- 8 Stop VCS and its components on node01 and node02.

```
# /sbin/init.d/vcs stop  
# /sbin/init.d/vxfen stop  
# /sbin/gabconfig -U  
# /sbin/lltconfig -Uo
```

- 9 Edit the following files on the first subcluster to make the values of LLT\_START, GAB\_START and VXFEN\_START equal to 0. (By default these values are 1.)

```
# /etc/rc.config.d/lltconf  
# /etc/rc.config.d/gabconf  
# /etc/rc.config.d/vxfenconf
```

Performing this step prevents LLT, GAB, and VxFEN from restarting if the nodes reboot accidentally. If the node01 and node02 reboot accidentally, it may cause node03 and node04 to panic.

- 10 Unfreeze the service groups on the second subcluster (node03 and node04).

```
# haconf -makerw  
# hagrps -unfreeze sg1 -persistent  
# hagrps -unfreeze sg2 -persistent  
# hagrps -unfreeze sg3 -persistent  
# hagrps -unfreeze sg4 -persistent  
# haconf -dump -makero
```

### Installing the first subcluster

Perform the following steps for installing VCS on the first subcluster (node01 and node02).

### To install the first subcluster

- 1 On node01 and node02, start the 5.0.1 RP2 installer for VCS.

```
# ./installrp -rsh
```

- 2 Enter the names of the nodes that you want to upgrade:

```
Enter the system names separated by spaces on which to install  
5.0.1 RP2: node01 node02
```

---

**Note:** Do not reboot the machines now.

---

- 3 Edit the following files on node01 and node02 to make the values of LLT\_START, GAB\_START and VXFEN\_START equal to 1.

```
# /etc/rc.config.d/lltconf
```

```
# /etc/rc.config.d/gabconf
```

```
# /etc/rc.config.d/vxfenconf
```

### Installing the second subcluster

Perform the following steps for installing VCS on the second subcluster (node03 and node04).

### To install the second subcluster

- 1 Take the service groups offline on node03 and node04.

---

**Warning:** Failure to perform this step can cause data corruption if you try to bring the failover groups (sg3 and sg4) online on the upgraded subcluster (node01 and node02.)

---

```
# hagrps -offline sg1 -sys node03
# hagrps -offline sg1 -sys node04
# hagrps -offline sg2 -sys node03
# hagrps -offline sg2 -sys node04
# hagrps -offline sg3 -sys node03
# hagrps -offline sg4 -sys node04
```

- 2 Stop VCS and its components on node03 and node04.

```
# /sbin/init.d/vcs stop
# /sbin/init.d/vxfen stop
# /sbin/gabconfig -U
# /sbin/lltconfig -Uo
```

---

**Note:** Your downtime starts from here.

---

- 3 Edit the following files on node03 and node04 to make the values of LLT\_START, GAB\_START and VXFEN\_START equal to 0. (By default it is 1).

```
# /etc/rc.config.d/lltconf
# /etc/rc.config.d/gabconf
# /etc/rc.config.d/vxfenconf
```

---

**Warning:** Performing this step prevents LLT, GAB, and VxFEN from restarting if the nodes reboot accidentally. If node03 and node04 reboot accidentally, it may cause node01 and node02 to panic.

---



- 4 Reboot node01 and node02.
- 5 On node01 and node02, start VCS and its components.

```
# /sbin/init.d/llt start
# /sbin/init.d/gab start
# /sbin/gabconfig -cx
# /sbin/init.d/vxfen start
# /sbin/init.d/vcs start
```
- 6 Unfreeze the service groups on the first sub cluster (node01 and node02).

```
# haconf -makerw
# hagrps -unfreeze sg1 -persistent
# hagrps -unfreeze sg2 -persistent
# hagrps -unfreeze sg3 -persistent
# hagrps -unfreeze sg4 -persistent
# haconf -dump -makero
```
- 7 On node01 and node02, bring the service groups online.

```
# hagrps -online sg1 -sys node01
# hagrps -online sg1 -sys node02
# hagrps -online sg2 -sys node01
# hagrps -online sg2 -sys node02
# hagrps -online sg3 -sys node01
# hagrps -online sg4 -sys node02
```

---

**Note:** Your downtime ends here.

---

**8** Verify the state of the service groups.

```
# hagrps -state
```

The output resembles the following:

```
#Group Attribute System Value
sg1 State node01 |ONLINE|
sg1 State node02 |ONLINE|
sg1 State node03 |OFFLINE|
sg1 State node04 |OFFLINE|
sg2 State node01 |ONLINE|
sg2 State node02 |ONLINE|
sg2 State node03 |OFFLINE|
sg2 State node04 |OFFLINE|
sg3 State node01 |ONLINE|
sg3 State node02 |OFFLINE|
sg3 State node03 |OFFLINE|
sg3 State node04 |OFFLINE|
sg4 State node01 |OFFLINE|
sg4 State node02 |ONLINE|
sg4 State node03 |OFFLINE|
sg4 State node04 |OFFLINE|
```

**9** On node03 and node04, start the 5.0.1 RP2 installer for VCS.

```
# ./installrp -rsh
```

**10** Enter the names of the nodes that you want to upgrade:

```
Enter the system names separated by spaces on which to install
5.0.1 RP2: node03 node04
```

---

**Note:** Do not reboot the machines now.

---

- 11 Edit the following files on node03 and node04 to make the values of LLT\_START, GAB\_START and VXFEN\_START equal to 1.

```
# /etc/rc.config.d/lltconf  
# /etc/rc.config.d/gabconf  
# /etc/rc.config.d/vxfenconf
```

- 12 Reboot node03 and node04.

After you have rebooted the nodes, all the four nodes now run the latest version of VCS, which is 5.0.1 RP2.

In this example, you have performed a phased upgrade of VCS because the service groups were down only from the point when you took them offline on node01 and node02, to the time VCS brought them online on node03 or node04 as appropriate.

## Installing Veritas Storage Foundation Cluster File System 5.0.1 Rolling Patch 2

You can perform the installation in one of the following ways:

- Full installation  
See [“Veritas Storage Foundation Cluster File System full installation”](#) on page 51.
- Phased installation  
See [“Veritas Storage Foundation Cluster File System phased installation”](#) on page 54.

### Veritas Storage Foundation Cluster File System full installation

Installing the 5.0.1 Rolling Patch 2 on a cluster requires you to stop the cluster during the installation. However, CVM services remain available.

To install 5.0.1 Rolling Patch 2 on an SFCFS cluster

- 1 Log in as the root user.
- 2 Verify that the path is `/opt/VRTS/bin` is in the PATH variable.
- 3 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 4 Run the following command to freeze HA service group operations on each node:

```
# hasys -freeze -persistent nodename
```

- 5 Change the cluster configuration to read-only mode on any of the nodes:

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

- 6 Stop all applications on the cluster that are not configured under VCS. Use native application commands to stop the application.
- 7 Unmount the VxFS and CFS file systems that are not managed by VCS. Make sure 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 vx  
# fuser -cu /mount_point  
# umount /mount_point
```

- 8 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
```

- 9 Stop VCS on the cluster from any of the nodes:

```
# hastop -all
```

- 10 Stop the VCS command server on the cluster:

```
# ps -ef | grep CmdServer  
# kill -9 pid_of_CmdServer
```

where *pid\_of\_CmdServer* is the process ID of the VCS command server.

**11** Verify that only ports a, b and d are open:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 4d3c08 membership 0123
Port b gen 4d3c0c membership 0123
Port d gen 4d3c0b membership 0123
```

**12** Mount the 5.0.1 Rolling Patch 2 product disc and navigate to the directory that contains the installation script. Run the installrp script and specify the names of the nodes:

```
# ./installrp galaxy nebula
```

**13** Restart the nodes:

```
# shutdown -r now
```

**14** Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

**15** Run the following command on each node to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent nodename
```

**16** Change the cluster configuration to read-only mode from any of the nodes:

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

**17** Bring the VCS service groups online from any of the nodes:

```
# hagrps -online group_name -any
```

---

**Note:** Downtime for failover service groups ends here.

---

- 18 Start the applications on the cluster that are not configured under VCS. Use native application commands to start the application.
- 19 To list the SFCFS 5.0.1 Rolling Patch 2 patches installed on your system, run the following command:

```
# swlist| egrep 'PHCO_41194|
|PHKL_41074|PVCO_03916|PHCO_41192|PHCO_40772
|PHCO_40771|PHNE_41196|PHCO_40770|PHKL_41193
|PHNE_41197|PHKL_40443|PHCO_40690'
```

## Veritas Storage Foundation Cluster File System phased installation

A phased installation involves installing Veritas Storage Foundation Cluster File System 5.0.1 Rolling Patch 2 on half of the nodes in the cluster at a time. The examples in the procedures assume a three-node SFCFS 5.0.1 cluster with the nodes galaxy and nebula constituting the first half of the cluster and the node jupiter constituting the second half of the cluster.

---

**Note:** Before starting the installation on the first half of the cluster, back up the configuration files.

---

### To install 5.0.1 RP2 on an SFCFS cluster

- 1 Log in as the root user.
- 2 Switch failover groups from a node in the first half of the cluster (galaxy) to the node in the second half of the cluster (jupiter):

```
# hagrps -switch failover_group -to jupiter
```

- 3 Change the cluster configuration to the read-write mode:

```
# haconf -makerw
```

- 4 Run the following commands on each node in the first half of the cluster to freeze HA service group operations :

```
# hasys -freeze -persistent galaxy
```

```
# hasys -freeze -persistent nebula
```

- 5 Change the cluster configuration to the read-only mode on any of the nodes:

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

- 6 On the first half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application.

