

Veritas Storage Foundation and High Availability Solutions Release Notes

Solaris

5.1 Rolling Patch 1



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

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

Product version: 5.1 RP1

Document version: 5.1RP1.0

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.

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

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

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

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

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

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

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

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Managed Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Education Services	Education Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about enterprise services, please visit our web site at the following URL:

www.symantec.com/business/services/

Select your country or language from the site index.

Contents

Technical Support	4
Chapter 1 Release Notes	11
Introduction	12
System Requirements	12
Supported Solaris operating systems	12
Database requirements	13
List of products	14
Fixed issues	14
Veritas Volume Manager fixed issues in 5.1 RP1 release	15
Veritas File System fixed issues in 5.1 RP1 release	16
Veritas Storage Foundation fixed issues in 5.1 RP1	19
Veritas Storage Foundation Cluster File System fixed issues in 5.1 RP1 release	20
Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 RP1	20
Veritas Cluster Server fixed issues in 5.1 RP1	21
Veritas Cluster Server agents for Veritas Volume Replicator fixed issues in 5.1 RP1	22
Storage Foundation Manager fixed issues in 5.1 RP1	23
VEA fixed issues in 5.1 RP1	25
Known issues	26
Veritas Storage Foundation known issues in 5.1 RP1 release	26
Veritas Volume Manager known issues in 5.1 RP1 release	28
Veritas File System known issues in 5.1 RP1 release	28
Veritas Storage Foundation Cluster File System known issues in 5.1 RP1 release	29
Veritas Storage Foundation for Oracle RAC known issues in 5.1 RP1	29
Veritas Cluster Server known issues in 5.1 RP1	30
Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 RP1	30
Software limitations	30
Veritas Storage Foundation software limitations in 5.1 RP1 release	31

Veritas Volume Manager software limitations in 5.1 RP1 release	31
Veritas Storage Foundation for Oracle RAC software limitations in 5.1 RP1	32
Changes in Storage Foundation High Availability	32
About the new installrp script	33
CVM master node needs to assume the logowner role for VCS managed VVR resources	35
Downloading the rolling patch archive	36
List of patches	36
Installing the Veritas software for the first time	38
Installing with JumpStart	40
Overview of JumpStart installation tasks	41
Generating the finish scripts	41
Preparing installation resources	46
Adding language pack information to the finish file	47
Installing 5.1 RP1 using the web-based installer	48
About the Web-based installer	48
Before using the Veritas Web-based installer	49
Starting the Veritas Web-based installer	49
Obtaining a security exception on Mozilla Firefox	49
Performing a pre-installation check with the Veritas Web-based installer	50
Installing products with the Veritas Web-based installer	50
Prerequisites for upgrading to 5.1 RP1	51
Supported upgrade paths	52
Upgrading 5.1 to 5.1 RP1	52
Performing a full upgrade to 5.1 RP1 on a cluster	52
Performing a phased upgrade to 5.1 RP1 on a cluster	65
Upgrading Veritas product with the Veritas Web-based installer	81
Performing a rolling upgrade using the installer	82
Performing a rolling upgrade manually	84
Upgrading to 5.1 RP1 on a standalone system	90
Verifying software versions	92
Removing and rolling back	92
Removing 5.1 RP1 from Veritas Cluster Server	92
Removing 5.1 RP1 on SF or SFCFS	97
Removing 5.1 RP1 on Storage Foundation for Oracle RAC	100
Documentation addendum	104
Disk agent	104
Using the preonline_vvr trigger for RVGLogowner resources	106

Appendix A	Installing or upgrading to Oracle RAC 11g Release 2	107
	Installing Oracle RAC 11g Release 2	107
	Preparing to install Oracle RAC using the SF Oracle RAC installer or manually	108
	Installing Oracle Grid Infrastructure using the Oracle Universal Installer	118
	Configuring LLT links in the GPnP profile	121
	Installing the Oracle RAC database using the Oracle Universal Installer	122
	Linking the ODM library	123
	Configuring the CSSD resource manually	124
	Preventing automatic startup of Oracle Clusterware	125
	Creating the Oracle RAC database	126
	Preventing automatic database startup	126
	Configuring VCS service groups for Oracle RAC	127
	Upgrading to Oracle RAC 11g Release 2	133
	Supported upgrade paths	133
	Preparing to upgrade Oracle RAC	134
	Upgrading Oracle RAC binaries	135
	Migrating the Oracle RAC database	135
	Performing post-upgrade tasks	135

Release Notes

This chapter includes the following topics:

- [Introduction](#)
- [System Requirements](#)
- [List of products](#)
- [Fixed issues](#)
- [Known issues](#)
- [Software limitations](#)
- [Changes in Storage Foundation High Availability](#)
- [Downloading the rolling patch archive](#)
- [List of patches](#)
- [Installing the Veritas software for the first time](#)
- [Installing with JumpStart](#)
- [Installing 5.1 RP1 using the web-based installer](#)
- [Prerequisites for upgrading to 5.1 RP1](#)
- [Supported upgrade paths](#)
- [Upgrading 5.1 to 5.1 RP1](#)
- [Verifying software versions](#)
- [Removing and rolling back](#)
- [Documentation addendum](#)

Introduction

This document provides information about the Storage Foundation and High Availability Solutions 5.1 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/334829>

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

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

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

System Requirements

This section describes the system requirements for this release

Supported Solaris operating systems

This release of the Veritas products is supported on the following Solaris operating systems:

- Solaris 9 (SPARC Platform 32-bit and 64-bit) with Update 8 or later
- Solaris 10 (SPARC or x64 Platform 64-bit) with Update 6 or later

If necessary, upgrade Solaris before you install the Veritas products.

Install all the latest required Solaris patches listed in this *Release Notes*.

See “[Required Solaris patches](#)” on page 12.

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

<http://entsupport.symantec.com/docs/334829>

Required Solaris patches

Before installing Veritas SFHA, ensure that the correct Solaris patches are installed.

See <http://sunsolve.sun.com> for the latest Solaris patch updates.

The following patches (or a later revision of those patches) are required for Solaris SPARC:

Table 1-1 Solaris SPARC patches

Operating system	Sun patch number
Solaris 9	114477-04 122300-29 - required for Live Upgrade
Solaris 10	119254-06 119042-02 125731-02 128306-05 127111-01

The following patches (or a later revision of those patches) are required for Solaris x64:

Table 1-2 Solaris x64 patches

Operating system	Sun patch number
Solaris 10	118344-14 118855-36 119043-11 119131-33 120012-14 125732-05 127128-11

Database requirements

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

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

Note: SF and SFCFS support running Oracle, DB2, and Sybase on VxFS and VxVM. SF and SFCFS do not support running SFDB tools with DB2 and Sybase.

List of products

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

- Veritas Storage Foundation (SF)
- Veritas Storage Foundation and High Availability (SFHA)
- Veritas Storage Foundation Cluster File System (SFCFS)
- Veritas Storage Foundation Cluster File System and High Availability (SFCFSHA)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)
- Veritas Volume Manager (VM)
- Veritas File System (FS)
- Veritas Cluster Server (VCS)

Fixed issues

The following sections describe the Veritas Storage Foundation High Availability issues that were fixed in this release.

- [Veritas Storage Foundation fixed issues in 5.1 RP1](#)
- [Veritas Volume Manager fixed issues in 5.1 RP1 release](#)
- [Veritas File System fixed issues in 5.1 RP1 release](#)
- [Veritas Storage Foundation Cluster File System fixed issues in 5.1 RP1 release](#)
- [Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 RP1](#)
- [Veritas Cluster Server fixed issues in 5.1 RP1](#)
- [Veritas Cluster Server agents for Veritas Volume Replicator fixed issues in 5.1 RP1](#)
- [Storage Foundation Manager fixed issues in 5.1 RP1](#)
- [VEA fixed issues in 5.1 RP1](#)

Veritas Volume Manager fixed issues in 5.1 RP1 release

Table 1-3 Veritas Volume Manager 5.1 RP1 fixed issues

Fixed issues	Description
1972852, 1972848	vxconfigd dumped core in dg_config_compare() while upgrading to 5.1.
1955693	VxVM 5.0MP3RP3 patch 122058-13 disables vxfsldlic service and prevents boot multi-user mode after jumpstart
1938484	EFI: Prevent multipathing don't work for EFI disk
1937841	VxVM: checkin the fmrshowmap utility
1915356	I/O stuck in vxvm caused cluster node panic
1935297	vxconfigd dumps core in get_prop()
1907796	Corrupted Blocks in Oracle after Dynamic LUN expansion and vxconfigd core dump
1901827	vxvg move failed silently and drops disks.
1899688	[VVR] Every I/O on smartsync enabled volume under VVR leaks memory
1889747	vxlustart customer is unable to do live upgrade with Solaris Zone on vxfs
1886007	vxconfigd lose license information, vxesd leaking File descriptors
1884070	When running iotest on volume, primary node runs out of memory
1881336	VVR: Primary Panic in vol_ru_replica_sent()
1872743	Layered volumes not startable due to duplicate rid in vxrecover global volume list.
1870049	Dump device changed to none after boot disk encapsulation
1860892	Cache Object corruption when replaying the CRECs during recovery
1857729	CVM master in the VVR Primary cluster panic when rebooting the slave during VVR testing
1850166	vxvm vxdisk error v-5-1-8643 device 0_bpcs001_fra: resize failed:
1846165	Data corruption seen on cdsdisks on Solaris-x86 in several customer cases
1840832	vxrootadm does not update the partition table while doing a grow operation

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

Fixed issues	Description
1840673	After adding new luns one of the nodes in 3 node CFS cluster hangs
1835139	CERT : pnate test hang I/O greater than 200 seconds during the filer giveback
1834848	TP:Solaris:reclamation causes data corruption
1826088	After pulling out FC cables of local site array, plex became DETACHED/ACTIVE
1825516	Unable to initialize and use ramdisk for VxVM use
1825270	Need for dmp_revive_paths(in dmp reconfiguration/restore_demon code path.
1792795	supportability feature/messages for plex state change, DCO map clearance, usage of fast re-sync by vxplex
1766452	VVR: VRAS: AIX: vradmind dumps core during collection of memory stats.
1664952	Refreshing private region structures degrades performance during "vxdisk listtag" on a setup of more than 400 disks.
1528160	An ioctl interrupted with EINTR causes frequent vxconfigd exit()'s on 4.1MP4RP3
1479735	CVR: I/O hang on slave if master (logowner) crashes with DCM active.

Veritas File System fixed issues in 5.1 RP1 release

Table 1-4 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number)

Fixed issues	Description
1979429	mount usage error message is thrown while using from /opt/VRTS/bin
1978029	fsadm -R returns EFAULT in certain scenarios.
1976287	conform.backup test is skipped even if no local zone is available.
1973739	vxdisk reclaim fails on dg version 140 works only with ver 150-Cannot reclaim space on dg/vols created on 5.0MP3 after 5.1 upgrade.
1973539	Wrong boundary for reclamation on Hitachi AMS2000 Series

Table 1-4 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number) (*continued*)

Fixed issues	Description
1972882	use separate structures for ioctl interfaces and CFS messages
1972207	full fsck is very slow on an fs with many ilist holes.
1969334	CFS-Command "vxupgrade" test failed.
1967027	LM-conform test odm hits an assert of "f:fdd_advreload:2"
1961790	LM.CMDS->fsck->full->scripts->fextop_12 fails
1960436	CFS.Comform->revnlookup hit "vx_msgprint" via "vx_cfs_iread" on the slave node
1958198	cfs odm stress/noise tests failed due to "bcmp error"
1957365	CFS-Conformance test failed
1957296	CFS-Conformance-Reconfig test hit assert "f:vx_validate_cistat:3"
1957043	LM -conformance/fcl/fcl_fsetquota.3 is failing
1957035	CFS cmds:fsck is failing
1957032	fsqa lm vxmssnap.9 test fails
1956926	cfs-cmds aborting due to fsck, mount, fsted, libtst 64-bit binaries
1954897	CFS-Conformance test hit assert "f:vx_mark_fset_clean:2"
1953913	LM-Command "vxedquota" test failed.
1952827	LM / CFS - Cmds-> alerts test failed.
1952818	LM / CFS -Cmds-> vxtunefs test failed.
1949962	Fix the vxrsh and vxproxyrshd processes for cfs reconfig testing.
1949077	cfs.conform.dbed hits assert "..f:vx_imap_process_inode:4a". by "..vx_workitem_process".
1948451	kernel-conform "sunppriv" and "getattr" tests are missing
1947359	mkdstfs fails to add new volumes
1947356, 1883938	Due to incorrect Makefile 'make clobber' is removing mkdstfs