- 7 On the first half of the cluster, unmount the VxFS and CFS file systems that are not managed by VCS. Make sure that no processes are running that make use of mounted shared file system or shared volume:

```
# mount | grep vx  
  
# fuser -kuc /mount_point  
  
# umount /mount_point
```

- 8 Stop the VCS daemons on each node of the first half of the cluster (galaxy and nebula) by using the following commands:

```
# hastop -local -force  
  
# vxclustadm stopnode  
  
# ps -eaf | grep vxfsckd  
  
# kill -9 pid_of_vxfsckd
```

where `pid_of_vxfsckd` is the process ID of `vxfsckd`.

```
# fsclustadm cfsdeinit  
  
# /sbin/init.d/odm stop (If applicable)  
  
# /sbin/init.d/vxfen stop  
  
# /sbin/init.d/gab stop  
  
# /sbin/init.d/llt stop
```

- 9 Verify that none of the ports is running on the cluster:

```
# gabconfig -a
```

- 10** Stop the VCS command server on the cluster:

```
# ps -ef | grep CmdServer  
  
# kill -9 pid_of_Cmdserver
```

where `pid_of_CmdServer` is the process ID of the VCS command server.

- 11** Mount the 5.0.1 RP2 product disc and navigate to the folder that contains the installation script. On the first half of the cluster, enter the following command:

```
# ./installrp -rsh galaxy nebula
```

---

**Caution:** Do not reboot the first half of the cluster when prompted.

---

**Note:** The cluster downtime starts here.

---

- 12** On the second half of the cluster (jupiter), change the cluster configuration to the read-write mode:

```
# haconf -makerw
```

- 13** On the second half of the cluster, run the following command to freeze HA service group operations:

```
# hasys -freeze -persistent jupiter
```

- 14** On the second half of the cluster, change the cluster configuration to the read-only mode:

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

- 15** Unmount the devices mounted on the first half of the cluster by using the following commands:

```
# mount -p  
  
# umount mount_point
```

- 16** On the second half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application.



- 17** On the second half of the cluster, unmount the VxFS or CFS file systems that are not managed by VCS. Make sure that no processes are running that make use of the mounted shared file system or shared volumes. Verify that no processes use the VxFS or CFS mount point by using the following commands:

```
# fuser -c /mount_point

# kill -9 pid_number

# umount /mount_point
```

- 18** Stop all the VCS daemons to shut down the second half of the cluster by using the following commands:

```
# hstop -local -force

# vxclustadm stopnode

# ps -eaf | grep vxfsckd

# kill -9 pid_of_vxfsckd
```

where `pid_of_vxfsckd` is the process ID of `vxfsckd`.

```
# fsclustadm cfsdeinit

# /sbin/init.d/odm stop (If applicable)

# /sbin/init.d/vxfen stop

# /sbin/init.d/gab stop

# /sbin/init.d/llt stop
```

- 19** Verify that none of the ports is running on the cluster:

```
# gabconfig -a
```

- 20** Stop the VCS command server on the second half of the cluster:

```
# ps -ef | grep CmdServer

# kill -9 pid_of_CMDserver
```

where `pid_of_CmdServer` is the process ID of the VCS command server.

21 After the second half of the cluster is down, reboot the first half of the cluster

```
# shutdown -r now
```

22 Bring the first half of the cluster online by enabling the seed port:

```
# gabconfig -cx
```

23 Start the HA daemon on the first half of the cluster:

```
# hstart
```

24 Unfreeze the first half of the cluster:

```
# haconf -makerw
```

```
# hasys -unfreeze -persistent galaxy
```

```
# hasys -unfreeze -persistent nebula
```

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

25 Make online the CVM group on all the nodes of the first half of the cluster:

```
# hgrp -online cvm -sys galaxy
```

```
# hgrp -online cvm -sys nebula
```

26 Verify that all the ports are up and running on the first half of the cluster by using the following command:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
```

```
Port a gen a2b701 membership 01  
Port b gen a2b704 membership 01  
Port f gen a2b70b membership 01  
Port h gen a2b703 membership 01  
Port v gen a2b707 membership 01  
Port w gen a2b709 membership 01
```

---

**Note:** The cluster downtime ends here.

---

- 27** Navigate to the folder that contains the installation script. On the second half of the cluster, enter the following command:

```
# ./installrp -rsh jupiter
```

- 28** Reboot the second half of the cluster:

```
# shutdown -r now
```

- 29** After the node in the second half of the cluster is up, bring the cluster online by manually seeding the port:

```
# gabconfig -cx
```

- 30** Start the HA daemon on the second half of the cluster:

```
# hastart
```

- 31** Unfreeze the second half of the cluster:

```
# haconf -makerw
```

```
# hasys -unfreeze -persistent jupiter
```

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

- 32** Make online the CVM group on the node in the second half of the cluster:

```
# hagrps -online cvm -sys jupiter
```

- 33** Verify that all the ports are up and running on the second half of the cluster by using following command:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen    a2b701 membership 012
Port b gen    a2b704 membership 012
Port f gen    a2b70b membership 012
Port h gen    a2b703 membership 012
Port v gen    a2b707 membership 012
Port w gen    a2b709 membership 012
```

---

**Note:** The node in the second half of the cluster (jupiter) joins the first half of the cluster here.

---

- 34** Bring the failover service group online on the second half of the cluster:

```
# hagrpf -online failover_group -sys jupiter
```

- 35** Switch the failover service groups back from the node in the second half of the cluster (jupiter) to a node in the first half of the cluster (galaxy):

```
# hagrpf -switch failover_group -to galaxy
```

- 36** On the second half of the cluster, manually mount the VxFS and CFS file systems that are not managed by VCS.

- 37** On the second half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

## Installing SF Oracle RAC 5.0.1 Rolling Patch 2

You can perform the installation in one of the following ways:

- Full installation  
See [“SF Oracle RAC full installation”](#) on page 61.
- Phased installation  
See [“SF Oracle RAC phased installation”](#) on page 62.

For instructions on installing or upgrading to Oracle RAC 11g Release 2, see the *Appendix: Installing or upgrading to Oracle RAC 11g Release 2* in this document.

## SF Oracle RAC full installation

Perform the following steps to install 5.0.1 Rolling Patch 2.

### To install 5.0.1 Rolling Patch 2 on an SF Oracle RAC cluster

- 1 Log in as root user.
- 2 On the each node of the cluster, stop Oracle database. If the Oracle RAC instance is managed by VCS run the following commands:

```
# hagr -offline oracle_group -sys galaxy
# hagr -offline oracle_group -sys nebula
# hagr -offline oracle_group -sys jupiter
# hagr -offline oracle_group -sys mercury
```

If the Oracle RAC instance is not managed by VCS, log in as the Oracle user and shut down the database instances:

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

- 3 Stop the Veritas Enterprise Administrator if it is running:
- 4 Stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 5 Unmount the VxFS and CFS file systems that are not managed by VCS. Ensure that no processes, that use the mounted shared file system or shared volumes, are running. To verify that no processes use the VxFS or CFS mount points, run the following commands:

```
# cat /etc/mnttab | grep vxfs
# fuser -uc /mount_point
# umount /mount_point
```

- 6 Stop all VxVM and CVM volumes for each disk group that is not managed by VCS:

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

Verify that no volumes remain open:

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

7 Stop VCS:

```
# hastop -all
```

8 Verify that only ports a, b, d and o are open:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 4d3c08 membership 0123
Port b gen 4d3c0c membership 0123
Port d gen 4d3c0b membership 0123
Port o gen 4d3c27 membership 0123
```

9 Mount the 5.0.1 Rolling Patch 2 product disc and navigate to the folder that contains the installation script. Enter the installrp script:

```
# ./installrp [-rsh] galaxy nebula jupiter mercury
```

10 Relink the SF Oracle RAC libraries with the Oracle libraries.

See [“Relinking the SF Oracle RAC libraries with Oracle RAC manually”](#) on page 97.

11 Copy the new MultiPrivNIC.cf file from /etc/VRTSvcs/conf to /etc/VRTSvcs/conf/config

12 Restart the nodes:

```
# shutdown -r now
```

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

```
$ srvctl start instance -d database_name -i instance_name
```

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

## SF Oracle RAC phased installation

A phased installation involves installing SF Oracle RAC 5.0.1 Rolling Patch 2 on half of the nodes in the cluster at a time. The examples in the procedures assume a four-node SF Oracle RAC cluster with the nodes galaxy and nebula constituting the first half of the cluster and the nodes jupiter and mercury constituting the second half of the cluster.

---

**Note:** Before you start the installation on the first half of the cluster, back up the configuration files.

---

### To install 5.0.1 Rolling Patch 2 on an SF Oracle RAC cluster

- 1 Log in as the root user.
- 2 Switch failover groups from a node in the first half of the cluster(galaxy/nebula) to the nodes in the second half of the cluster(jupiter and mercury). For example:

```
# hagrps -switch failover_group -to jupiter
# hagrps -switch failover_group -to mercury
```

- 3 On the first half of the cluster, stop Oracle database. If the Oracle RAC instance is managed by VCS run the following commands:

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

If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the first half of the cluster and shut down the instances:

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

- 4 On the first half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 5 On the first half of the cluster, unmount the VxFS and CFS file systems that are not managed by VCS:

```
# cat /etc/mnttab | grep vxfs
# fuser -cu /mount_point
# umount /mount_point
```

- 6 On first half of the cluster stop all VxVM and CVM volumes for each diskgroup that is not managed by VCS:

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

Verify that no volumes remain open:

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

- 7 On first half of the cluster, stop VCS:

```
# hastop -local
```

- 8 Verify that only ports a, b, d and o are open:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 6b5901 membership 01
Port b gen 6b5904 membership 01
Port d gen 6b5907 membership 01
Port o gen 6b5905 membership 01
```

- 9 Mount the 5.0.1 Rolling Patch 2 product disc and navigate to the folder that contains the installation script. On the first half of the cluster, enter the `installrp` script:

```
# ./installrp [-rsh] galaxy nebula
```

---

**Note:** After you complete the installation on the first half of the cluster, no GAB ports appear in the output when you run the `gabconfig -a` command.

---

- 10 On the first half of the cluster, relink the SF Oracle RAC libraries with the Oracle libraries.

See [“Relinking the SF Oracle RAC libraries with Oracle RAC manually”](#) on page 97.

- 11 Copy the new `MultiPrivNIC.cf` file from `/etc/VRTSvcs/conf` to `/etc/VRTSvcs/conf/config`

- 12 On the first half of the cluster, restart the nodes:

```
# shutdown -r now
```

---

**Note:** After first half of the cluster restarts, ports a, b, d, and o form membership with the second half of the cluster. HAD(port h) does not form membership because it reports an engine version mismatch.

---



- 13 On the second half of the cluster, take all the failover service groups offline:

```
# hagrps -offline group_name -any
```

- 14 On the second half of the cluster, bring down port h:

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

---

**Note:** You must run this command on every node in the second half of the cluster. This command does not stop the CVM, CFS, or the applications managed by VCS and there is no downtime for applications.

---

- 15 On the first half of the cluster, start VCS:

```
# hstart
```

---

**Note:** You must run this command on every node in the first half of the cluster.

---

- 16 On the first half of the cluster, bring the VCS service groups online. For failover service groups:

```
# hagrps -online group_name -any
```

Once the cvm service group comes online, the GAB ports (v, w, and f) and all the service groups pertaining to the CFS mounts also come online automatically. The failover service groups must be brought online manually using the above command.

- 17 On the first half of the cluster, manually mount the VxFS or CFS file systems that are not managed by VCS.
- 18 On the first half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.
- 19 On the first half of the cluster, bring down port h:

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

---

**Note:** You must run this command on every node in the first half of the cluster. This command does not stop the CVM, CFS, or the applications managed by VCS and there is no downtime for applications.

---

- 20 On the second half of the cluster, start VCS:

```
# hastart
```

---

**Note:** There is no downtime for applications in this process.

---

- 21 On the second half of the cluster, stop the Oracle database. If the Oracle RAC instance is managed by VCS:

```
# hagrpl -offline oracle_group -sys jupiter  
# hagrpl -offline oracle_group -sys mercury
```

If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the second half of the cluster and shut down the instances:

```
$ srvctl stop instance -d database_name -i instance_name  
$ CRS_HOME/bin/crsctl stop crs
```

- 22 On the second half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 23 On the second half of the cluster, unmount the VxFS or CFS file systems that are not managed by VCS. Make sure that no processes are running that make use of the mounted shared file system or shared volumes.

To verify that no processes use the VxFS or CFS mount point:

```
# cat /etc/mnttab | grep vxfs  
# fuser -cu /mount_point  
# umount /mount_point
```

- 24 On the second half of the cluster, 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:

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

- 25 On each node in the second half of the cluster, stop VCS:

```
# hastop -local
```

26 On each node in the first half of the cluster, start VCS:

```
# hastart
```

27 Navigate to the folder that contains the installation script. On the second half of the cluster, enter the `installrp` script:

```
# ./installrp [-rsh] jupiter mercury
```

28 On the second half of the cluster, relink the SF Oracle RAC libraries with the Oracle libraries.

See [“Relinking the SF Oracle RAC libraries with Oracle RAC manually”](#) on page 97.

29 Copy the new `MultiPrivNIC.cf` file from `/etc/VRTSvcs/conf` to `/etc/VRTSvcs/conf/config`

30 Restart the nodes in the second half of the cluster:

```
# shutdown -r now
```

When the nodes in the second half of the cluster come online, all the GAB ports (a, b, d, o, h, v, w, and f) the CFSSMount service groups also come online automatically.

31 On the second half of the cluster, manually mount the VxFS and CFS file systems that are not managed by VCS.

32 On the second half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

## Uninstalling the patches

This section provides instructions for uninstalling 5.0.1 Rolling Patch 2.

Depending on the product you want to uninstall, perform the steps in one of the following sections:

- Veritas Storage Foundation:  
See [“Uninstalling Veritas Storage Foundation 5.0.1 Rolling Patch 2”](#) on page 68.
- Veritas Cluster Server:  
See [“Uninstalling VCS 5.0.1 Rolling Patch 2”](#) on page 69.
- Veritas Storage Foundation Cluster File System:  
See [“Uninstalling Veritas Storage Foundation Cluster File System 5.0.1 Rolling Patch 2”](#) on page 71.

- Veritas Storage Foundation for Oracle RAC  
See “[Uninstalling SF Oracle RAC 5.0.1 Rolling Patch 2](#)” on page 73.

## Uninstalling Veritas Storage Foundation 5.0.1 Rolling Patch 2

Perform the following procedure to remove Veritas Storage Foundation 5.0.1 Rolling Patch 2. It is recommended to refer to the Installation Summary Files for the list of patches that get installed during installation.

### To uninstall Veritas Storage Foundation 5.0.1 Rolling Patch 2

- 1 Log in as root.
- 2 On all the nodes, stop the currently running VxPAL agents. See `vxpalctrl` (1M):
  - Stop the storage agent:

```
# /opt/VRTSobc/pa133/bin/vxpalctrl -a StorageAgent -c stop
```
  - Check the status of the storage agent:

```
# /opt/VRTSobc/pa133/bin/vxpalctrl -a StorageAgent -c status
```
  - Stop the action agent:

```
# /opt/VRTSobc/pa133/bin/vxpalctrl -a actionagent -c stop
```
  - Check the status of the action agent:

```
# /opt/VRTSobc/pa133/bin/vxpalctrl -a actionagent -c status
```
  - Stop the gridnode agent:

```
# /opt/VRTSobc/pa133/bin/vxpalctrl -a gridnode -c stop
```
  - Check the status of the gridnode agent:

```
# /opt/VRTSobc/pa133/bin/vxpalctrl -a gridnode -c status
```
- 3 On all the nodes, remove all the patches by using the `swremove` command:

```
# swremove -x autoreboot=true patch_name1, patch_name2 ...
```

Symantec recommends that all the patches installed during Veritas Storage Foundation 5.0.1 Rolling Patch 2 installation be removed through a single command line. The system automatically reboots after removing the patches.

## Uninstalling VCS 5.0.1 Rolling Patch 2

Perform the following procedure to uninstall VCS 5.0.1 Rolling Patch 2.

### To uninstall VCS 5.0.1 Rolling Patch 2

- 1 List the service groups in the cluster along with their status.

```
# hagrps -state
```

- 2 Offline the ClusterService group if it is configured.

```
# hagrps -offline -force ClusterService -sys system
```

- 3 Make the VCS configuration writable.

```
# haconf -makerw
```

- 4 Freeze all the service groups.

```
# hagrps -freeze service_group -persistent
```

where `service_group` is the name of the service group.

---

**Note:** The ClusterService group cannot be frozen.

---

- 5 Save the configuration (main.cf) file with the frozen groups.

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

- 6 Take a backup of the current main.cf and all types.cf configuration files.

For example, on one of the nodes in the cluster, type:

```
# 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/types.cf \
/etc/VRTSvcs/conf/types.cf.save
```

- 7 Shut down VCS and the VCS CmdServer.

```
# hastop -all -force  
  
# CmdServer -stop
```

- 8 Verify that VCS has shut down.

```
# gabconfig -a
```

The output resembles

```
GAB Port Memberships  
  
Port a gen 23dc0001 membership 01
```

---

**Note:** The output shows no membership for port h.

---

- 9 Stop vxfen on each cluster node, if the VCS cluster uses fencing option.

```
# /sbin/init.d/vxfen stop
```

- 10 Unconfigure GAB.

```
# /sbin/gabconfig -U
```

- 11 Unconfigure LLT.

```
# /sbin/lltconfig -Uo
```

- 12 Remove the VCS patches from each node in the cluster.

```
#swlist -x autoreboot=true -x enforce_dependencies=false  
  
PHNE_41196 PHNE_41197
```

- 13 Restore the types.cf configuration files from the location where you saved them, or manually edit the /etc/VRTSvcs/conf/config/types.cf file to remove the newly added attributes.

- 14 Restart all the nodes in the cluster.

```
# shutdown -ry now
```

- 15 After VCS has started, perform the following steps:

- Verify whether all the resources have been probed.

```
# hastatus -summary
```

- Unfreeze all the service groups.

```
# haconf -makerw
```

```
# hagrps -unfreeze service_group -persistent
```

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

where `service_group` is the name of the service group.

- Bring the ClusterService group online.

```
# hagrps -online service_group -sys system
```

```
# hagrps -online ClusterService -sys system
```

where `system` is the system name.

## Uninstalling Veritas Storage Foundation Cluster File System 5.0.1 Rolling Patch 2

Perform the following procedure to remove 5.0.1 Rolling Patch 2.