Table 1-4 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number) (*continued*)

Fixed issues	Description
1946442	tot build setup machine running LM-cmds -> fspadm got failures.
1946433	Enhance mkdstfs for explicitly selecting the purpose of volumes data vs metadata
1946431	mkdstfs only adds to existing volume tags
1946134	fsadm keeps relocating and copying already relocated and copied reorg-ed regions of a file in subsequent passes.
1944283	fcl close is not happening properly
1943116	mkdstfs uses wrong perl instead of /opt/VRTS/bin/perl
1943116	mkdstfs uses wrong perl instead of /opt/VRTS/bin/perl
1940870	Documentation/Test discrepancies
1940409	CFS support for cached ODM
1940390	Cached ODM needs improvements for async requests codm
1934107, 1891400	[VxFS][281-815-793]SLES10 VxFS 5.0MP3 - Incorrect ACL inheritance
1934103	[PRI-1] Performance issues with mmap VxFS 4.1MP4 RP2
1934101	cfs Test cfs-stress-enterprise hit the same assert :f:vx_cwfrz_wait:2
1934098, 1860701	clone removal can block resive ops
1934096, 1746491	NASGW:core dump of fsvmap.
1934095, 1838468	Data page fault at vx_qiostats_update due to fiostats structure already free'd
1934094, 1846461	[VxFS] Customer requests vxfsstat metrics to monitor UNHASHED entries in the dcache
1934085, 1871935	secondaries ias_ilst not updated fully.
1933975, 1844833	fsadm shrink fs looping in vx_reorg_emap due to VX_EBMAPMAX from vx_reorg_enter_zfod

Table 1-4 Veritas File System 5.1 RP1 fixed issues (listed incident number/parent number) (*continued*)

Fixed issues	Description
1933844	[VxFS 5.0MP2RP4] [281-803-975] bad mutex panic in VxFS
1933798	File System with Disk Layout Version 4 and 5 cannot be mounted on VxFS 5.1
1933635, 1914625	[VxFS] Behavior of DST Access age-based file placement policy with preferred files
1931973	fsppadm gives spurious messages when run from multiple CFS nodes found only from 5.1 onwards
1908776	[VxFS][320-219-830] UX.vxfs mount: ERROR: V-3-22168: Cannot open portal device...
1906521	CFS-conform/quotas test hit assert vx_populate_pnq via vx_detach_fset
1902241	9-15a driver regression observed on SFCFSORA TPCC test
1897458, 1805046	wrong alert generation from vxfs when file system usage threshold is set
1895454	Sol10x86 lm.conform->ts some TCs fail
1878583	CFS: getattr call optimization to speedup the case when binaries are being mmapped from many nodes on CFS.

Veritas Storage Foundation fixed issues in 5.1 RP1

Table 1-5 Veritas Storage Foundation fixed issues in 5.1 RP1

Fixed issues	Description
1974086	reverse_resync_begin fails after successfully unmount of clone database on same node when primary and secondary host names do not exactly match.
1940409, 471276	Enhanced support for cached ODM
1901367, 1902312	dbed_vmclonedb failed to umount on secondary server after a successful VM cloning in RAC when the primary SID string is part of the snapplan name.

Table 1-5 Veritas Storage Foundation fixed issues in 5.1 RP1 (*continued*)

Fixed issues	Description
1896097	5.1 GA Patch:dbed_vmclonedb -o recoverdb for offhost get failed
1873738, 1874926	dbed_vmchecksnap fails on standby database, if not all redologs from primary db are present.
1810711, 1874931	dbed_vmsnap reverse_resync_begin failed with server errors.

Veritas Storage Foundation Cluster File System fixed issues in 5.1 RP1 release

Table 1-6 Veritas Storage Foundation Cluster File System 5.1 RP1 fixed issues (listed incident number, parent number)

Fixed issues	Description
1980842, 1983222	Fixed issue in cfsadmin command for RVG volumes.
1961790, 1986445	Fixed issue in the mount(1M) command to correctly set the master node.
1878583, 1544221	getattr call optimization to speedup the case when binaries are being mmapped from many nodes on CFS.

Veritas Storage Foundation for Oracle RAC fixed issues in 5.1 RP1

Table 1-7 Veritas Storage Foundation for Oracle RAC 5.1 RP1 fixed issues

Fixed issues	Description
1938797	LMX should register with NULL canput for performance.
1934892	Issue: PrivNIC Agent support for Sun 10GbE NICs (nxge interfaces) with native 64k MTU default value Resolution: Change the mtu of the active link (where the ip address is currently plumbed). Then change the mtu of the other links (other than active link) to high value and bring the active link down.

Table 1-7 Veritas Storage Foundation for Oracle RAC 5.1 RP1 fixed issues
(continued)

Fixed issues	Description
1932014	Issue: Remove internal name built during linking libvcsmm (for Oracle 11gR2) Resolution: Remove internal name recorded by the vcsmm build script while building vcsmm library.
1908916	Fixed an issue in LMX where it may cause panic due to reuse of a freed buffer.
1891389	Issue: cssd agent support for Oracle 11gR2 Resolution: cssd agent is modified to support 11gR2 framework. 1908916 Issue: Panic lmx buffer modified after being freed Resolution: Fix the manipulation of the work queue tail pointer/ done queue tail pointer whenever the request is removed.
1853839	Issue: MultiPrivNIC resource state change to UNKNOWN once member node shutdown Resolution: The sum of the number nodes that are visible from all the devices would be zero if there is no valid LLT device. The code has been changed to handle this case.

Veritas Cluster Server fixed issues in 5.1 RP1

Table 1-8 Veritas Cluster Server 5.1 RP1 fixed issues

Fixed issues	Description
1975424	[IPMultiNICB][410-647-713][AIG] In a zone, if there are multiple IPMultiNICB resources on same subnet, issues with source address of IP pkts.
1972789	[VCS ASMinstAgent]ASMinstAgent Monitor problem inside a zone where the oracle home directory is NFS / NAS mounted inside the zone
1972770	[VCSOR]ASMinstAgent does not detect the state of the instance correctly inside local zones.
1968572	Postpatch script throws an error while installing in the non-global zone though ll and gab are hollow packages
1962548	ASMinstAgent dumpipng core.

Table 1-8 Veritas Cluster Server 5.1 RP1 fixed issues (*continued*)

Fixed issues	Description
1954723	[VCS Oracle Agent] [410-989-573] Oracle Agent Monitor problem inside a zone where the oracle home directory is NFS / NAS mounted inside the zone
1950427	[VCSOR] ASMDGAgent should disable and enable diskgroups in offline and online EPs for 11g R2.
1941647	haalert CLI hangs if engine is not in running state.
1922411	vxfentsthdw should detect storage arrays which interpret NULL keys as valid for registrations/reservations
1916022	[VCSOR][240-998-619] Changes made to Oracle agent via e1722109 do not honour ContainerName attribute
1916004	ASMAgent connecting as sysdba instead of sysasm for 11gR2
1915909	[VCS][281-889-442] hares allows to create resources which has "." special character
1911287	group stuck at OFFLINE STOPPING state when there is no ip to be cleaned in IPMultiNICB.
1874267	[ENGINE] Don't set MonitorOnly to 0 if ExternalStateChange does not have "OfflineGroup" value
1870424	LLT should give error if an attempt is made to configure more than 8 links (LLT_MAX_LINK) under LLT
1848114	SxRT5.1:Oakmont:IPMultiNICB:Resource does not come online when failover from panicked node
1504123	[SFW-HA 5.1 GCO] Symantec SE - GCO failover does not work when user account has "!" in name.

Veritas Cluster Server agents for Veritas Volume Replicator fixed issues in 5.1 RP1

No additional fixed issues exist for Veritas Cluster Server agents for Veritas Volume Replicator in the 5.1 RP1 release.

Storage Foundation Manager fixed issues in 5.1 RP1

Table 1-9 Storage Foundation Manager 5.1 RP1 fixed issues

Fixed issues	Description
1934914	Configuration fails if 2.1 CS is not configured and directly upgraded to 2.1RP1 CS
1931017	Copyright year for Windows, Solaris and HP-UX patches are 2009
1918582	Licenses not getting discovered in case default locale is non-English
1917308	when had is stopped/started vcs based monitoring should continue to function
1910997	Checkpoint size showing zero in Webgui
1904090	LDR fails to display deployment summary
1897156	Paths are not shown for one of the array ports whereas Luns information is shown
1894441	'Refresh host' needed to populate the MHs info, after upgrading package/patch through sysaddon
1893699	Unable to add a host to the management server. V-39-4095-903 401 Unauthorized User Error
1893244	Unable to add a host to the management server. V-39-4095-803 401 Unauthorized User Error
1889739	LoP hosts get list out in 'Not Installed Hosts', when deployed the sysaddon for Linux x86 MH
1888082	After deploying sysaddon patch the operation status pop up is not having host details
1887241	remove use of threads in Perl discovery
1878876	vxlist core dumping after server firmware upgrade
1878266	too many hareg processes seen on a machine where sfmh is installed
1873461	DCLI does not properly handle 2 vdirs for one OShandle
1872805	prtdiag and psrinfo -v not supported in Solaris 8, causing LDR not to display correct results
1869752	Add support for DB2 9.x support

Table 1-9 Storage Foundation Manager 5.1 RP1 fixed issues (*continued*)

Fixed issues	Description
1865225	IPv6 address not discovered in SFM gui for AIX hosts
1861664	Fix the library path for gvidid to work in case of HP 11.11
1858963	SFMH is uninstalled even if it was installed prior to install of SFW/SFWHA
1857468	VEA/vxpal continuously generate errors 0xc1000039 in vm_vxisis.log with no apparent reason
1855466	When a VVR RVG goes offline it is reported as at risk, however when it goes online again the state does not change in the UI
1855087	vxlist incorrectly shows nolabel flag for labeled disks
1854459	db2exp process is frequently core dumping on cluster node
1853081	vxship missing in VRTSsfmh for Linux
1850797	DMP Connectivity Summary view slow and causes high db CPU
1839795	Path type is empty on HP for SF 5.0 on 11.31-IA/PA
1831711	Volume Migration fails because it cannot find a target enclosure
1831697	Managing Storage Enclosure Summary reports 1 enclosure when actually 3 exist
1827451	Addhost log information is off by one month
1826556	dcli vdid can fail on HPUX LVM disks
1826409	SFM needs vxsvc service running to administer but service is not started
1825858	CS showing wrong gab port information
1809918	Servlet Exception error after adding Opteron MH to CS
1804496	postremove error messages on SFM uninstall
1797382	SFM is reporting numerous could not set locale correctly messages in error.log
1791528	VRTSsfmh error log reporting numerous errors from managed hosts
1791063	dclisetup.sh needs to be run again after upgrade to VxVM 5.1

Table 1-9 Storage Foundation Manager 5.1 RP1 fixed issues (*continued*)

Fixed issues	Description
1712298	WEBUI shows MH status as "Faulted - VEA: vxsvc or StorageAgent is not running" though all services running

VEA fixed issues in 5.1 RP1

Table 1-10 VEA 5.1 RP1 fixed issues

Fixed issues	Description
1961540	vmprov does not calculate disk nolabel state correctly.
1961519	vxsvc running as a daemon shows stderr and stdout printf's
1958763	isisd wont start, core file generated.
1958351	VEA gui fails to show controller-enclosures mapping.
1954150	Appropriate message should be display while creating Multiple Volume when size is incorrect
1954118	Not able to edit Log Settings for Alert/Task log.
1954101	While launching Gui, VEA Error message thrown "creating an instance of a class vrts.vvr.ce.REntryPoint failed"
1954047	Incorrect host version in VEA gui for 5.1RP1.
1953701	vxsvc does not start after installing RP1.
1925365	the replicated data size is showing with a negative value in VEA. (>TB)
1879928	Finish button for Break-off Snapshot for a Vset does nothing
1873583	VVR event notification sending 2 messages per event
1857207	Enabling FastResync has no effect when creating a RAID-5 volume
1846581	Core generated while downloading extension using client utility.
1840050	Core got generated while performing Volume Set operation.
1635720	Need to support volume tagging related operations of GUI in VMPROVIDER

Known issues

The following are new additional Storage Foundation and High Availability known issues in this 5.1 RP1 release.