### To uninstall SFCFS 5.0.1 Rolling Patch 2

- 1 List the service groups in the cluster along with their status. On any node, type:

```
# hagrps -state
```

- 2 Take the ClusterService group offline if it is configured.

```
# hagrps -offline -force ClusterService -sys system
```

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

```
# haconf -makerw
```

- 4 Freeze all the service groups. On any node, type:

```
# hagrps -freeze service_group -persistent
```

where `service_group` is the name of the service group.

---

**Note:** The ClusterService group cannot be frozen.

---

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

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

- 6 Take a backup of the current main.cf and all types.cf configuration files. For example, on one of the nodes 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
```

- 7 Shut down VCS and the VCS CmdServer. On any node, type:

```
# hastop -all -force  
  
# CmdServer -stop
```

- 8 Verify that VCS has shut down. On each node, type:

```
# gabconfig -a
```

The output resembles

```
GAB Port Memberships  
Port a gen 23dc0001 membership 01
```

---

**Note:** The output shows no membership for port h.

---

- 9 Stop vxfen on each cluster node, if the VCS cluster uses fencing option.

```
# vxfenconfig -U
```

- 10 Unconfigure GAB. On each node, type:

```
# gabconfig -U
```

- 11 Unconfigure LLT. On each node, type:

```
# lltconfig -Uo
```



- 12 Remove the SFCFS patches from each node in the cluster. Type the following command:

```
# swremove -x autoreboot=true -x enforced_dependencies=false  
patch_name1, patch_name2 ...
```

- 13 Restore the types.cf configuration files from the location where you saved them, or manually edit the `/etc/VRTSvcs/conf/config/types.cf` file to remove the newly added attributes.

- 14 Restart all the nodes in the cluster.

```
# shutdown -ry now
```

- 15 After SFCFS has started, perform the following steps:

- Verify whether all the resources have been probed. On each node, type:

```
# hastatus -summary
```

- Unfreeze all the service groups. On any node, type:

```
# haconf -makerw
```

```
# hagrps -unfreeze service_group -persistent
```

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

where `service_group` is the name of the service group.

- If necessary, bring the ClusterService group online. On any node type:

```
# hagrps -online ClusterService -sys  
system
```

where `system` is the system name.

## Uninstalling SF Oracle RAC 5.0.1 Rolling Patch 2

The following procedure describes the steps involved in uninstalling 5.0.1 Rolling Patch 2.

The examples in the procedures assume a four-node SF Oracle RAC cluster with the nodes `galaxy`, `nebula`, `jupiter` and `mercury`.

---

**Note:** Before uninstalling the SF Oracle RAC 5.0.1 Rolling Patch 2, back up your existing configuration files.

---

## To uninstall the SF Oracle RAC 5.0.1 Rolling Patch 2

- 1 Log in as the root user.
- 2 On each node of the cluster, stop the Oracle database. If the Oracle RAC instance is managed by VCS:

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

If the Oracle RAC instance is not managed by VCS, log in as the Oracle user and shut down the Oracle database instances:

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

- 3 Stop the Veritas Enterprise Administrator if it is running:

```
# /opt/VRTSob/bin/vxsvcctrl stop
```

- 4 Stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 5 Unmount the VxFS and CFS file systems that are not managed by VCS. Ensure that no processes, that make use of the mounted shared file system or shared volumes, are running.

To verify that no processes use the VxFS or CFS mount points, run the following commands:

```
# cat /etc/mnttab | grep vxfs
# fuser -cu /mount_point
# umount /mount_point
```

- 6 Stop all VxVM and CVM volumes for each disk group that is not managed by VCS:

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

Verify that no volumes remain open:

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

- 7 Stop VCS:

```
# hastop -all
```

**8** Unmount ODM:

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

**9** Stop fencing:

```
# /sbin/init.d/vxfen stop
```

**10** Stop VCSMM, GAB, LMX, and LLT:

```
# vcsmmconfig -U  
# gabconfig -U  
# lmxconfig -U  
# lltconfig -U
```

**11** Remove the following patches from each node of cluster:

---

**Note:** This command reboots the node after you remove the patch.

---

```
# swremove -x autoreboot=true patch_list
```

where *patch\_list* is the following list of patches:

```
PVKL_03904 PVCO_03916 PHNE_41196 PHNE_41197  
PHCO_41194 PHKL_41074 PHCO_41192 PHKL_41193  
PHKL_40772 PHCO_40771 PHCO_40770 PHKL_40443  
PHCO_40690 PVCO_03905 PVCO_03906 PHCO_40520  
PHCO_39783 PHCO_40673 PHCO_40674 PHCO_40782  
PHCO_41221 PHCO_41321 PHCO_40672
```

**12** After the nodes restart verify that the patches have been removed from the nodes:

```
# swlist| egrep 'patch_list' | \  
wc -l
```

The output must be 0.

**13** Verify that all the ports are up:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 70d901 membership 01
Port b gen 70d906 membership 01
Port d gen 70d905 membership 01
Port f gen 70d911 membership 01
Port h gen 70d909 membership 01
Port o gen 70d904 membership 01
Port v gen 70d90c membership 01
Port w gen 70d90e membership 01
```

**14** For relinking Oracle binaries with SF Oracle RAC 5.0.1 binaries, stop the Oracle database and Oracle Clusterware on the each node of the cluster.

If the Oracle RAC instance is managed by VCS:

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

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

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

**15** Relink the SF Oracle RAC libraries with the Oracle libraries.

See [“Relinking the SF Oracle RAC libraries with Oracle RAC manually”](#) on page 97.

**16** On the each node of the cluster, start the Oracle database. If the Oracle RAC instance is managed by VCS:

```
# hagrpl -online oracle_group -sys galaxy
# hagrpl -online oracle_group -sys nebula
# hagrpl -online oracle_group -sys jupiter
# hagrpl -online oracle_group -sys mercury
```

If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes and start the Oracle database instances:

```
$ srvctl start instance -d database_name -i instance_name
```

## Reference documentation

The documentation for this release is available on the software disc in the PDF format. Symantec recommends copying documentation from the disc to your system directory. This release includes the following document.

[Table 1-13](#) lists the document included in this release.

**Table 1-13** Documentation in 5.0.1 Rolling Patch 2

Title	File Name
<i>Veritas Storage Foundation and High Availability Solutions Release Notes</i>	sfha_notes_501rp2_hpux.pdf

To refer to the product documentation for Veritas Storage Foundation and High Availability Solutions 5.0.1, go to:

<http://www.symantec.com/business/support/overview.jsp?pid=15107> and click on the **HP-UX** tab.



# Installing or upgrading to Oracle RAC 11g Release 2

This appendix includes the following topics:

- [Installing Oracle RAC 11g Release 2](#)
- [Upgrading to Oracle RAC 11g Release 2](#)

## Installing Oracle RAC 11g Release 2

This section provides instructions for installing Oracle RAC 11g Release 2.

Perform the following steps to install Oracle RAC 11g Release 2:

1. Complete the preparatory tasks before starting the Oracle RAC installation.  
See [“Preparing to install Oracle RAC using the SF Oracle RAC installer or manually”](#) on page 80.
2. Install Oracle RAC 11g Release 2.  
See [“Installing Oracle Grid Infrastructure using the Oracle Universal Installer”](#) on page 92.  
See [“Configuring LLT links in the GPnP profile”](#) on page 95.  
See [“Installing the Oracle RAC database using the Oracle Universal Installer”](#) on page 95.
3. Complete the following post-installation tasks:
  - Relink the SF Oracle RAC libraries with Oracle RAC:  
See [“Relinking the SF Oracle RAC libraries with Oracle RAC manually”](#) on page 97.
  - Configure the CSSD resource.

- See [“Configuring the CSSD resource manually”](#) on page 97.
- Modify the Oracle RAC configuration to prevent automatic startup of Oracle Clusterware.  
See [“Preventing automatic startup of Oracle Clusterware”](#) on page 99.
- Create the Oracle RAC database.  
See [“Creating the Oracle RAC database”](#) on page 99.
- Modify the Oracle RAC configuration to prevent automatic startup of the Oracle database.  
See [“Preventing automatic database startup”](#) on page 99.
- Configure VCS service groups for Oracle RAC.  
See [“Configuring VCS service groups for Oracle RAC”](#) on page 100.

## Preparing to install Oracle RAC using the SF Oracle RAC installer or manually

This section provides instructions for performing the pre-installation tasks for Oracle RAC 11g Release 2.

The following preparatory tasks are similar to those documented in the version 5.0.1 guide and are therefore not described in this section. For instructions, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide (5.0.1)*, Chapter "Preparing to install Oracle RAC".

- Identify the public virtual IP addresses for use by Oracle.
- Set the kernel parameters.
- Verify the user "nobody" exists.
- Create Oracle user and groups.  
In Oracle RAC 11g Release 2, you need to create the grid user and the user home directory on each system in addition to the Oracle user and group.
- Create storage for OCR and voting disk.  
The Oracle Clusterware files for Oracle Cluster Registry (OCR) and voting disk in Oracle RAC 11g Release 2 must reside in directories in the cluster file system (CFS) or on ASM disk groups. You can create ASM disk groups using CVM raw volumes.
- Configure private IP addresses for Oracle RAC.

Perform the following tasks using the instructions in this section:

- Identify the SCAN IP address.  
See [“Identifying the SCAN IP address”](#) on page 81.



- Create Oracle Grid Infrastructure and Oracle database home directories manually.  
See [“Creating Oracle Grid Infrastructure and Oracle database home directories manually”](#) on page 81.
- Set up user equivalence on cluster nodes.  
See [“Setting up user equivalence”](#) on page 91.
- Synchronize time settings on cluster nodes.  
See [“Synchronizing time settings on cluster nodes”](#) on page 91.
- Edit the user profiles.  
See [“Editing the user profiles”](#) on page 91.

## Identifying the SCAN IP address

You need a public IP address available and configured as a Single Client Access Name (SCAN) address on the domain name server (DNS) for round robin resolution to three addresses (recommended) or at least one address. SCAN is used to connect to databases within the cluster irrespective of the nodes they are running on.

---

**Note:** The virtual IP addresses that are used for SCAN IP resolution must be on the same subnet. Oracle RAC does not support their configuration on different subnets.

---

## Creating Oracle Grid Infrastructure and Oracle database home directories manually

You can create the Oracle directories on the local file system, on a local Veritas file system, or on a Veritas cluster file system.

---

**Note:** Symantec and Oracle recommend that you install the Oracle Grid Infrastructure and Oracle database binaries local to each node in the cluster.

---

[Table A-1](#) lists the Oracle directories you need to create:

**Table A-1** List of directories

Directory	Description
<p>Oracle Grid Infrastructure Home Directory (GRID_HOME)</p>	<p>The path to the home directory that stores the Oracle Grid Infrastructure binaries. The Oracle Universal Installer (OUI) installs Oracle Clusterware and Oracle ASM into this directory, also referred to as GRID_HOME.</p> <p>The directory must be owned by the installation owner of Oracle Grid Infrastructure (oracle or grid), with the permission set to 775.</p> <p>The path to the Grid home directory must be the same on all nodes. As the root user, create a path compliant with Oracle Optimal Flexible Architecture (OFA) guidelines, so that the OUI can select that directory during installation.</p>
<p>Oracle base directory (ORACLE_BASE)</p>	<p>The base directory that contains all the Oracle installations. Create separate Oracle base directories for the grid user and the Oracle user.</p> <p>It helps to ensure that installations of multiple databases maintain an Optimal Flexible Architecture (OFA) configuration.</p> <p>The path to the Oracle base directory must be the same on all nodes.</p>
<p>Oracle home directory (ORACLE_HOME)</p>	<p>The directory in which the Oracle RAC software is installed. The path to the Oracle home directory must be the same on all nodes.</p> <p>You can create the Oracle home directory on the local file system or on a cluster file system.</p>

Use one of the following options to create the directories:

- Local file system      See [“To create the directories on the local file system”](#) on page 83.
- Veritas File System      See [“To create the file system and directories on the Veritas File system \(VxFS\)”](#) on page 83.
- Cluster File System      See [“To create the file system and directories on cluster file system for Oracle Grid Infrastructure and Oracle database”](#) on page 86.

### To create the directories on the local file system

- 1 Log in as the root user on each node.
- 2 Create the Oracle base directory (ORACLE\_BASE):

For grid user:

```
# mkdir -p /u02/app/grid
# chown -R grid:oinstall /u02/app/grid
# chmod -R 775 /u02/app/grid
```

For Oracle user:

```
# mkdir -p /u01/app/oracle
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

- 3 Create the Oracle Grid Infrastructure home directory (GRID\_HOME):

---

**Note:** Ensure the grid home directory is not a subdirectory of the Oracle base directory. Installing Oracle Clusterware in an Oracle base directory causes installation errors.

---

```
# mkdir -p /u02/app/11.2.0/grid
# chown -R grid:oinstall /u02/app/11.2.0/grid
```

- 4 Create the Oracle database home directory (ORACLE\_HOME):

```
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

### To create the file system and directories on the Veritas File system (VxFS)

The sample commands in the procedure are for node galaxy. Repeat the steps on each node of the cluster.

- 1 As the root user, create a VxVM local disk group `bindg_hostname` on each node.

Make sure that the size of the disk group is approximately 12 GB; if not, add another disk to the disk group.

```
# vxdg init bindg_galaxy Disk_1
```

- 2 Create separate volumes for Oracle Grid Infrastructure (`crsbinvol`) and Oracle database (`orabinvol`):

```
# vxassist -g bindg_galaxy make crsbinvol 5G  
# vxassist -g bindg_galaxy make orabinvol 7G
```

- 3 Create file systems with the volumes `crsbinvol` and `orabinvol`:

```
# mkfs -F vxfs /dev/vx/rdisk/bindg_galaxy/crsbinvol  
# mkfs -F vxfs /dev/vx/rdisk/bindg_galaxy/orabinvol
```

- 4 Mount the file systems. Perform this step on each node.

```
# mount -F vxfs /dev/vx/dsk/bindg_galaxy/crsbinvol \  
/u02  
# mount -F vxfs /dev/vx/dsk/bindg_galaxy/orabinvol \  
/u01
```

- 5 Create the following directories for Oracle, `ORACLE_BASE`, `GRID_HOME`, `ORACLE_HOME`.

The file system and directories created on shared storage in this procedure are based on the following layout:

```
$ORACLE_BASE      For grid user:
                   /u02/app/grid
                   For Oracle user:
                   /u01/app/oracle

$GRID_HOME        /u02/app/11.2.0/grid

$ORACLE_HOME      /u01/app/oracle/product/11.2.0/dbhome_1

# mkdir -p /u02/app/11.2.0/grid
# mkdir -p /u01/app/oracle
# mkdir -p /u02/app/grid
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
```

- 6 Change the ownership and permissions on all nodes of the cluster.

---

**Note:** The ownership and permissions must be changed on all nodes of the cluster because `/u02/app` must be owned by `grid:oinstall`, otherwise `/u02/app/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

---

```
# chown -R grid:oinstall /u02/app/grid
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
# chmod -R 775 /u02/app/grid
```

- 7 In `/etc/fstab` file add:

```
/dev/vx/dsk/bindg_galaxy/orabinvol \
/u01 vxfs
delaylog 0 2
/dev/vx/rdisk/bindg_galaxy/crsbinvol \
/u02 vxfs
delaylog 0 2
```

- 8 Repeat all the steps on each node of the cluster.

**To create the file system and directories on cluster file system for Oracle Grid Infrastructure and Oracle database**

Perform the following steps on the CVM master node in the cluster.

- 1 As the root user, create a VxVM shared disk group `bindg`:

```
# vxdg -s init bindg Disk_1
```

- 2 Set the activation mode for the disk group on all the nodes:

```
# vxdg -g bindg set activation=sw
```

- 3 From the CVM master node, start the disk group:

```
# vxvol -g bindg startall
```

- 4 Create separate volumes for Oracle Grid Infrastructure (`crsbinvol`) and Oracle database (`orabinvol`):

```
# vxassist -g bindg make crsbinvol 5G
```

```
# vxassist -g bindg make orabinvol 7G
```

**5** Create the following directories for Oracle, `ORACLE_BASE`, `GRID_HOME`, `ORACLE_HOME`.

The file system and directories created on shared storage in this procedure are based on the following layout:

<code>\$ORACLE_BASE</code>	For grid user: <code>/u02/app/grid</code> For Oracle user: <code>/u01/app/oracle</code> <code>/u02/app/grid</code> and <code>/u01/app/oracle</code> are on local storage.
<code>\$GRID_HOME</code>	<code>/u02/app/11.2.0/grid</code> <code>/u02/app/11.2.0</code> is on local storage. <code>/u02/app/11.2.0/grid</code> is on shared storage.
<code>\$ORACLE_HOME</code>	<code>/u01/app/oracle/product/11.2.0/dbhome_1</code> <code>/u01/app/oracle/product/11.2.0</code> is on local storage. <code>/u01/app/oracle/product/11.2.0/dbhome_1</code> is on shared storage.

```
# mkdir -p /u02/app/11.2.0/grid
# mkdir -p /u01/app/oracle
# mkdir -p /u02/app/grid
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
```

**6** Create file systems with the volumes `crsbinvol` and `orabinvol`:

```
# mkfs -F vxfs /dev/vx/rdisk/bindg/crsbinvol
# mkfs -F vxfs /dev/vx/rdisk/bindg/orabinvol
```

**7** Mount the file systems. Perform this step on each node.

```
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/crsbinvol \
/u02/app/11.2.0/grid
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/orabinvol \
/u01/app/oracle/product/11.2.0/dbhome_1
```

## 8 Change the ownership and permissions on all nodes of the cluster.

---

**Note:** The ownership and permissions must be changed on all nodes of the cluster because `/u02/app` must be owned by `grid:oinstall`, otherwise `/u02/app/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

---

```
# chown -R grid:oinstall /u02/app/grid
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
# chmod -R 775 /u02/app/grid
```

## 9 Add the CVMVolDg and CFSSMount resources to the VCS configuration.

See [“To add the CFSSMount and CVMVolDg resources to the VCS configuration using CLI”](#) on page 88.