- [Veritas Storage Foundation known issues in 5.1 RP1 release](#)
- [Veritas Volume Manager known issues in 5.1 RP1 release](#)
- [Veritas File System known issues in 5.1 RP1 release](#)
- [Veritas Storage Foundation Cluster File System known issues in 5.1 RP1 release](#)
- [Veritas Storage Foundation for Oracle RAC known issues in 5.1 RP1](#)
- [Veritas Cluster Server known issues in 5.1 RP1](#)
- [Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 RP1](#)

Veritas Storage Foundation known issues in 5.1 RP1 release

The following are new additional Storage Foundation known issues in this 5.1 RP1 release.

Live Upgrade may fail on Solaris 2.10 x86 (1984664)

This Live Upgrade issue is specific to Solaris 2.10 x86 operating system. If you run the VxVM `vxlustart` script and choose a disk for the destination boot environment (BE) that is missing the `fdisk` "SOLARIS System" partition, the Live Upgrade operation fails to properly setup the destination BE.

Workaround:

Run the `fdisk` command on the destination BE disk before performing the Live Upgrade operation. If the `fdisk` table is missing or does not contain a "SOLARIS System" partition, it must be setup.

To setup the fdisk table

- ◆ Run the `fdisk` command on the root of the destination BE:

```
# fdisk /dev/rdisk/c1t1d0p0
No fdisk table exists. The default partition for the disk is:

a 100% "SOLARIS System" partition

Type "y" to accept the default partition,
otherwise type "n" to edit the partition table.
Please answer with "y" or "n": y
```

On Solaris x86 the disks must be initialized with a fdisk table and Solaris system partition before it can be used.

dbed_clonedb of offline checkpoint fails with ORA-00600 with Oracle 11gR2 when ODM is enabled (1982674)

When performing offline checkpoint database cloning on Oracle 11gR2 and ODM is enabled, the `dbed_clonedb` command fails with error:

```
$ dbed_clonedb -S mofc1n1 -m /tmp/mofc1n1 -c \
Checkpoint_1267604996
SFORA dbed_clonedb ERROR V-81-4920 Database mofc1n1 is still in
recovery mode.
SFORA dbed_clonedb ERROR V-81-4881 Log file is at /tmp/oralog.out.10392.
```

The `/tmp/oralog.out.10392` file indicates an error.

Sample output of the `/tmp/oralog.out.10392` file:

```
ALTER DATABASE OPEN RESETLOGS
*
ERROR at line 1:
ORA-00600: internal error code, arguments: [ksfdgmsn4],
[ODM ERROR V-41-4-2-207-1 Operation not permitted],
[], [], [], [], [], [], [], [], []
ORA-00318: log 1 of thread 1, expected file size 512 doesn't match 512
ORA-00312: online log 1 thread 1:
'/tmp/mofc1n1/snap_data11r2/FLAS11r2/redo01.log'
```

Note: This issue may occur in a VVR environment.

Workaround:

Perform the offline checkpoint cloning for 11gR2 on another `ORACLE_HOME` where ODM is disabled.

Dbed_ckptrollback fails for -F datafile option for Oracle database version 11gr2 (1959400)

On Oracle 11gr2 database, `dbed_ckptrollback` fails with following error "SFORA rb.file ERROR V-81-3038 Error occurred while querying Oracle Database." The root cause of this problem is an Oracle 11GR2 defect (8367917).

Workaround:

To manually recover the datafile

- 1 Take the corrupt data file offline.
- 2 Mount the checkpoint using `dbed` utilities.
- 3 Restore the corrupt file manually.
- 4 Recover the datafile.
- 5 Bring the datafile online.

Veritas Volume Manager known issues in 5.1 RP1 release

The following are new additional Veritas Volume Manager known issues in this 5.1 RP1 release.

Changing naming scheme fails (1958711)

Changing naming scheme fails for devices controlled by MPxIO driver on Solaris. There is no workaround at this time.

vxesd dump core when it starts (1897007)

This issue happens during the case when the system is connected to a switch with more than 64 ports.

Workaround: To fix the issue, change the switch to lesser port number.

Veritas File System known issues in 5.1 RP1 release

No additional known issues exist for Veritas File System in the 5.1 RP1 release.

Veritas Storage Foundation Cluster File System known issues in 5.1 RP1 release

The following are new additional Veritas Storage Foundation Cluster File System known issues in this 5.1 RP1 release.

NFS issues with VxFS checkpoint (1974020)

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

```
Stale NFS file handle
```

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

There is no workaround at this time.

installrp recognizes SFCFSHA 5.1RP1 as SFCFS after installing SFCFSHA 5.1RP1 using JumpStart (1991079)

The `installrp` script recognizes SFCFSHA 5.1RP1 as SFCFS after installing SFCFSHA 5.1RP1 using JumpStart.

Workaround:

After you finish running JumpStart, use the `installcfs -license -ha` command to license SFCFSHA. Run `installrp` to finish the configuration.

Veritas Storage Foundation for Oracle RAC known issues in 5.1 RP1

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

Message about `mmp1_reconfig_iocti` in system log

If the Veritas Cluster Server Membership Module (VCSMM) calls the `mmp1_reconfig_iocti` function of the fencing module (VxFEN) at the time of system startup, the call fails displaying the following error message on the console and the `/var/adm/messages` file:

```
mmp1_reconfig_iocti: dev_iocti failed, vxfen may not be configured
```

You may ignore this message.

Veritas Cluster Server known issues in 5.1 RP1

The following are new additional Veritas Cluster Server known issues in this 5.1 RP1 release.

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

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

Workaround:

Set MonitorOption attribute for Oracle resource to 0.

VCS agent for Oracle: Make sure that the ohasd has an entry in the init scripts (1985093)

Make sure that the ohasd process has an entry in the init scripts so that when the process is killed or the machine is rebooted, this automatically restarts the process.

VCS agent for Oracle: Intentional Offline does not work

Intentional Offline does not work for the VCS agent for Oracle.

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

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

Workaround:

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

Veritas Cluster Server agents for Veritas Volume Replicator known issues in 5.1 RP1

No known issues exist for Veritas Storage Foundation Cluster File System in the 5.1 RP1 release.

Software limitations

The following are additional Veritas Storage Foundation and High Availability software limitations in this release.

- [Veritas Storage Foundation software limitations in 5.1 RP1 release](#)
- [Veritas Volume Manager software limitations in 5.1 RP1 release](#)
- [Veritas Storage Foundation for Oracle RAC software limitations in 5.1 RP1](#)

Veritas Storage Foundation software limitations in 5.1 RP1 release

The following are additional Veritas Storage Foundation software limitations in this release.

Thin reclamation support limitations

The thin reclamation feature has the following limitations:

- Thin reclamation only supports VxFS file systems on VxVM volumes. Other file systems are not supported.
- Thin reclamation is only supported for mounted volumes.
The file system map is not available to reclaim the unused storage space on unmounted file systems.
- Thin reclamation is not supported on raw VxVM volumes.
VxVM has no knowledge of application usage on raw volumes. Therefore, VxVM cannot perform the reclamation on raw volumes. The application must perform the reclamation on raw volumes.
- Thin reclamation is not supported on the RAID-5 layout.
The thin reclamation is storage dependent and the space underneath may or may not be reclaimed fully. Thin reclamation is not supported in a RAID-5 layout, because data consistency cannot be ensured.
- Thin Reclamation is not supported on volumes with snapshots or snapshots themselves. Any reclamation requests on such volumes or snapshots or their corresponding mount points will not result in any reclamation of their underlying storage.

Veritas Volume Manager software limitations in 5.1 RP1 release

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

Cluster Volume Manager (CVM) fail back behavior for non-Active/Active arrays (1441769)

This describes the fail back behavior for non-Active/Active arrays in a CVM cluster. This behavior applies to A/P, A/PF, APG, A/A-A, and ALUA arrays.

When all of the Primary paths fail or are disabled in a non-Active/Active array in a CVM cluster, the cluster-wide failover is triggered. All hosts in the cluster start using the Secondary path to the array. When the Primary path is enabled, the hosts fail back to the Primary path. However, suppose that one of the hosts in the cluster is shut down or brought out of the cluster while the Primary path is disabled. If the Primary path is then enabled, it does not trigger failback. The remaining hosts in the cluster continue to use the Secondary path. When the disabled host is rebooted and rejoins the cluster, all of the hosts in the cluster will continue using the Secondary path. This is expected behavior.

For A/P,APG, A/A-A, and ALUA arrays, if the disabled host is rebooted and rejoins the cluster before the Primary path is enabled, enabling the path does trigger the failback. In this case, all of the hosts in the cluster will fail back to the Primary path.

DMP settings for NetApp storage attached environment

To minimize the path restoration window and maximize high availability in the NetApp storage attached environment, set the DMP restore daemon cycle to 60 seconds. The default value of this tunable is 300 seconds. The change is persistent across reboots.

Issue the following command at the prompt:

```
# vxddmpadm settune dmp_restore_internal=60
```

To verify the new setting, use the following command:

```
# vxddmpadm gettune dmp_restore_internal
```

Veritas Storage Foundation for Oracle RAC software limitations in 5.1 RP1

The following are additional Veritas Storage Foundation for Oracle RAC software limitations in this release.

CRSResource agent

CRSResource agent is not supported for Oracle 11g Release 2.

Changes in Storage Foundation High Availability

The following sections describe changes in product behavior in this release.

About the new installrp script

Veritas Storage Foundation and High Availability Solutions 5.1 RP1 provides a new upgrade script. To upgrade from Veritas Storage Foundation and High Availability Solutions version 5.1 or later, the recommended upgrade method is to use the new upgrade script. The `installrp` script allows you to upgrade all the patches associated with the packages installed and then starts all the processes.

installrp script options

Table 1-11 shows command line options for the product upgrade script

Command Line Option	Function
[<i>system1 system2...</i>]	Specifies the systems on which to run the upgrade options. If not specified, the command prompts for a system name.
[<code>-precheck</code>]	The <code>-precheck</code> option is used to confirm that systems meet the products install requirements before installing.
[<code>-logpath log_path</code>]	The <code>-logpath</code> option is used to select a directory other than <code>/opt/VRTS/install/logs</code> as the location where <code>installrp</code> log files, summary file, and response file are saved.
[<code>-responsefile response_file</code>]	The <code>-responsefile</code> option is used to perform automated installations or uninstallations using information stored in a file rather than prompting for information. <code><response_file></code> is the full path of the file that contains configuration definitions.
[<code>-tmppath tmp_path</code>]	The <code>-tmppath</code> option is used to select a directory other than <code>/var/tmp</code> as the working directory for <code>installrp</code> . This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
[<code>-hostfile hostfile_path</code>]	The <code>-hostfile</code> option specifies the location of a file containing the system names for installer.

Table 1-11 shows command line options for the product upgrade script
(continued)

Command Line Option	Function
[<code>-jumpstart <i>jumpstart_path</i></code>]	The <code>-jumpstart</code> option is used to generate finish scripts which can be used by Solaris JumpStart Server for automated installation of all packages and patches for every product, an available location to store the finish scripts should be specified as a complete path. The <code>-jumpstart</code> option is supported on Solaris only.
[<code>-keyfile <i>ssh_key_file</i></code>]	The <code>-keyfile</code> option specifies a key file for SSH. When this option is used, <code>-i <ssh_key_file></code> is passed to every SSH invocation.
[<code>-patchpath <i>patch_path</i></code>]	The <code>-patchpath</code> option is used to define the complete path of a directory available to all install systems (usually NFS mounted) that contains all patches to be installed by <code>installrp</code> .
[<code>-rootpath <i>root_path</i></code>]	The <code>-rootpath</code> option is used to re-root the install of all packages to the given path. On Solaris, <code>-rootpath</code> passes <code>-R <root_path></code> to <code>pkgadd</code> .

Table 1-11 shows command line options for the product upgrade script
(continued)

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

CVM master node needs to assume the logowner role for VCS managed VVR resources

If you use VCS to manage VVR resources in a SFCFS or SF Oracle RAC environment, Symantec strongly recommends that you perform the steps in the section “Using the `preonline_vvr` trigger for RVGLogowner resources.” These steps ensure that the CVM master node always assumes the logowner role. Not doing this can result

in unexpected issues. These issues are due to a CVM slave node that assumes the logowner role.

See “[Using the preonline_vvr trigger for RVGLogowner resources](#)” on page 106.

Downloading the rolling patch archive

The patches included in the 5.1 RP1 release are available for download from the Symantec website. After downloading the 5.1 RP1 file, use the gunzip and tar to uncompress and extract.

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

<http://entsupport.symantec.com/docs/334829>

List of patches

This section lists the patches and packages.