To add the CFSSMount and CVMVolDg resources to the VCS configuration using CLI

### 1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

### 2 Configure the CVM volumes under VCS:

```
# hares -add crsorabin_voldg CVMVolDg cvm
# hares -modify crsorabin_voldg Critical 0
# hares -modify crsorabin_voldg CVMDiskGroup bindg
# hares -modify crsorabin_voldg CVMVolume -add crsbinvol
# hares -modify crsorabin_voldg CVMVolume -add orabinvol
# hares -modify crsorabin_voldg CVMActivation sw
```



### 3 Set up the file system under VCS:

```
# hares -add crsbin_mnt CFMount cvm

# hares -modify crsbin_mnt Critical 0

# hares -modify crsbin_mnt MountPoint "/u02/app/11.2.0/grid"

# hares -modify crsbin_mnt BlockDevice \
"/dev/vx/dsk/bindg/crsbinvol"

# hares -add orabin_mnt CFMount cvm

# hares -modify orabin_mnt Critical 0

# hares -modify orabin_mnt MountPoint \
"/u01/app/oracle/product/11.2.0/dbhome_1"

# hares -modify orabin_mnt BlockDevice \
"/dev/vx/dsk/bindg/orabinvol"
```

### 4 Link the parent and child resources:

```
# hares -link crsorabin_voldg cvm_clus

# hares -link crsbin_mnt crsorabin_voldg

# hares -link crsbin_mnt vxfsckd

# hares -link orabin_mnt crsorabin_voldg

# hares -link orabin_mnt vxfsckd
```

### 5 Enable the resources:

```
# hares -modify crsorabin_voldg Enabled 1

# hares -modify crsbin_mnt Enabled 1

# hares -modify orabin_mnt Enabled 1

# haconf -dump -makero
```

## 6 Verify the resource configuration in the main.cf file.

```
CFSMount crsbin_mnt (  
    Critical = 0  
    MountPoint = "/u02/app/11.2.0/grid"  
    BlockDevice = "/dev/vx/dsk/bindg/crsbinvol"  
)  
  
CFSMount orabin_mnt (  
    Critical = 0  
    MountPoint = "/u01/app/oracle/product/11.2.0/dbhome_1"  
    BlockDevice = "/dev/vx/dsk/bindg/orabinvol"  
)  
  
CVMVoldg crsorabin_voldg (  
    Critical = 0  
    CVMDiskGroup = bindg  
    CVMVolume = { crsbinvol, orabinvol }  
    CVMActivation = sw  
)  
  
crsbin_mnt requires crsorabin_voldg  
crsbin_mnt requires vxfsckd  
orabin_mnt requires crsorabin_voldg  
orabin_mnt requires vxfsckd  
crsorabin_voldg requires cvm_clus
```

## 7 Verify that the resources are online on all systems in the cluster.

```
# hares -state crsorabin_voldg  
  
# hares -state crsbin_mnt  
  
# hares -state orabin_mnt
```

---

**Note:** At this point, the crsorabin\_voldg resource is reported offline, and the underlying volumes are online. Therefore, you need to manually bring the resource online on each node.

---

To bring the resource online manually:

```
# hares -online crsorabin_voldg -sys galaxy  
  
# hares -online crsorabin_voldg -sys nebula
```

## Setting up user equivalence

You must establish Oracle user and grid user equivalence on all nodes to allow the Oracle Universal Installer to securely copy files and run programs on the nodes in the cluster without requiring password prompts.

Set up passwordless SSH communication between the cluster nodes for the Oracle user and the grid user.

For more information, see the Oracle documentation.

## Synchronizing time settings on cluster nodes

Synchronize the time settings on all cluster nodes. For more information, see the Oracle documentation.

## Editing the user profiles

Edit the Oracle user and grid user profile files to set the paths to ORACLE\_BASE and ORACLE\_HOME on each node.

In the following sample procedure, the shell environment is `ksch`, the Oracle user home directory is `/home/oracle`, and the Oracle grid user home directory is `/home/grid`.

### To edit the Oracle user profile

- 1 As the Oracle user, set the proper environment variables on each node.

```
export ORACLE_BASE=/u01/app/oracle
export ORACLE_HOME=/u01/app/oracle/product/11.2.0/dbhome_1
export GRID_HOME=/u01/grid
export LIBPATH=$ORACLE_HOME/lib:$ORACLE_HOME/lib32:$GRID_HOME/\
lib:$GRID_HOME/lib32:/usr/lib:/lib
export PATH=$PATH:$GRID_HOME/bin:$ORACLE_HOME/bin:/opt/VRTS/bin
export CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME\
/jlib:$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib
```

- 2 Apply the profile changes:

```
$ . /home/oracle/.profile
```

### To edit the Oracle grid user profile

- 1 As the Oracle grid user, set the proper environment variables on each node.

```
export GRID_HOME=/u02/app/11.2.0/grid
export LIBPATH=$ORACLE_HOME/lib:$ORACLE_HOME/lib32:$GRID_HOME/\
lib:$GRID_HOME/lib32:/usr/lib:/lib
export PATH=$PATH:$GRID_HOME/bin:$ORACLE_HOME/bin:/opt/VRTS/bin
export CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME\
/jlib:$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib
```

- 2 Apply the profile changes:

```
$ . /home/grid/.profile
```

## Installing Oracle Grid Infrastructure using the Oracle Universal Installer

This section provides instructions for installing the Oracle Grid Infrastructure software using the Oracle Universal Installer. The software is installed on each node in the Oracle Grid Infrastructure home directory.

### To install Oracle Grid Infrastructure using the Oracle Universal Installer

- 1 Log in as the Oracle grid user. On the first node, set the DISPLAY variable.

- For Bourne Shell (bash), type:

```
$ DISPLAY=10.20.12.150:0.0;export DISPLAY
```

where 10.20.12.150 is the IP address of X client where you want to export the display for the installer.

- For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY 10.20.12.150:0.0
```

where 10.20.12.150 is the IP address of X client where you want to export the display for the installer.

- 2 Start the Oracle Universal Installer on the first node.

```
$ cd /dvd_mount
```

```
$ ./runInstaller
```

- 3 Enter the following information when prompted by the Oracle Universal Installer:

- Select installation option  
Select the option **Install and Configure Grid Infrastructure for a Cluster**.
- Select installation type  
Select the option **Advanced Installation**.
- Specify cluster configuration  
Enter the SCAN name for the cluster that will be used by the database clients to connect to databases within the cluster.  
Review the list of cluster nodes displayed.
- Grid Plug and Play information  
Provide the following information:
  - Name of the cluster
  - SCAN name  
The SCAN address on the domain name server (DNS) must resolve to three addresses (recommended) or at least one address.
  - SCAN port
- Specify network interface usage  
Identify the planned use for each interface: Public, Private, or Do Not use.

---

**Note:** Make sure that the same private interfaces that you specified at the time of configuring PrivNIC and MultiPrivNIC are listed on the screen.

---

---

**Note:** Mark the interfaces for the subnet containing the private IP addresses managed by the PrivNIC/MultiPrivNIC agents as 'Private'.

---

The interfaces that are **Private** are stored in GPNP profile as a 'cluster\_interconnect' for Oracle Clusterware communication and database cache fusion traffic.

- Storage option information  
Select the option **Shared File System**.
- OCR storage option  
Enter the full path of the location where you want to store the OCR information.  
For example, if you are storing the OCR information on CFS, enter:  
`/ocrvote/ocr.`

---

**Note:** Select the option **External Redundancy**. OCR mirroring is performed by CVM.

---

- Voting Disk storage option

Enter the full path of the location where you want to store the voting disk information.

For example, if you are storing the voting disk information on CFS, enter:

```
/ocrvote/vote
```

---

**Note:** Select the option **External Redundancy**. Voting disk redundancy is provided by CVM.

---

- Specify installation location

Enter the full path to the Oracle base directory and the Oracle Grid Infrastructure home directory.

- Create inventory

Enter the full path to the Oracle inventory directory where you want to store the installation files.

- 4 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle Grid Infrastructure installation.
- 5 Run the `oraInstRoot.sh` script as prompted by the Oracle Universal Installer.
- 6 Run the `root.sh` script on each node as prompted by the Oracle Universal Installer:

```
# cd $GRID_HOME
```

```
# ./root.sh
```

The Oracle Clusterware daemons are started on the node.

---

**Note:** The Oracle Cluster Verification Utility fails during the installation of the Oracle Grid Infrastructure software. You may ignore this error. This is a known issue with Oracle.

---

## Configuring LLT links in the GPnP profile

Update the GPnP profile to include the remaining LLT links that were not added to the profile during the Oracle Grid Infrastructure installation.

### To configure the LLT links in the GPnP profile

- 1 View the currently configured interfaces:

```
# $GRID_HOME/bin/oifcfg getif
lan0 10.2.156.0          global      public
lan1 192.168.12.0       global      cluster_interconnect
```

The interfaces that are currently stored in the GPnP profile, their subnets, and their role (public or cluster\_interconnect) are displayed.

- 2 Add the remaining LLT links to the GPnP profile:

```
# $GRID_HOME/bin/oifcfg setif -global \
lan2/192.168.12.0:cluster_interconnect
```

If you are using multiple IP addresses on different subnet for cluster interconnect (for load balancing), add the remaining interface subnets to the GPnP profile.

```
# $GRID_HOME/bin/oifcfg setif -global \
lan2/192.168.2.0:cluster_interconnect
# $GRID_HOME/bin/oifcfg setif -global \
lan1/192.168.2.0:cluster_interconnect
```

- 3 Verify that the correct interface subnet is in use:

```
# $GRID_HOME/bin/oifcfg getif
lan0 10.2.156.0          global      public
lan1 192.168.12.0       global      cluster_interconnect
lan2 192.168.12.0       global      cluster_interconnect
lan1 192.168.2.0        global      cluster_interconnect
lan2 192.168.2.0        global      cluster_interconnect
```

Make sure all the LLT links are configured and listed in the GPnP profile.

## Installing the Oracle RAC database using the Oracle Universal Installer

The following procedure describes how to install the Oracle RAC database using the Oracle Universal Installer. Symantec recommends that you install the Oracle RAC database locally on each node.

### To install Oracle RAC database using the Oracle Universal Installer

- 1 Log in as the Oracle user. On the first node, set the DISPLAY variable.

- For Bourne Shell (bash), type:

```
$ DISPLAY=10.20.12.150:0.0;export DISPLAY
```

- For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY 10.20.12.150:0.0
```

- 2 Start the Oracle Universal Installer.

```
$ cd /dvd_mount
```

```
$ ./runInstaller
```

- 3 Enter the following information when prompted by the Oracle Universal Installer:

Select installation option	Select the option <b>Install database software only</b> .
Node selection	Select <b>Real Application Clusters database installation</b> . Select the nodes on which the Oracle RAC database software must be installed.
Select database edition	Select <b>Enterprise Edition</b> .
Specify installation location	Review or enter the ORACLE_BASE and ORACLE_HOME directory paths.

The Oracle Universal Installer runs product-specific prerequisite checks. Any items that are flagged must be manually checked and configured.

- 4 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle database installation.
- 5 Run the root.sh script as prompted by the Oracle Universal Installer.

```
# cd $ORACLE_HOME
```

```
# ./root.sh
```



## Relinking the SF Oracle RAC libraries with Oracle RAC manually

If you added or upgraded the Oracle patches, you must relink the SF Oracle RAC libraries to Oracle. You must enable Oracle to use the Veritas ODM and cluster membership (VCSMM) libraries.

Symantec recommends that you relink the SF Oracle RAC libraries only after completing all the required patch additions.

To relink the SF Oracle RAC libraries with Oracle RAC, run the following command as the oracle user:

```
$ /opt/VRTSvcs/rac/bin/linkrac oracle_version
```

Where `oracle_version` is 10gR2, 11gR1, or 11gR2 depending on the version of Oracle RAC you installed.

For example, to relink the SF Oracle RAC libraries with Oracle RAC 10g Release 2, run the following command as the oracle user:

```
$ /opt/VRTSvcs/rac/bin/linkrac 10gR2
```

## Configuring the CSSD resource manually

Add the `cssd` resource to the VCS configuration and set CSSD dependencies on the resources that manage OCR and voting disk and the private IP addresses for Oracle Clusterware.

---

**Note:** It is recommended that the OCR, voting disk, and PrivNIC/MultiPrivNIC resources be configured in the same VCS group as that of the `cssd` resource. If the resources are not in the same group, set the appropriate dependencies between the service groups.

---

### To configure the CSSD resource

- 1 Change the permission on the VCS configuration file to read-write mode:

```
# haconf -makerw
```

- 2 Add the CSSD resource to the `cvm` group:

```
# hares -add cssd Application cvm
```

**3** Modify the CSSD resource attributes:

```
# hares -modify cssd StartProgram /opt/VRTSvcs/rac/bin/cssd-online
# hares -modify cssd StopProgram /opt/VRTSvcs/rac/bin/cssd-offline
# hares -modify cssd MonitorProgram /opt/VRTSvcs/rac/bin/cssd-monitor
# hares -modify cssd CleanProgram /opt/VRTSvcs/rac/bin/cssd-clean
# hares -modify cssd Critical 0
# hares -override cssd OnlineWaitLimit
# hares -modify cssd OnlineWaitLimit 5
```

**4** Enable the CSSD resource:

```
# hares -modify cssd Enabled 1
```

**5** Set the dependency of the CSSD resource on the CFMount or CVMVoldg resources that manage OCR and voting disk.

If you configured OCR and voting disk on CVM raw volumes:

```
# hares -link cssd ocr_voldg_ocrvotedg
# hares -link cssd vote_voldg_ocrvotedg
```

If you configured OCR and voting disk on CFS:

```
# hares -link cssd ocrvote_mnt_ocrvotedg
```

**6** Set the dependency of the CSSD resource on the PrivNIC or MultiPrivNIC resources that manage the private IP address for Oracle Clusterware.

If you configured the PrivNIC resource:

```
# hares -link cssd ora_priv
```

If you configured the MultiPrivNIC resource:

```
# hares -link cssd multi_priv
```

- 7 If the Oracle Clusterware and the Oracle database binaries are on CFS, set the dependencies between the CSSD resource and the CFSSMount resources for the binaries manually:

```
# hares -link cssd crsbin_mnt
# hares -link cssd orabin_mnt
```

- 8 Change the permission on the VCS configuration file to read-only mode:

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

## Preventing automatic startup of Oracle Clusterware

The use of the CSSD agent is mandatory to ensure adequate handling of service group inter-dependencies and thereby prevent the premature startup of Oracle Clusterware. Therefore, disable automatic startup of Oracle Clusterware when the system starts.

### To prevent automatic startup of Oracle Clusterware

- 1 Log in as the root user on each node in the cluster.
- 2 Disable automatic startup of Oracle Clusterware:

```
# $GRID_HOME/bin/crsctl disable crs
```

Where \$GRID\_HOME is the Oracle Grid Infrastructure home directory.

## Creating the Oracle RAC database

Create the Oracle RAC database on shared raw volumes or cluster file systems.

To create the requisite database storage, see the instructions in the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide (5.0.1)*.

For instructions on creating the database, see the Oracle RAC documentation.

## Preventing automatic database startup

Configure the Oracle RAC database for manual startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent. Before you configure the VCS service groups for Oracle, you need to prevent the Oracle database from starting automatically. The Oracle Clusterware and Oracle agent may attempt to start the database instance at the same time if the database mount is available. To prevent the Oracle database from starting automatically, you must change the management policy for the database from automatic to manual using the Oracle

`SRVCTL` command. The command changes the `AUTO_START` attribute of the Oracle database and instance resources.

### To prevent automatic database startup

- 1 Register the database, if not already registered:

```
$ srvctl add database -d db-name -o oracle_home \  
-p location-of-parameterfile -y manual
```

- 2 Once the database is registered, change the management policy for the database to manual:

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

- 3 Start the database:

```
$ srvctl start database -d db-name
```

## Configuring VCS service groups for Oracle RAC

You can set up the Oracle database to be managed by VCS or by Oracle Clusterware. Symantec recommends that the Oracle database be configured under VCS.

---

**Note:** The VCS Oracle agent is supported only for administrator-managed databases. For policy-managed databases, use the service group configuration without the VCS Oracle agent.

---

When the database is configured under VCS:

- You can choose to configure the service group in a way that insulates all the databases from failure in any of the databases in the group.
- VCS manages the start and stop sequence of the applications and the database.

[Figure A-1](#) illustrates a service group configuration with the VCS Oracle agent.

Figure A-1 Service group configuration with the VCS Oracle agent

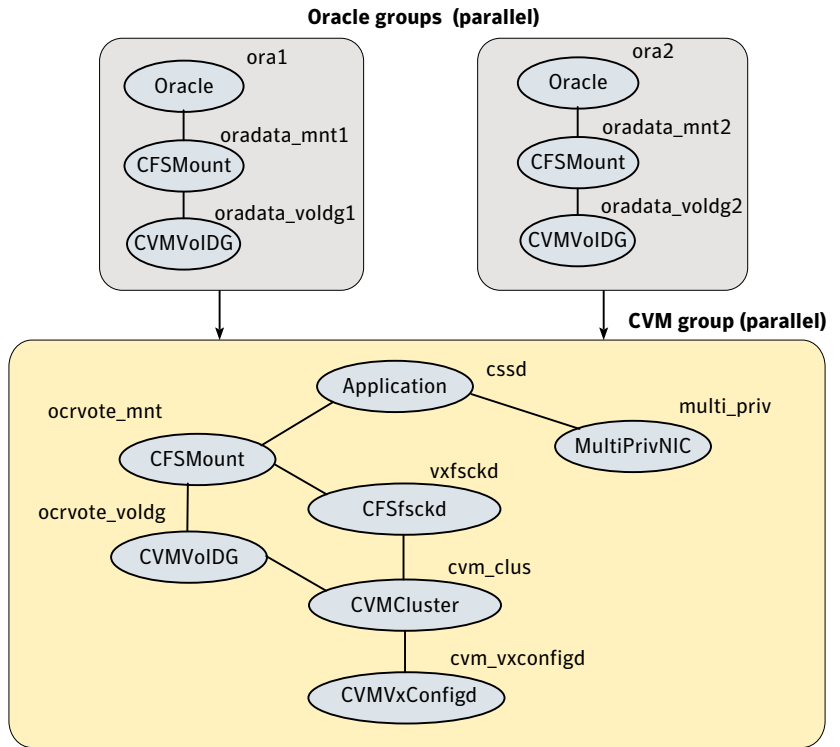



Figure A-2 illustrates a service group configuration without the VCS Oracle agent.

**Figure A-2** Service group configuration without the VCS Oracle agent



In a service group configuration without the VCS Oracle agent, Oracle Clusterware controls the database. An online local firm dependency exists between the Oracle Clusterware group and the CVM group. When the system starts, the CVM group brings up the volume and mount points for the databases. The Oracle Clusterware group brings up the OCR and voting disk, configures the private IP address for Oracle Clusterware, and starts Oracle Clusterware. Oracle Clusterware starts the database and the application is brought online.

---

**Note:** In a service group configuration without the VCS Oracle agent, when the system starts, all volumes and mount points **MUST** be online for the dependent service groups to be online.

---

To configure the Oracle database under VCS, create Oracle service groups after installing Oracle RAC and creating a database.

You can create Oracle service groups by modifying the VCS configuration file in one of the following ways:

- Edit the main.cf file
  - See [“Editing the main.cf file to configure VCS service groups for Oracle RAC”](#) on page 103.

- Use the CLI  
See [“Using the CLI to configure VCS service groups for Oracle RAC”](#) on page 104.