Table 1-12 Patches and packages for Solaris 9 on SPARC

Patch ID	5.1 package names	Products affected	Patch size
142633-02	VRTSvxfs	FS	35912 KB
142631-02	VRTSdbed	SF, SFHA, SFCFS, SFCFSHA, SF Oracle RAC	39 MB
143696-01	VRTSdbac	SF Oracle RAC	7.1 MB
142629-02	VRTSvxvm	VM	71 MB
143270-02	VRTSodm	SF, SFCFS	1540 KB
143273-02	VRTScavf	SFCFS	695 KB
143260-02	VRTSilt	VCS	839 KB
143262-02	VRTSgab	VCS	812 KB
143706-02	VRTSvxfen	VCS	114 KB
143264-02	VRTSvcs	VCS	9497 KB
143265-02	VRTSvcsag	VCS	8 KB

Table 1-12 Patches and packages for Solaris 9 on SPARC
(continued)

Patch ID	5.1 package names	Products affected	Patch size
143279-02	VRTScps	VCS	890 KB
143276-02	VRTSvcsea	VCS	151 KB
143687-01	VRTSob	VEA	27566 KB
141270-02	VRTSsfmh	SFMH	9531 KB

Table 1-13 Patches and packages for Solaris 10 on SPARC

Patch ID	5.1 package names	Products affected	Patch size
142634-02	VRSTvxfs	FS	46436 KB
142631-02	VRTSdbed	SF, SFHA, SFCFS, SFCFSHA, SF Oracle RAC	39 MB
143697-01	VRTSdbac	SF Oracle RAC	5.8 MB
142629-02	VRTSvxvm	VM	73 KB
143271-02	VRTSodm	SF, SFCFS	1540 KB
143274-02	VRTScavf	SFCFS	697 KB
143261-02	VRTSilt	VCS	723 KB
143263-02	VRTSgab	VCS	1533 KB
143707-02	VRTSvxfen	VCS	933 KB
143264-02	VRTSvcscs	VCS	18916 KB
143265-02	VRTSvcscsag	VCS	755 KB
143279-02	VRTScps	VCS	1781 KB
143276-02	VRTSvcsea	VCS	299 KB
143687-01	VRTSob	VEA	45.6 MB
141270-02	VRTSsfmh	SFMH	26 MB

Table 1-14 Patches and packages for Solaris 10 on x64

Patch ID	5.1 package names	Products affected	Patch size
142635-02	VRTSvxfs	FS	28663 KB
142632-02	VRTSdbed	SF, SFHA, SFCFS, SFCFSHA, SF Oracle RAC	19 MB
143698-01	VRTSdbac	SF Oracle RAC	4.98 MB
142630-02	VRTSvxvm	VM	423 MB
143272-02	VRTSodm	SF, SFCFS	972 KB
143275-02	VRTScavf	SFCFS	701 KB
143266-02	VRTSllt	VCS	739 KB
143267-02	VRTSgab	VCS	581 KB
143708-02	VRTSvxfen	VCS	916 KB
143268-02	VRTSvcs	VCS	20061 KB
143269-02	VRTSvcsag	VCS	1046 KB
143280-02	VRTScps	VCS	1848 KB
143277-02	VRTSvsea	VCS	4343 KB
143693-01	VRTSob	VEA	52955 KB
141752-02	VRTSsfmh	SFMH	23552 KB

Installing the Veritas software for the first time

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

To install the Veritas software for the first time

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

```
# ./installsf node1 node2 ... nodeN
```

- For Storage Foundation HA:

```
# ./installsf -ha node1 node2 ... nodeN
```

- For Storage Foundation Cluster File System:

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

- For Storage Foundation Cluster File System HA:

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

- For Storage Foundation for Oracle RAC:

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

- For Veritas Cluster Server:

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

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

See [“Prerequisites for upgrading to 5.1 RP1”](#) on page 51.

- 3 Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program.

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

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

See [“About the new installrp script”](#) on page 33.

- If the 5.1 product is installed and not configured, run the `installrp` script to install 5.1 RP1 and configure the product.

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

See [“About the new installrp script”](#) on page 33.

The `installrp` script will give you an option to configure the product. If you choose not to configure the product at the time of the 5.1 RP1 installation, then proceed to step 4.

- 4 Mount the 5.1 product disc and navigate to the folder that contains the installation program. Run the same 5.1 installation script that you used in step 1, this time specifying the `-configure` option to configure the software.

- For Storage Foundation:

```
# ./installsf -configure node1 node2 ... nodeN
```

- For Storage Foundation HA:

```
# ./installsf -ha -configure node1 node2 ... nodeN
```

- For Storage Foundation Cluster File System:

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

- For Storage Foundation Cluster File System HA:

```
# ./installsfcfs -ha -configure node1 node2 ... nodeN
```

- For Storage Foundation for Oracle RAC:

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

- For Veritas Cluster Server:

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

See the 5.1 Installation for your product.

Installing with JumpStart

These JumpStart instructions assume a working knowledge of JumpStart. See the JumpStart documentation that came with your operating system for details on using JumpStart. Only fresh installations of Veritas product are supported using JumpStart. Upgrading is not supported. The following procedure assumes a stand-alone configuration.

For the language pack, you can use JumpStart to install packages. You add the language packages in the script, and put those files in the JumpStart server directory.

Overview of JumpStart installation tasks

The following instructions apply to the following Veritas products:

- Storage Foundation
- Storage Foundation for Oracle RAC
- Storage Foundation Cluster File System (HA)
- Veritas Cluster Server

Review the summary of tasks before you perform the JumpStart installation.

Summary of tasks

- 1 Add a client (register to the JumpStart server). See the JumpStart documentation that came with your operating system for details.
- 2 Read the JumpStart installation instructions.
- 3 Generate the finish scripts.
See [“Generating the finish scripts”](#) on page 41.
- 4 Prepare installation resources.
See [“Preparing installation resources”](#) on page 46.
- 5 Run JumpStart to install the Veritas product.

Note: JumpStart may reboot systems after product installation.

- 6 Run the `installer` command from the disc or from directory `/opt/VRTS/install` directory to configure the Veritas software.

```
# /opt/VRTS/install/installer installprod -configure
```

Where *installprod* is the product's installation command.

- 7 Modify the rules file for JumpStart.
See the JumpStart documentation that came with your operating system for details.

Generating the finish scripts

Perform these steps to generate the finish script to install Veritas product.

To generate the script

- 1 Run the `installrp` program to generate the scripts.

```
# installrp -jumpstart directory_to_generate_scripts
```

Where the *directory_to_generate_scripts* is where you want to put the scripts.

For example:

```
# ./installrp -jumpstart /js_scripts
```

- 2 When you are prompted to encapsulate the root disk automatically, choose **yes** to do so. If you do not want to encapsulate it automatically, choose **no** and go to step 6.
- 3 Specify a disk group name for the root disk.

```
Specify the disk group name of the root disk to be encapsulated:  
rootdg
```

- 4 Specify private region length.

```
Specify the private region length of the root disk to be  
encapsulated: (65536)
```

- 5 Specify the disk's media name of the root disk to encapsulate.

```
Specify the disk media name of the root disk to be encapsulated:  
(rootdg_01)
```

6 JumpStart finish scripts, installer scripts, uninstaller scripts of Veritas products, and encapsulation scripts are generated in the directory you specified in step 1. Output resembles:

```
The finish scripts for AT50 is generated at
/js4/jumpstart_at50.fin
The installer script to configure AT is generated at
/js4/installat
The installer script to uninstall AT is generated at
/js4/uninstallat
The finish scripts for FS51 is generated at
/js4/jumpstart_fs51.fin
The installer script to configure FS is generated at
/js4/installfs
The installer script to uninstall FS is generated at
/js4/uninstallfs
The finish scripts for SF51 is generated at
/js4/jumpstart_sf51.fin
The installer script to configure SF is generated at
/js4/installsf
The installer script to uninstall SF is generated at
/js4/uninstallsf
The finish scripts for SFCFS51 is generated at
/js4/jumpstart_sfcfs51.fin
The installer script to configure SFCFS is generated at
/js4/installsfdfs
The installer script to uninstall SFCFS is generated at
/js4/uninstallsfcfs
The finish scripts for SFCFSHA51 is generated at
/js4/jumpstart_sfcfsha51.fin
The installer script to configure SFCFSHA is generated at
/js4/installsfcfsha
The installer script to uninstall SFCFSHA is generated at
/js4/uninstallsfcfsha
The finish scripts for SFHA51 is generated at
/js4/jumpstart_sfha51.fin
The installer script to configure SFHA is generated at
/js4/installsfha
The installer script to uninstall SFHA is generated at
/js4/uninstallsfha
The finish scripts for SFRAC51 is generated at
/js4/jumpstart_sfrac51.fin
The installer script to configure SF Oracle RAC is generated at
```

```
/js4/installsfrac
```

The installer script to uninstall SF Oracle RAC is generated at

```
/js4/uninstallsfrac
```

The finish scripts for VCS51 is generated at

```
/js4/jumpstart_vcs51.fin
```

The installer script to configure VCS is generated at

```
/js4/installvcs
```

The installer script to uninstall VCS is generated at

```
/js4/uninstallvcs
```

The finish scripts for VM51 is generated at

```
/js4/jumpstart_vm51.fin
```

The installer script to configure VM is generated at

```
/js4/installvm
```

The installer script to uninstall VM is generated at

```
/js4/uninstallvm
```

The encapsulation boot disk script for VM is generated at

```
/js4/encap_bootdisk_vm51001000.fin
```

List the `js_scripts` directory.

```
# ls /js_scripts
```

You could select scripts according to the products. For example.

For SF:

```
encap_bootdisk_vm51001000.fin installsf jumpstart_sf51.fin uninstallsf
```

For SF Oracle RAC:

```
encap_bootdisk_vm51001000.fin installsfrac jumpstart_sfrac51.fin  
uninstallsfrac
```

For SFHA:

```
encap_bootdisk_vm51001000.fin installsfha jumpstart_sfha51.fin  
uninstallsfha
```

For VCS:

```
encap_bootdisk_vm51001000.fin jumpstart_vcs51.fin installvcs  
uninstallvcs
```

- 7** Modify the JumpStart script according to your requirements. You must modify the `BUILDSRC` and `ENCAPSRC` values. Keep the values aligned with the resource location values.

```
BUILDSRC="hostname_or_ip:/path_to_pkgs_patches_scripts"
// If you don't want to encapsulate the root disk automatically
// comment out the following line.
ENCAPSRC="hostname_or_ip:/path_to_encap_script"
```

- 8** If you want to install different products, use the following command to get the sequence for the product. In the following commands, replace the variable *prod* with the product's acronym. See the product documentation for more information.

- For the minimum set of packages, use:

```
# installprod -minpkgs
```

- For the recommended set of packages, use:

```
# installprod -recpkgs
```

An example of this command is:

For SF:

```
# ./installsf -minpkgs
SF: PKGS: VRTSvlic VRTSperl VRTSvxvm VRTSaslapm VRTSvxfs
```

For SF Oracle RAC:

```
# ./installsfrac -minpkgs
SF Oracle RAC: PKGS: VRTSvlic VRTSperl VRTSvxvm VRTSaslapm VRTSvxfs
VRTSllt VRTSgab VRTSvxfen VRTSvcs VRTSvcsag VRTSat VRTSvcssea VRTSdbed
VRTSglm VRTScavf VRTSgms VRTSodm VRTSdbac
```

For SFCFS:

```
# ./installsfcfs -minpkgs
SFCFS: PKGS: VRTSvlic VRTSperl VRTSvxvm VRTSaslapm VRTSvxfs VRTSllt
VRTSgab VRTSvxfen VRTSvcs VRTSvcsag VRTSat VRTSglm VRTScavf
```

For VCS:

```
# ./installvcs -minpkgs
VCS: PKGS: VRTSvlic VRTSperl VRTSllt VRTSgab VRTSvxfen VRTSvcs
VRTSvcsag VRTSat
```

Use the list of packages that is generated to replace the package list in the finish scripts.

Preparing installation resources

Prepare resources for the JumpStart installation.

To prepare the resources

- 1 Copy the contents of the installation disc to the shared storage.

```
# cd /cdrom/cdrom0  
# cp -r * BUILDSRC
```

Note: After you copied the patches, you must unzip and untar them.