## Editing the main.cf file to configure VCS service groups for Oracle RAC

This section describes how to manually edit the `/etc/VRTSvcs/conf/config/main.cf` file to configure VCS service groups for Oracle RAC.

---

**Note:** Before you edit the main.cf file, change the cluster configuration to read-write mode: `# haconf -makerw`

---

You need to perform the steps in the following procedures:

- Configure VCS service groups for Oracle RAC  
See [“To configure VCS service groups for Oracle RAC”](#) on page 103.
- Start VCS after modifying the configuration file  
See [“To start VCS after modifying the configuration file”](#) on page 104.

### To configure VCS service groups for Oracle RAC

- 1 Log in as the root user to one of the systems.
- 2 As root user, save your existing configuration before you modify main.cf:

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

If the configuration is not writable, a warning appears: "Cluster not writable."  
You may safely ignore the warning.

- 3 Stop the VCS engine on all systems and leave the resources available:

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

- 4 Back the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
```

```
# cp main.cf main.orig
```

- 5 Using `vi` or another text editor, edit the main.cf file to add the Oracle service groups.

6 Save and close the main.cf file.

7 Verify the syntax of the file :

```
# cd /etc/VRTSvcs/conf/config  
  
# hacf -verify .
```

For additional information and instructions on modifying the VCS configuration by editing the main.cf file:

See the *Veritas Cluster Server User's Guide*.

### Starting VCS after modifying the configuration file

You need to start VCS for the new configuration to take effect.

#### To start VCS after modifying the configuration file

1 Start VCS on the node on which you modified the configuration file and check its status:

```
# hstart  
  
# hastatus
```

2 When "LOCAL\_BUILD" is listed in the message column, start VCS on the other systems:

```
# hstart
```

3 On one of the systems, enter the following command to verify that the service groups for Oracle RAC are brought online:

```
# hagr -display
```

4 Check the status of the groups.

```
# hagr -state
```

5 Check the status of the resources.

```
# hares -state
```

### Using the CLI to configure VCS service groups for Oracle RAC

This section describes how to configure the Oracle service group using the CLI.



The following procedure assumes that you have created the database.

**To configure the Oracle service group using the CLI**

- 1 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the service group to the VCS configuration:

```
# hagrps -add oradb1_grp
```

- 3 Modify the attributes of the service group:

```
# hagrps -modify oradb1_grp Parallel 1
```

```
# hagrps -modify oradb1_grp SystemList galaxy 0 nebula 1
```

```
# hagrps -modify oradb1_grp AutoStartList galaxy nebula
```

- 4 Add the CVMVolDg resource for the service group:

```
# hares -add oradata_voldg CVMVolDg oradb1_grp
```

- 5 Modify the attributes of the CVMVolDg resource for the service group:

```
# hares -modify oradata_voldg CVMdiskGroup oradatadg
```

```
# hares -modify oradata_voldg CVMActivation sw
```

```
# hares -modify oradata_voldg CVMVolume oradatavol
```

- 6 Add the CFMount resource for the service group:

```
# hares -add oradata_mnt CFMount oradb1_grp
```

- 7 Modify the attributes of the CFMount resource for the service group:

```
# hares -modify oradata_mnt MountPoint "/oradata"
```

```
# hares -modify oradata_mnt BlockDevice \
```

```
"/dev/vx/dsk/oradatadg/oradatavol"
```

- 8 Add the Oracle RAC database instance to the service group:

```
# hares -add oral Oracle oradb1_grp
```

**9** Modify the attributes of the Oracle resource for the service group:

```
# hares -modify ora1 Owner oracle
# hares -modify ora1 Home "/app/oracle/orahome"
# hares -modify ora1 StartUpOpt SRVCTLSTART
# hares -modify ora1 ShutDownOpt SRVCTLSTOP
```

**10** Localize the Sid attribute for the Oracle resource:

```
# hares -local ora1 Sid
```

**11** Set the Sid attributes for the Oracle resource on each system:

```
# hares -modify ora1 Sid vrts1 -sys galaxy
# hares -modify ora1 Sid vrts2 -sys nebula
```

**12** Set the dependencies between the CFSSMount resource and the CVMVolDg resource for the Oracle service group:

```
# hares -link oradata_mnt oradata_voldg
```

**13** Set the dependencies between the Oracle resource and the CFSSMount resource for the Oracle service group:

```
# hares -link ora1 oradata_mnt
```

**14** Create an online local firm dependency between the oradb1\_grp service group and the cvm service group:

```
# hagrps -link oradb1_grp cvm online local firm
```

**15** Enable the Oracle service group:

```
# hagrps -enableresources oradb1_grp
```

**16** Change the cluster configuration to the read-only mode:

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

**17** Bring the Oracle service group online on all the nodes:

```
# hagrps -online oradb1_grp -any
```

For more information and instructions on configuring the service groups using the CLI:

See the *Veritas Cluster Server User's Guide*.

## Location of VCS log files

You may want to review the log files at `/var/VRTSvcs/log/engine_A.log` for errors or status messages. When large amounts of data are written, multiple log files may be written, such as `engine_B.log`, `engine_C.log`, and so on. The `engine_A.log` contains the most recent data.

# Upgrading to Oracle RAC 11g Release 2

This section provides instructions on upgrading to Oracle RAC 11g Release 2.

---

**Note:** Starting with Oracle RAC 11g Release 2, the Oracle Clusterware software must be upgraded to run in a new directory called the Oracle Grid Infrastructure home. If you plan to upgrade ASM to Release 11g R2, make sure that ASM also resides in the same Oracle Grid Infrastructure home directory.

Even if the existing Oracle Clusterware software is installed in a shared home directory, you can use a local home directory for installing the Oracle Grid Infrastructure software.

---

For information on supported upgrade paths:

See “[Supported upgrade paths](#)” on page 107.

Complete the following steps to upgrade to Oracle RAC 11g Release 2:

1. Prepare to upgrade Oracle RAC.  
See “[Preparing to upgrade Oracle RAC](#)” on page 108.
2. Upgrade Oracle RAC binaries.  
See “[Upgrading Oracle RAC binaries](#)” on page 108.
3. Migrate the Oracle RAC database.  
See “[Migrating the Oracle RAC database](#)” on page 109.
4. Complete the post-upgrade tasks.  
See “[Performing post-upgrade tasks](#)” on page 109.

## Supported upgrade paths

[Table A-2](#) lists the upgrade paths for Oracle RAC.

**Table A-2** Supported upgrade paths for Oracle RAC

From current version	Upgrade to
Oracle RAC 10g Release 2	Oracle RAC 11g Release 2
Oracle RAC 11g Release 1	Oracle RAC 11g Release 2

**Note:** When you upgrade to a different version of Oracle RAC, make sure that the full path of the Oracle Grid Infrastructure home directory and the Oracle database home directory is different from the path where the existing version of Oracle RAC resides.

The upgrade procedure assumes that the beginning configuration includes the following components, and that these components are running on the cluster nodes:

- SF Oracle RAC 5.0.1 RP2
- A supported version of the operating system

## Preparing to upgrade Oracle RAC

Perform the following pre-upgrade tasks before upgrading Oracle RAC.

### To prepare for upgrading Oracle RAC

- 1 Take a hot or cold backup of the existing database.
- 2 Back up the existing Oracle home and central inventory.
- 3 If the Oracle RAC database is under VCS control, freeze the Oracle service groups to prevent VCS from reporting the resource as faulted when Oracle RAC stops and starts the database during the upgrade:

```
# haconf -makerw
# hagrpl -freeze oracle_group -persistent
# haconf -dump -makero
```

## Upgrading Oracle RAC binaries

Review your Oracle installation manuals and the appropriate Oracle support Web sites before upgrading Oracle RAC.

---

**Note:** Make sure that you upgrade to Oracle RAC 11g Release 2 with the same user credentials as that of the existing installation.

---

### To upgrade Oracle RAC binaries

- 1 Upgrade Oracle Clusterware to a new directory called the Oracle grid infrastructure home (GRID\_HOME).

---

**Note:** Starting with Oracle RAC 11g Release 2, ASM must reside in the Oracle Grid Infrastructure home directory. If you plan to upgrade ASM to Release 11g R2, make sure that you upgrade it to run in the Oracle Grid Infrastructure home directory.

---

For instructions, see the Oracle RAC documentation.

- 2 Make sure that Oracle Clusterware is running.
- 3 Install the Oracle RAC database binaries.

For instructions, see the Oracle RAC documentation.

- 4 Relink the SF Oracle RAC libraries with Oracle RAC.

See [“Relinking the SF Oracle RAC libraries with Oracle RAC manually”](#) on page 97.

## Migrating the Oracle RAC database

For instructions on migrating the existing Oracle RAC database, see the Oracle metalink documentation.

After migrating the database, complete the post-upgrade tasks:

See [“Performing post-upgrade tasks”](#) on page 109.

## Performing post-upgrade tasks

Perform the steps in the following procedure to complete the upgrade.

### To perform post-upgrade tasks

- 1 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Modify the Oracle RAC configuration to prevent automatic startup of Oracle Clusterware.

See [“Preventing automatic startup of Oracle Clusterware”](#) on page 99.

- 3 Modify the Oracle RAC database configuration to prevent automatic database startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent.

See [“Preventing automatic database startup”](#) on page 99.

- 4 Unfreeze the VCS service groups that were frozen earlier.

As root user, enter:

```
# hagrps -unfreeze oracle_group -persistent
```

```
# hagrps -unfreeze cvm_group -persistent
```

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