- 2 Generate the response file for the package list that you found in [Generating the finish scripts](#) step 8. In this example the packages are:

For SF:

```
VRTSvlic VRTSperl VRTSspt VRTSvxvm VRTSaslapm VRTSob VRTSsfmh  
VRTSvxfs VRTSfssdk VRTSdbed VRTSodm VRTSat
```

For SF Oracle RAC:

```
VRTSvlic VRTSperl VRTSvxvm VRTSaslapm VRTSvxfs VRTSllt VRTSgab  
VRTSvxfen VRTSvcs VRTSvcsag VRTSat VRTSvcssea VRTSdbed VRTSglm  
VRTScavf VRTSgms VRTSodm VRTSdbac
```

For SFCFS:

```
VRTSvlic VRTSperl VRTSspt VRTSvxvm VRTSaslapm VRTSob VRTSsfmh  
VRTSvxfs VRTSfssdk VRTSat VRTSllt VRTSgab VRTSvxfen VRTSvcs VRTScps  
VRTSvcsag VRTScutil VRTSvcssea VRTSdbed VRTSglm VRTScavf VRTSgms  
VRTSodm
```

For VCS:

```
VRTSvlic VRTSperl VRTSspt VRTSat VRTSllt VRTSgab VRTSvxfen VRTSvcs  
VRTScps VRTSvcsag VRTScutil VRTSvcssea
```

```
# cd BUILDSRC/pkgs/  
# pkgask -r package_name.response -d BUILDSRC/pkgs/packages_name.pkg
```

- 3 Create the adminfile file under *BUILDSRC/pkgs/* directory. The adminfile file's contents follow:

```
mail=  
instance=overwrite  
partial=nocheck  
runlevel=quit  
idepend=quit  
rdepend=nocheck  
space=quit  
setuid=nocheck  
conflict=nocheck  
action=nocheck  
basedir=default
```

- 4 If you want to encapsulate the root disk automatically when perform the JumpStart installation, copy the scripts `encap_bootdisk_vm51001000.fin` generated in [Generating the finish scripts](#) step 6 to *ENCAPSRC*
- 5 Modify the rules file as required.

For example:

```
any - - profile_sf jumpstart_sf51.fin
```

For detailed instructions, see the *Sun Microsystems' JumpStart* documentation.

Adding language pack information to the finish file

For the language pack, copy the language packages from the language pack installation disc to the shared storage.

```
# cd /cdrom/cdrom0/pkgs  
# cp -r * BUILDSRC/pkgs
```

Add lines for the language packages in the finish script. If the finish file resembles:

```
. . .  
for PKG in VRTSperl VRTSvlic VRTSicsco . . .  
  
do  
. . .  
done
```

Add the following lines for the language pack after the patch information for VCS. Copy the command syntax between the "do" and "done" lines and add that for the language pack lines as well. Note that the line that starts "for PKG" is on three lines in this guide, but should be on a single line in the file.

```
. . .
for PKG in VRTSmulic VRTSatJA VRTSjacav VRTSjacs VRTSjacse
VRTSjacsu VRTSjadba VRTSjafs VRTSjavm VRTSjadbe VRTSjaodm
VRTSatZH VRTSzhvm

do
.
.
.
done
```

Installing 5.1 RP1 using the web-based installer

This section describes how to install 5.1 RP1 using the web-based installer.

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

About the Web-based installer

The `webinstaller` script is used to start and stop the Veritas XPortal Server `xprtlwid` process. The `webinstaller` script can also be used to check the status of the XPortal Server.

When the `webinstaller` script starts the `xprtlwid` process, the script displays a URL. Use this URL to access the Web-based installer from a Web browser such as Internet Explorer or FireFox.

The Web installer creates log files whenever the Web installer is operating. While the installation processes are operating, the log files are located in a session-based directory under the `/var/tmp` directory. After the install process completes, the log files are located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

The location of the Veritas XPortal Server configuration file is `/var/opt/webinstaller/xprtlwid.conf`.

Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

Table 1-15 Web-based installer requirements

System	Function	Requirements
Target system	The system(s) where the Veritas products will be installed.	Must be a supported platform for Veritas product 5.1 RP1
Installation server	The server from which to initiate the installation. The installation media is mounted and accessible from the installation server.	Must be the same OS as the system(s) on which to install.
Administrative system	The system on which you run the web browser to perform the installation.	Web browser

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

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

```
# ./webinstaller start
```

The `webinstaller` script displays a URL.

- 2 On the administrative server, start the Web browser.

- 3 Navigate to the URL displayed from step 1.

- 4 The browser may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

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

Obtaining a security exception on Mozilla Firefox

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

To obtain a security exception

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

Performing a pre-installation check with the Veritas Web-based installer

This section describes performing a pre-installation check with the Veritas Web-based installer.

To perform a pre-installation check

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 49.
- 2 On the Select a task and a product page, select **Perform a Pre-installation check** from the **Task** drop-down list.
- 3 Indicate the systems on which to perform the precheck. Enter one or more system names, separated by spaces. Click **Validate**.
- 4 The installer performs the precheck and displays the results.
- 5 If the validation completes successfully, click **Next**. The installer prompts you to begin the installation. Click **Yes** to install Veritas product on the selected system. Click **No** to install later.
- 6 Click **Finish**. The installer prompts you for another task.

Installing products with the Veritas Web-based installer

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

To install Veritas product

- 1 Perform preliminary steps.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 49.

- 3 On the Select a task and product page, select **Install RP1** from the **Task** drop-down list.
- 4 Select Veritas product or Veritas product High Availability from the Product drop-down list, and click Next.
- 5 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 6 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.
- 7 After the validation completes successfully, click **Next** to install Veritas product on the selected system.
- 8 For Storage Foundation, click Next to complete the configuration and start the product processes.

For Storage Foundation High Availability, the installer prompts you to configure the cluster.

Note that you are prompted to configure only if the product is not yet configured.

If you select n, you can exit the installer. You must configure the product before you can use Veritas product.

See your Veritas product 5.1 Installation Guide to configure your product.

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

- 9 Select the checkbox to specify whether you want to send your installation information to Symantec.

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

Click **Finish**. The installer prompts you for another task.

Prerequisites for upgrading to 5.1 RP1

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

- For any product in the Storage Foundation stack, regardless of your operating system, you must have the 5.1 release installed before you can upgrade that product to the 5.1 RP1 release.
- Each system must have sufficient free space to accommodate patches.

Supported upgrade paths

This section describes the supported upgrade paths for this release.

- 5.1 to 5.1 RP1
- 5.1 P1 to 5.1 RP1
- 5.1 to 5.1 P1 to 5.1 RP1

Upgrading 5.1 to 5.1 RP1

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

- [Performing a full upgrade to 5.1 RP1 on a cluster](#)
Use the procedures to perform a full upgrade to 5.1 RP1 on a cluster that has VCS, SFHA, SFCFS, or SF Oracle RAC installed and configured.
- [Performing a phased upgrade to 5.1 RP1 on a cluster](#)
Use the procedures to perform a phased upgrade to 5.1 RP1 on a cluster that has VCS, SFHA, SFCFS, or SF Oracle RAC installed and configured.
- [Upgrading Veritas product with the Veritas Web-based installer](#)
Use the procedure to upgrade your Veritas product with the Web-based installer.
- [Performing a rolling upgrade using the installer](#)
Use the procedure to upgrade your Veritas product with a rolling upgrade.
- [Performing a rolling upgrade manually](#)
Use the procedure to upgrade your Veritas product manually with the rolling upgrade.
- [Upgrading to 5.1 RP1 on a standalone system](#)
Use the procedure to upgrade to 5.1 RP1 on a system that has SF and VCS installed.

Performing a full upgrade to 5.1 RP1 on a cluster

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

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

- Freeze service group operations and stop VCS on the cluster.
- Take the nodes offline and install the software patches.

- Bring the nodes online to restart cluster failover services.

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

- [Performing a full upgrade to 5.1 RP1 for VCS](#)
- [Performing a full upgrade to 5.1 RP1 on a SFHA cluster](#)
- [Performing a full upgrade to 5.1 RP1 on a SFCFS cluster](#)
- [Performing a full upgrade to 5.1 RP1 on a SF Oracle RAC cluster](#)

Performing a full upgrade to 5.1 RP1 for VCS

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

To upgrade VCS

- 1 Review the installation prerequisites.
See [“Prerequisites for upgrading to 5.1 RP1”](#) on page 51.
- 2 Check the readiness of the nodes where you plan to upgrade. Start the pre-upgrade check:

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

See [“About the new installrp script”](#) on page 33.

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

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

- 5 After the upgrade, review the log files.
- 6 Verify the upgrade.
See [“Verifying software versions”](#) on page 92.

Performing a full upgrade to 5.1 RP1 on a SFHA cluster

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

To perform a full upgrade to 5.1 RP1 on a SFHA cluster

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

- 3 Make the VCS configuration writable on a node that is being upgraded:

```
# haconf -makerw
```

- 4 Freeze the HA service group operations. Enter the following command on each node, if you selected a group of nodes on which to upgrade the operating system:

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

- 5 Make the VCS configuration read-only:

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

- 6 Close any instance of VCS GUI that is running on the node.

- 7 Stop VCS:

```
# hastop -local
```

- 8 Stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of CmdServer.

- 9 Stop cluster fencing, GAB, and LLT.

```
# /etc/init.d/vxfen stop  
# /etc/init.d/gab stop  
# /etc/init.d/llt stop
```

- 10 If required, apply the OS kernel patches.

See [“System Requirements”](#) on page 12.

See *Sun Microsystems’* documentation for the procedures.

- 11 Repeat step 7 through step 9 if the system reboots after upgrading the operating system. You need to perform this to stop the components that started by the init scripts, if any.

- 12 Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program. Enter the `installrp` script. If ssh key authentication is configured then enter:

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

If ssh is not configured then enter:

```
# ./installrp -rsh node1 node2
```

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

- 13 After all of the nodes in the cluster are upgraded, shut down and reboot each of the nodes. After the nodes come up, application failover capability is available.

- 14 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

- 15 Unfreeze the service group operations on each node:

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

- 16 Make the VCS configuration read-only:

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

Performing a full upgrade to 5.1 RP1 on a SFCFS cluster

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

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

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

```
# haconf -makerw
```

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

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

- 5 Make the configuration read-only:

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

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

```
# df -F vxfs
```

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

```
# umount /checkpoint_name
```

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

```
# df -F vxfs
```

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

```
# umount /filesystem
```

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

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

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

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

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

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

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

- 9 Stop activity to all VxVM volumes.

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

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

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

Verify that no volumes remain open:

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

- 11 Stop VCS:

```
# hastop -all
```

- 12 On each node, stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of CmdServer.

- 13 On each node, stop ODM, cluster fencing, GAB, and LLT in the following order:

- Solaris 9:

```
# /etc/init.d/odm stop  
# /etc/init.d/vxfen stop  
# /etc/init.d/gab stop  
# /etc/init.d/llt stop
```

- Solaris 10:

```
# svcadm disable -t vxfen  
# svcadm disable -t vxodm  
# svcadm disable -t gab  
# svcadm disable -t llt
```

- 14 If required, apply the OS kernel patches.

See “[System Requirements](#)” on page 12.

See *Sun Microsystems’* documentation for the procedures.

- 15 On each node, check if the VEA service is running:

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

If the VEA service is running, stop it:

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

- 16** Navigate to the folder that contains the `installrp` program and start the `installrp` program:

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

Review the output.

- 17** Start services for LLT, GAB, cluster fencing and ODM on all upgraded nodes:

For Solaris 9:

```
# /etc/init.d/llt start
# /etc/init.d/gab start
# /etc/init.d/vxfen start
# /etc/init.d/odm start
```

For Solaris 10:

```
# svcadm enable llc
# svcadm enable gab
# svcadm enable vxfen
# svcadm enable vxodm
```

- 18** Start vcs on all upgraded nodes:

```
# /opt/VRTSvcs/bin/hastart
```

- 19** If necessary, reinstate any missing mount points in the `/etc/vfstab` file on each node.

- 20** Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

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

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

- 22** Make the configuration read-only:

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

- 23** Bring the CVM service group online on each node:

```
# hagrps -online cvm -sys nodename
```

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

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

- 25 If you stopped any RVGs in step 10, restart each RVG:

```
# vxrvrg -g diskgroup start rvg_name
```

- 26 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

- 27 Remount all Storage Checkpoints on all nodes:

```
# mount /checkpoint_name
```

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

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

To upgrade to 5.1 RP1 on a SF Oracle RAC cluster

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

```
# haconf -makerw
```
- 4 Enter the following command to freeze HA service group operations on each node:

```
# hasys -freeze -persistent nodename
```
- 5 Make the configuration read-only:

```
# haconf -dump -makero
```
- 6 If CRS is not controlled by VCS, enter the following command on each node of the cluster to stop CRS:

```
# /etc/init.d/init.crs stop
```

7 Stop VCS.

```
# hastop -all
```

8 Stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of CmdServer.

9 Stop VCSMM and LMX if they are running:

For Solaris 9:

```
# /etc/init.d/vcsmm stop  
# /etc/init.d/lmx stop
```

For Solaris 10:

```
# svcadm disable -t vcsmm  
# svcadm disable -t lmx
```

10 Stop cluster fencing, ODM, and GAB:

For Solaris 9:

```
# /etc/init.d/vxfen stop  
# /etc/init.d/odm stop  
# /etc/init.d/gab stop
```

For Solaris 10:

```
# svcadm disable -t vxfen  
# svcadm disable -t vxodm  
# svcadm disable -t gab
```

11 On each node, unload the vxfen, LMX, GAB, VCSMM, GMS, and GLM kernel modules if they are still loaded:

- Verify if the `vxfen` kernel module is loaded. For example:

```
# modinfo|grep vxfen  
210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.1)
```

If the `vxfen` kernel module is loaded then unload it. For example:

```
# modunload -i 210
```

- Verify if the LMX kernel module is loaded. For example:

```
# modinfo | grep lmx
239 ffffffff1253000 13a30 236 1 lmx (LLT Mux '5.1')
```

If the LMX kernel module is loaded then unload it. For example:

```
# modunload -i 239
```

- Verify if the VCSMM kernel module is loaded. For example:

```
# modinfo | grep vcsmm
312 78bc0000 43ae8 293 1 vcsmm (VRTSvcsmm 5.1)
```

If the VCSMM kernel module is loaded then unload it. For example:

```
# modunload -i 312
```

- Verify if the GMS kernel module is loaded. For example:

```
# modinfo | grep gms
311 78289c91 4867 292 1 vxgms (VxGMS 5.1 (SunOS))
```

If the GMS kernel module is loaded then unload it. For example:

```
# modunload -i 311
```

- Verify if the GLM kernel module is loaded. For example:

```
# modinfo | grep glm
310 78b68000 24268 291 1 vxglm (VxGLM 5.1 (SunOS 5.10))
```

If the GLM kernel module is loaded then unload it. For example:

```
# modunload -i 310
```

- Verify if the GAB kernel module is loaded. For example:

```
# modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.1)
```

If the GAB kernel module is loaded then unload it. For example:

```
# modunload -i 149
```

12 Stop LLT:

For Solaris 9:

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

For Solaris 10:

```
# svcadm disable -t llt
```

- Verify if the LLT kernel module is loaded. For example:

```
# modinfo | grep llt
147 50ca4000 d6bc 110 1 llt (LLT 5.1)
```

If the LLT kernel module is loaded then unload it. For example:

```
# modunload -i 147
```

- 13 If required, apply the OS kernel patches.

See “[System Requirements](#)” on page 12.

See *Sun Microsystems*’ documentation for the procedures.

Note: If you are upgrading a SF Oracle RAC cluster, you must upgrade the nodes of the cluster at this stage to one of the operating system versions that this release supports.

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

```
# df -F vxfs
```

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

```
# umount /filesystem
```

- On each node of the cluster, verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | fsdb -F vxfs filesystem | grep clean
flags 0 mod 0 clean clean_value
```

A *clean_value* value of 0x5a indicates the file system is clean, 0x3c indicates the file system is dirty, and 0x69 indicates the file system is dusty. A dusty file system has pending extended operations.

- If a file system is not clean, enter the following commands for that file system:

```
# fsck -F vxfs filesystem
# mount -F vxfs filesystem mountpoint
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly. There may be a pending large fileset clone removal extended operation if the unmount command fails with the following error:

```
file system device busy
```

You know for certain that an extended operation is pending if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system
file system still in progress.
```

- If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
- Repeat the following command to verify that the unclean file system is now clean:

```
# echo "8192B.p S" | fsdb -F vxfs filesystem | grep clean
flags 0 mod 0 clean clean_value
```

- 15 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
- 16 On each node of the cluster, stop all VxVM volumes by entering the following command for each disk group:

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

Verify that no volumes remain open:

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

- 17 Check if the VEA service is running:

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

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

- 18** Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program. Enter the `installrp` script. If ssh key authentication is configured then enter:

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

If ssh is not configured then enter:

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

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

- 19** After the entire cluster is upgraded, reboot all of the nodes of the cluster.

```
# /usr/sbin/shutdown -g0 -y -i6
```

- 20** If necessary, reinstate any missing mount points in the `/etc/vfstab` file on each node.

- 21** Run the following commands to start the Storage Foundation for Oracle RAC processes:

For Solaris 9:

```
# /etc/init.d/llt start  
# /etc/init.d/gab start  
# /etc/init.d/odm start  
# /etc/init.d/vxfen start  
# /etc/init.d/vcsmm start  
# /etc/init.d/lmx start  
# /opt/VRTSvcs/bin/hastart
```

For Solaris 10:

```
# svadm enable llc  
# svadm enable gab  
# svadm enable vxodm  
# svadm enable vxfen  
# svadm enable vcsmm  
# svadm enable lmx  
# /opt/VRTSvcs/bin/hastart
```

- 22** From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```


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

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

24 Make the configuration read-only:

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

25 Enter the following command on each node to take service groups online:

```
# hagrps -online service_group -sys nodename
```

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

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

27 If CRS is not controlled by VCS, enter the following command on each node to start CRS.

```
# /etc/init.d/init.crs start
```

28 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

29 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

Performing a phased upgrade to 5.1 RP1 on a cluster

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

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

- Freeze service group operations and stop VCS on the cluster.

- Select a group of one or more cluster nodes to upgrade (group A), and leave a group of one or more nodes running (group B).
- Take offline the nodes in group A and install the software patches on those nodes.
- Take offline the nodes in group B and bring online the nodes in group A to restart cluster failover services.
- Upgrade the nodes in group B, then bring those nodes online to join. The cluster is fully restored.

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

- [Performing a phased upgrade to 5.1 RP1 for VCS](#)
- [Performing a phased upgrade to 5.1 RP1 on an SFHA cluster](#)
- [Performing a phased upgrade to 5.1 RP1 on an SFCFS cluster](#)
- [Performing a phased upgrade to 5.1 RP1 on an SF Oracle RAC cluster](#)

Performing a phased upgrade to 5.1 RP1 for VCS

The following procedure describes performing a phased upgrade for VCS.

To perform a phased upgrade to 5.1 RP1 for VCS

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.
- 3 Switch the service group to another node that is running.

```
# hagrps -switch service_group -to nodename
```

- 4 Make the VCS configuration writable on a node that is being upgraded:

```
# haconf -makerw
```

- 5 Freeze the service group operations. Enter the following command on each node, if you selected a group of nodes on which to upgrade the operating system:

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

- 6 Make the VCS configuration read-only:

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

- 7 Close any instance of VCS GUI that is running on the node.
- 8 Select the group of nodes that are to be upgraded first, and follow step 9 through step 19 for these nodes.

- 9 Stop VCS on each node in the selected group:

```
# hastop -local
```

- 10 Stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of CmdServer.

- 11 Stop cluster fencing, GAB, and LLT.

For Solaris 9:

```
# /etc/init.d/vxfen stop  
# /etc/init.d/gab stop  
# /etc/init.d/llt stop
```

For Solaris 10:

```
# svcadm disable -t vxfen  
# svcadm disable -t gab  
# svcadm disable -t llt
```

- 12 If required, you can upgrade the operating system on the nodes in the selected group at this stage and patch them to a supported kernel version.

See “[System Requirements](#)” on page 12.

- 13 Repeat step 9 through step 11 if the system reboots after upgrading the operating system. You need to perform this to stop the components that started by the init scripts, if any.
- 14 Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program. On the first sub-cluster, use `installrp` command to upgrade VCS to 5.1 RP1:

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

- 15** Stop VCS, I/O fencing, GAB, and LLT on the nodes that you plan to upgrade next.

For Solaris 9:

```
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

For Solaris 10:

```
# svcadm disable -t vxfen
# svcadm disable -t gab
# svcadm disable -t llt
```

- 16** On the first sub-cluster, restart the nodes:

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

- 17** On the nodes that you have rebooted, seed the nodes.

```
# gabconfig -xc
```

- 18** Make the VCS configuration writable again from any node in the selected group:

```
# haconf -makerw
```

- 19** Unfreeze the service group operations on each node where you upgraded the operating system:

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

- 20** Make the VCS configuration read-only:

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

- 21** Bring the service group online on the original node:

```
# hagrps -online service_group -sys nodename
```

- 22 Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program. On the second sub-cluster, use `installrp` command to upgrade VCS to 5.1 RP1:

```
# ./installrp node3 node4 ...
```

- 23 Restart the nodes in the second sub-cluster.

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

Performing a phased upgrade to 5.1 RP1 on an SFHA cluster

The following procedure describes performing a phased upgrade on an SFHA cluster.

To perform a phased upgrade to 5.1 RP1 on an SFHA cluster

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.
- 3 Switch the service group to another node that is running.

```
# hagrps -switch service_group -to nodename
```

- 4 Make the VCS configuration writable on a node that is being upgraded:

```
# haconf -makerw
```

- 5 Freeze the HA service group operations. Enter the following command on each node, if you selected a group of nodes on which to upgrade the operating system:

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

- 6 Make the VCS configuration read-only:

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

- 7 Close any instance of VCS GUI that is running on the node.
- 8 Select the group of nodes that are to be upgraded first, and follow step 9 through step 19 for these nodes.
- 9 Stop VCS on each node in the selected group:

```
# hstop -local
```

- 10** Stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of *CmdServer*.

- 11** Stop cluster fencing, GAB, and LLT.

For Solaris 9:

```
# /etc/init.d/vxfen stop  
# /etc/init.d/gab stop  
# /etc/init.d/llt stop
```

For Solaris 10:

```
# svcadm disable -t vxfen  
# svcadm disable -t gab  
# svcadm disable -t llt
```

- 12** If required, apply the OS kernel patches on the nodes in the selected group.

See “[System Requirements](#)” on page 12.

See Sun Microsystems’ documentation for the procedures.

- 13** Repeat step 9 through step 11 if the system reboots after upgrading the operating system. You need to perform this to stop the components that started by the init scripts, if any.

- 14** Before you begin the upgrade, you can check the readiness of the nodes where you plan to upgrade. The command to start the pre-upgrade check is:

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

where *node1* is *galaxy* and *node2* is *nebula*, for example:

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

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

- 15** After all of the nodes in the selected group are upgraded, shut down and reboot each of the nodes. After the nodes come up, application failover capability is available for that group of nodes.

- 16 Make the VCS configuration writable again from any node in the selected group:

```
# haconf -makerw
```

- 17 Unfreeze the service group operations on each node for which you upgraded the operating system:

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

- 18 Make the VCS configuration read-only:

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

- 19 Switch the service group to the original node:

```
# hagrps -switch service_group -to nodename
```

- 20 Repeat step 9 through step 19 for the second group of nodes.

Performing a phased upgrade to 5.1 RP1 on an SFCFS cluster

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

To perform a phased upgrade to 5.1 RP1 on an SFCFS cluster

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` and `/opt/VRTSvcs/bin` is in your `PATH` so that you can execute all product commands.
- 3 If you have a failover service group, switch the service group to another node that is running.

```
# hagrps -switch service_group -to nodename
```

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

```
# haconf -makerw
```

- 5 Enter the following command to freeze HA service group operations on each node:

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

- 6 Make the configuration read-only:

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

- 7 Select a group of nodes that are to be upgraded first, and follow step 8 through step 17 and step 23 to step 36 for these nodes.

- 8 On each node in the selected group, enter the following command to check if any Storage Checkpoints are mounted:

```
# df -F vxfs
```

If any Storage Checkpoints are mounted, on each node in the selected group unmount all Storage Checkpoints

```
# umount /checkpoint_name
```

- 9 On each node in the selected group, enter the following command to check if any VxFS file systems are mounted:

```
# df | grep vxfs
```

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

```
# umount /filesystem
```

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

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

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

- 11 Stop VCS on each node in the selected group:

```
# hastop -local
```


12 Stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of CmdServer.

13 Stop ODM, cluster fencing, GAB, and LLT in the following order:

■ Solaris 9:

```
# /etc/init.d/odm stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

■ Solaris 10:

```
# umount /dev/odm
# svcadm disable vxodm
# svcadm disable vxfen
# svcadm disable gab
# svcadm disable llt
```

14 Check if the VEA service is running:

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

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

15 If required, apply the OS kernel patches on the nodes in the selected group.

See Sun Microsystems' documentation for the procedures.

16 Repeat step 11 through step 14 if the system reboots after upgrading the operating system. You need to perform this to stop the components that started by the init scripts, if any.

- 17** Before you begin the upgrade, you can check the readiness of the nodes where you plan to upgrade. The command to start the pre-upgrade check is:

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

where *node1* is galaxy and *node2* is nebula, for example:

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

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

If the precheck determines the nodes ready for upgrade, continue with the upgrade. If the precheck discovers issues with the nodes, resolve the issues, and return to the precheck.

- 18** On the second set of nodes (the ones that have not yet been upgraded), enter the following command to check if any Storage Checkpoints are mounted:

```
# df -F vxfs
```

If any Storage Checkpoints are mounted, on each node in this group unmount all Storage Checkpoints

```
# umount /checkpoint_name
```

- 19** On the second set of nodes (the ones that have not yet been upgraded), enter the following command to check if any VxFS file systems are mounted:

```
# df | grep vxfs
```

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

```
# umount /filesystem
```

- 20** Stop VCS on the second set of nodes (the nodes that you have not upgraded):

```
# hastop -local
```

- 21** Stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of CmdServer.

- 22** Stop ODM, cluster fencing, GAB, and LLT on the second set of nodes:

■ Solaris 9:

```
# /etc/init.d/odm stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

■ Solaris 10:

```
# umount /dev/odm
# svcadm disable vxodm
# svcadm disable vxfen
# svcadm disable gab
# svcadm disable llt
```

23 On the first set of nodes (the upgraded nodes), run the following commands to start the Storage Foundation High Availability processes:

■ Solaris 9:

```
# /etc/init.d/odm start
# /etc/init.d/vxfen start
# /etc/init.d/gab start
# /etc/init.d/llt start
```

■ Solaris 10:

```
# svcadm enable llt
# svcadm enable gab
# svcadm enable vxfen
# svcadm enable vxodm
```

24 Start VCS on an upgraded node:

```
# /opt/VRTSvcs/bin/hastart
```

25 On the nodes that you have upgraded, seed the nodes.

```
# gabconfig -xc
```

26 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.

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

28 Make the VCS configuration writable again from any node in the selected group:

```
# haconf -makerw
```

29 Enter the following command on each node in the selected group to unfreeze HA service group operations:

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

30 Make the configuration read-only:

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

31 Autoenable and bring online the failover service group to an upgraded node:

```
# hagrps -autoenable service_group -sys nodename  
# hagrps -online service_group -sys nodename
```

32 Bring the CVM service group online on each node in the selected group:

```
# hagrps -online cvm -sys nodename
```

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

```
# mount /filesystem
```

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

```
# mount /checkpoint_name
```

35 Repeat step 15 through step 17 for the second group of nodes.

36 Repeat step 23 through step 34 for the second group of nodes.

Performing a phased upgrade to 5.1 RP1 on an SF Oracle RAC cluster

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

The phased upgrade methodology involves upgrading 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 starting the upgrade on the first half of the cluster, back up the configuration files.

To upgrade to 5.1 RP1 on an SF Oracle RAC cluster

- 1 Log in as superuser.
- 2 Switch failover groups from the first half of the cluster from *galaxy* to the second half of the cluster to *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:

```
# 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
```

If CRS is not controlled by VCS, enter the following command on each node to stop CRS.

```
# /etc/init.d/init.crs stop
```

- 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.

Ensure that no processes are running that make use of mounted shared file system or shared volumes. To verify that no processes use the VxFS or CFS mount point, enter the following commands:

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

- 6 On first half of the cluster stop all VxVM and CVM volumes for each diskgroup that are 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 5.1 RP1 product disc and navigate to the folder that contains the installation program. On the first half of the cluster, enter the `installrp` script:

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

Note: After you complete the upgrade of the first half of the cluster, no GAB ports will be showing 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 Oracle.

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

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

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

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

- 12** On the second half of the cluster, stop the Oracle database:

If the Oracle RAC instance is managed by VCS:

```
# hagrps -offline oracle_group -sys jupiter
# hagrps -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
```

If CRS is not controlled by VCS, enter the following command on each node to stop CRS.

```
# /etc/init.d/init.crs stop
```

Note: The downtime starts now.

- 13** On the second half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 14** On the second half of the cluster, unmount the VxFS or CFS file systems that are not managed by VCS.

Ensure that no processes are running that make use of mounted shared file system or shared volumes. To verify that no processes use the VxFS or CFS mount point:

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

- 15** 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
```

- 16** On the second half of the cluster, stop VCS:

```
# hastop -local
```

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

```
# hastart
```

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

```
# hagrp -online group_name -any
```

Once the cvm service group comes online, the GAB ports v, w, and f come online; 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.

- 19 On the first half of the cluster, manually mount the VxFS or CFS file systems that are not managed by VCS.
- 20 On the first half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

If CRS is not controlled by VCS, enter the following command on each node of first of cluster to start CRS.

```
# /etc/init.d/init.crs start
```

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 start the instances:

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

Note: The downtime ends here.

- 21 Navigate to the folder that contains the installation program. On the second half of the cluster, enter the installrp script:

```
# ./installrp jupiter mercury
```

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

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

- 23 Restart the nodes in the second half of the cluster.

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

When the nodes in the second half of the cluster come up, all the GAB ports a, b, d, o, h, v, w and f will be online. All the CFS mount service groups also come online automatically.

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

If CRS is not controlled by VCS, enter the following command on each node of first of cluster to start CRS.

```
# /etc/init.d/init.crs start
```

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 start the instances:

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

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

Upgrading Veritas product with the Veritas Web-based installer

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

Note: Upgrading SF Oracle RAC with the Web-based installer is not supported.

To upgrade Veritas product

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 49.
- 3 Select **Install RP1**.
The installer detects the product that is installed on the specified system.
- 4 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Validate**.

- 5 Click **Next** to complete the upgrade.

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

- 6 Click **Finish**. After the upgrade, if the product is not configured, the web-based installer asks: "Do you want to configure this product?" If the product is already configured, it will not ask any questions.

If you are upgrading from 4.x, you may need to create new VCS accounts if you used native operating system accounts.

Performing a rolling upgrade using the installer

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime. Rolling upgrades are not compatible with phased upgrades. Do not perform "mixed" rolling upgrades with phased upgrades.

Rolling upgrades take two discrete phases. In the first, you upgrade the kernel packages with exception of VCS packages and agent packages. In the second, you upgrade the non-kernel packages, which are VCS packages and agents packages.

You can perform a rolling upgrade for the following products:

- Veritas Cluster Server
- Storage Foundation and High Availability
- Storage Foundation Cluster File System
- Storage Foundation Cluster File System and High Availability
- Storage Foundation for Oracle RAC

You can perform a rolling upgrade from 5.1 to 5.1 RP1 or from 5.1 P1 to 5.1 RP1.

Prerequisites for a rolling upgrades

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrades.
- Split up your clusters into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you logged in as superuser and have the media mounted.

- VCS must be running before performing the rolling upgrade.

Performing a rolling upgrade on kernel packages: phase 1

Note that in the following instructions that a sub-cluster can represent one or more nodes in a full cluster, but is represented by nodeA.

To perform the rolling upgrade on kernel packages: phase 1

- 1 On the first sub-cluster, start the installer for the rolling upgrade with the `-upgrade_kernelpkgs` option.

```
./installrp -upgrade_kernelpkgs nodeA
```

Review the EULA, if you accept its terms, enter **y** to proceed.

- 2 Note that if the boot-disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
- 3 Note the installation log location. The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.
- 4 The installer further replaces kernel components. Review the output.
- 5 The installer starts processes and brings all the service groups online.
- 6 Repeat step 1 to 5 on the second subcluster.

Performing a rolling upgrade on non-kernel packages: phase 2

You now upgrade the non-kernel packages..

To perform the rolling upgrade on non-kernel packages: phase 2

- 1 Start the installer for the rolling upgrade with the `-upgrade_nonkernelpkgs` option. Specify all the nodes in the cluster:

```
./installrp -upgrade_nonkernelpkgs nodeA nodeB nodeC...
```

Review the EULA, if you accept its terms, enter **y** to proceed.

- 2 Note the installation log location. The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1.
- 3 The installer upgrades non-kernel components. Review the output.

- 4 The installer starts processes and brings all the service groups online.
- 5 Manually check the cluster's status.

```
# hastatus -sum
```

Performing a rolling upgrade manually

You can perform a Split stack rolling upgrade—an independent upgrade of CFS/ODM/CVM/LLT/GAB/VXFEN/LMX/VCSMM and the VCS Engine ('had').

Split stack rolling upgrade—an independent upgrade of CFS/ODM/CVM/LLT/GAB/VXFEN/LMX/VCSMM and the VCS Engine ('had')

Review the following notes:

- It is possible to conduct Rolling Upgrade of one node at a time.
- Recommended for clusters of any number of nodes and Service Group distributions, including N+1 configurations.
- Failover Service Groups will incur downtime 2 times, during failover/failback.

To perform a split stack rolling upgrade—an independent upgrade of CFS/ODM/CVM/LLT/GAB/VXFEN/LMX/VCSMM and the VCS engine ('had')

- 1 Consider a four node SFRAC cluster. Identify sub-clusters to be upgraded together. A sub-cluster could even be just one of the nodes of the cluster.
- 2 Review cluster's system list. Confirm that each Service Group will eventually have a target node to run on, when sub-clusters are upgraded in a rolling fashion.
- 3 Verify that `/opt/VRTS/bin` and `/opt/VRTSodm/bin` are added to `PATH` variable.
- 4 Display the system list:

```
# hagrps -display ServiceGroup -attribute SystemList
```

- 5 On the sub-cluster to be upgraded, run module specific commands as below for LLT, GAB, VXFEN, LMX, VCSMM, CVM, CFS, ODM on one of the nodes of the sub-cluster to be upgraded, to get the current protocol version. This version need not be same for all modules.

```
# lltconfig -W
# gabconfig -W
# vxfenconfig -W
# lmxconfig -W
# vcsmmconfig -W
# vxctl protocolversion
# fsclustadm protoversion
# odmclustadm protoversion
```

- 6 On the sub-cluster to be upgraded, stop all the applications and resources that are not under VCS control but are still using CVM and CFS stack.
- 7 Switch the failover Service Groups from the sub-cluster to be upgraded, to the other sub-cluster. The following command needs to be run for each affected Service Group on each node where the Service Group is active, on the sub-cluster to be upgraded. You may also specify a target node for a given Service Group, as required. However there is a downtime to the failover Service Groups at this stage as part of the switch.

```
# hagrpswitch ServiceGroup -to target_system_name
```

- 8 Validate that the Service Groups are switched over as desired. In case the switch didn't succeed for any of the Service Groups, the user still has a window available to make any changes to the impacted Service Groups at this stage.
- 9 Unmount all vxfs file systems on the sub-cluster.
- 10 Stop 'had' on the sub-cluster to be upgraded, and switch any remaining failover Service Groups on this sub-cluster atomically.

```
# hastop -local -evacuate
```

Review the following notes:

- If all the Service Groups had switched over in step 6 itself, the 'evacuate' operation for the above command is idempotent.
- With the above step, it is ensured that if one of the nodes in the remaining sub-cluster goes down at this stage, the Service Groups that have already been moved to the remaining sub-cluster will not attempt to switch back to any of the nodes on the sub-cluster being upgraded. Any pending switches can also occur in this step.

- The parallel Service Groups on the nodes of the sub-cluster to be upgraded are brought down at this stage. They will continue to be available on the remaining sub-cluster.
 - CVM, CFS will also be stopped by VCS on the nodes of the sub-cluster being upgraded. They will continue to be available on the remaining sub-cluster.
- 11 Stop applications/resources that are outside VCS control and use VxFS, VxVM.
 - 12 Manually update the `/etc/vxfenmode`, `/etc/gabtab`, and `/etc/vcsmmtab` files to indicate the protocol version that the corresponding module in the new stack should talk to that on the older stack on each of the nodes. This protocol version is the same as the one obtained in step 5. For CVM, CFS and ODM, run the following commands on each of the nodes, to set the protocol version.

```
# vxdctl setversion N
# fsclustadm protoset N
# odmclustadm protoset N
```

where *N* is the protocol version derived in step 5.

This step ensures that the sub-clusters consistently communicate at the older protocol version should there be any intermediate node joins/leaves until the entire cluster is explicitly rolled over to communicate at the new version.

For example, for `/etc/vxfenmode`:

```
# cat /etc/vxfenmode
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3          - use scsi3 persistent reservation disks
# customized    - use script based customized fencing
# sybase        - use scsi3 disks in kernel but coordinate membership
# with Sybase ASE
# disabled      - run the driver but don't do any actual fencing
#
vxfen_mode=disabled
vxfen_protocol_version=10

# cat /etc/gabtab
/sbin/gabconfig -c -n4 -V33
```

- 13 Stop VXFEN, ODM, VCSMM, LMX, GAB and LLT in that order, on each of the nodes of the sub-cluster to be upgraded.
- 14 Simultaneously upgrade of all the components except the VCS Engine ('had') on the sub-cluster chosen for upgrade. VCS engine and agent related packages are not upgraded at this stage. CFS, ODM, CVM, LMX, VCSMM, GAB, LLT, VXFEN will be upgraded together.
 - Upgrade (use patchadd on Solaris) all the packages with new product version, except VCS and agent related packages on the sub-cluster being upgraded.
Some examples of the patch names for SPARC are: VRTSsfmh 141270-02, VRTSvxvm 142629-02, VRTSllt 143260-02/143261-02, VRTSgab 143262-02/143263-02, VRTSvxfen 143706-02/143707-02, VRTSdbac 143696-01/143697-01, VRTSob 43687-01, and VRTScps 143279-02.
 - Re-link oracle in case of SFRAC.
 - In reverse order, start the components that you previously stopped.
For example, on Solaris 10:

```
# svcadm enable llt
# svcs -a|grep llt
online          0:22:44  svc:/system/llt:default
# svcadm enable gab
# svcs -a|grep gab
online          0:25:07  svc:/system/gab:default
# svcadm enable vxodm
# svcs -a|grep vxodm
online          0:25:20  svc:/system/vxodm:default
# svcadm enable vxfen
# svcs -a|grep vxfen
online          0:25:37  svc:/system/vxfen:default
```

For Solaris 9, the commands are listed under /etc/init.d to start modules.
For example, to start LLT, use /etc/init.d/llt start, etc.

LLT and LMX start communicating with the new version automatically when they are started.

Start HA.

```
# hastart
```

- Once all the services are started, start use the `hastart` command to start HA. All ports should come up successfully and cluster start communication the other nodes.

- 15 Upgrade the remaining sub-cluster(s) one by one, per above procedure from step 4 onwards.
- 16 After each of the nodes are upgraded to the new product version, initiate a cluster-wide and across-the-stack rollover of the kernel stack to the new protocol version.
 - LLT and LMX are already at new protocol version at the end of step 14.
 - Run `gabconfig -R` on one of the nodes of the cluster being upgraded. This command will block until roll over is complete cluster wide. GAB also quiesces I/Os, which will result in flow control.
 - Run `vxfenconfig -R` on one of the nodes of the cluster being upgraded. Wait till the command returns.
 - Run `vcsmmconfig -R` on one of the nodes of the cluster being upgraded. Wait till the command returns.
 - Run `vxdtl upgrade` on the CVM master node of the cluster being upgraded.
 - Run `fsclustadm protoclear` to clear the set protocol version on all the nodes in the cluster.
 - Run `fsclustadm protoupgrade` from any node of cluster to upgrade the protocol version across the cluster.
 - Run `odmclustadm protoclear` to clear the set protocol version on all nodes.
 - Run `odmclustadm protoupgrade` on one of the nodes of the sub-cluster being upgraded.

While upgrading odmcluster protocol version, you might see a message like:

```
"Protocol upgrade precheck fails:  
    some nodes do not support multiple protocols"
```

You can ignore this message. The odm module is running on the latest version. You can verify this by using the following command on all the upgraded nodes:

```
# odmclustadm protoversion  
Cluster Protocol Versions:  
Node    #PROTOCOLS  CUR    PREF    FLAGS  
local:  3           3      -
```


- Reverse the changes done to `/etc/vxfenmode`, `/etc/gabtab`, and `/etc/vcsmntab` files in step 12 above.
- 17 Upgrade VCS engine ('had') to the new version. Perform one of the following procedures:
- Force stop 'had' and install the new version.
 - Force stop 'had' on all the nodes. There is no HA from this point onwards.

```
# hastop -all -force
```
 - Modify the VCS configuration to reflect version specific requirements if any.
 - Upgrade VRTSvcs and agent-related patches (the patches on SPARC):

```
VRTSvcs 143264-02  
VRTSvcsag 143265-02  
VRTScavf 143274-02  
VRTSvcsea 143276-02
```
 - Start VCS on all nodes. HA for the entire cluster is restored at this stage.
 - Upgrade 'had' in a phased manner. This procedure will reduce the overall HA downtime during the upgrade.
 - Divide the cluster into two sub-clusters. Upgrade the first sub-cluster.
 - Force stop VCS on the sub-cluster. There will be no HA for the sub-cluster being upgraded, from this step onwards.

```
# hastop -local -force
```
 - Modify the VCS configuration to reflect version specific requirements if any.
 - Upgrade VRTSvcs and agent-related patches (the patches on SPARC):

```
VRTSvcs 143264-02  
VRTSvcsag 143265-02  
VRTScavf 143274-02  
VRTSvcsea 143276-02
```
 - Force stop VCS on the remaining sub-cluster. There is no HA for the entire cluster from this point onwards.

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

- Start VCS on each of the nodes of the upgraded sub-cluster. VCS will not online the failover Service Groups at this time since they are autodisabled. Now HA is restored for the upgraded sub-cluster.

```
# hastart
```

Upgrading to 5.1 RP1 on a standalone system

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

To upgrade to 5.1 RP1 on a standalone system

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 3 If required, apply the OS kernel patches.
See “[System Requirements](#)” on page 12.
See *Sun Microsystems’* documentation for the procedures.
- 4 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df | grep vxfs
```

- 5 Unmount all Storage Checkpoints and file systems:

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

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

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

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

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

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

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

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

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

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

Verify that no volumes remain open:

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

9 Check if the VEA service is running:

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

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

10 Mount the 5.1 RP1 product disc and navigate to the folder that contains the installation program. Enter the `installrp` script:

```
# ./installrp nodename
```

11 If necessary, reinstate any missing mount points in the `/etc/vfstab` file.

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

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

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

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

14 Remount all VxFS file systems and Storage Checkpoints:

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

15 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

Verifying software versions

To verify if the Veritas patches are installed on your system, enter the following command:

```
# showrev -p|grep patch_id
```

Removing and rolling back

Roll back of the 5.1 RP1 to the release 5.1 version is not supported for certain products. It is recommended that you follow the steps in the following sections to remove all the installed Veritas software, and then perform a complete reinstallation of the release 5.1 software. You can roll back 5.1 RP1 to the release 5.1 version for Veritas Cluster Server.

Note: Symantec recommends using the following steps to roll back. There is no uninstallrp to roll back the patches.

- [Removing 5.1 RP1 from Veritas Cluster Server](#)
- [Removing 5.1 RP1 on SF or SFCFS](#)
- [Removing 5.1 RP1 on Storage Foundation for Oracle RAC](#)

Removing 5.1 RP1 from Veritas Cluster Server

Use the following procedure to remove VCS 5.1 RP1 from your cluster manually.

To remove 5.1 RP1 from VCS manually

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

```
# hagrps -state
```

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

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

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

```
# haconf -makerw
```

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

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

where *service_group* is the name of the service group. Note that 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 Make a backup copy of the current *main.cf* and all *types.cf* configuration files. For example, on one node in the cluster, type:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/types.cf.save
```

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

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

- 8 Shut down CmdServer. On each node, type:

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

- 9 Verify that VCS has shut down. On any node, type:

```
# /sbin/gabconfig -a
```

The output resembles: GAB Port Memberships Port a gen 23dc0001 membership 01 The output shows no membership for port h.

- 10 For Solaris 10, on nodes that run non-global zones, check if the non-global zones are in the running state. Boot the non-global zones that are not in the running state.

- Check the zone's state. On each node, type:

```
# zoneadm list -icv
```

- Boot the zone if it is not in the running state. On each node, type:

```
# zoneadm -z zone boot
```

where *zone* is the name of the non-global zone.

Note: Do not configure one or more Solaris zones to boot from the shared storage.

- 11 Unconfigure vxfen if the VCS cluster uses the fencing option. On each node, type:

```
# /sbin/vxfenconfig -U
```

- 12 Unload vxfen. On each node, perform the following steps:

- Identify the vxfen kernel module, for example:

```
# modinfo |grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.0MP3RP3)
```

- Unload vxfen using the module number.

```
# modunload -i 210
```

- 13 Unconfigure GAB. On each node, type:

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

- 14 Unload GAB. On each node, perform the following steps:

- Identify the GAB kernel module. For example:

```
# modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.0MP3RP3)
```

- Unload GAB using the module number:

```
# modunload -i 149
```

15 Unconfigure LLT. On each node, perform the following steps:

- Type:

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

- Type **y** on each node in response to the message.

16 Unload LLT. On each node, perform the following steps:

- Identify the LLT kernel module. For example:

```
# modinfo | grep llc
147 50ca4000 d6bc 110 1 llc (LLT 5.0MP3RP3)
```

- Unload LLT using the module number:

```
# modunload -i 147
```

17 Remove the VCS 5.1 RP1 patches. On each node, type:

- For Solaris SPARC 8:

```
# patchrm 139356-03
```

- For Solaris SPARC 9:

```
# patchrm 139357-03
```

- For Solaris SPARC 10:

```
# patchrm 142607-03
# patchrm 139359-03
# patchrm 139358-03
```

- For Solaris x64:

```
# patchrm 139361-03  
# patchrm 139360-03  
# patchrm 142608-03
```

Note: For Solaris SPARC 8, 9, 10, if you must remove the 5.1 RP1 Authentication Service patch (123722-02), you must uninstall the entire VCS product stack, then reinstall VCS.

- 18** Verify that the patches have been removed. On each node, type:

```
# showrev -p | grep VRTS
```

- 19** If the LLT, GAB, or VXFEN modules cannot be stopped or unloaded following the patch removal, reboot all nodes in the cluster.
- 20** If you do not perform step 19, start the VCS components manually. On each node, type:

```
# /sbin/lltconfig -c  
# /sbin/gabconfig -cx  
# /sbin/vxfenconfig -c  
# /opt/VRTSvcs/bin/hastart
```

You do not have to start vxfen unless you use the fencing option.

- 21** After VCS has started, perform the following steps:

- Verify all resources have been probed. On any node, type:

```
# hastatus -summary
```

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

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

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

- 22** Bring online the ClusterService service group, if necessary. On any node type:

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

where *system* is the node name.

Removing 5.1 RP1 on SF or SFCFS

You can use the following procedure to uninstall 5.1 RP1 on SF or SFCFS.

To uninstall 5.1 RP1 on SFor SFCFS

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 3 Unmount all Storage Checkpoints and file systems:

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

- 4 Check if the root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk. For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Note: Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the rootdg disk group in addition to the root disk for `vxunroot` to succeed.

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is restarted from the unencapsulated root disk.

- 5 Enter the following command to check if any VxFS file systems are mounted:

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems:

```
# umount /filesystem
```

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

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

- Use the vxrvrg stop command to stop each RVG individually:

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

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

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

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

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

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

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

To verify that no volumes remain open, enter the following command:

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

- 9 Stop VCS along with all its resources. Then, stop the remaining resources manually:

- For Solaris 9:

```
# /etc/init.d/vcs stop
```

- For Solaris 10:

```
# svcadm disable vcs
```

- 10** If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

```
# rm /etc/vxfenmode
```

- 11** Unmount /dev/odm:

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

- 12** Unload the ODM module:

```
# modinfo | grep odm  
# modunload -i odm_mod_id
```

- 13** Unload the cluster fencing (vxfen) module:

- For Solaris 9:

```
# /etc/init.d/vxfen stop  
# modinfo | grep vxfen  
# modunload -i vxfen_mod_id
```

- For Solaris 10:

```
# svcadm disable vxfen  
# modinfo | grep vxfen  
# modunload -i vxfen_mod_id
```

- 14** Stop GAB and LLT in the following order:

For Solaris 9:

```
# /etc/init.d/gab stop  
# /etc/init.d/llt stop
```

For Solaris 10:

```
# svcadm disable -t gab  
# svcadm disable -t llt
```

- 15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

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

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

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

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

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

Removing 5.1 RP1 on Storage Foundation for Oracle RAC

You can use the following procedure to uninstall the 5.1 RP1 on Storage Foundation for Oracle RAC systems.

To uninstall the 5.1 RP1 on SF Oracle RAC

- 1 Stop Oracle and CRS on each node of the cluster.
 - If CRS is controlled by VCS, log in as superuser on each system in the cluster and enter the following command:

```
# hastop -local
```

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

For 10gr2:

```
# /etc/init.d/init.crs stop
```

For 11gr2:

```
# /etc/init.d/init.ohasd stop
```

Unmount all VxFS file system used by a database or application and enter the following command to each node of the cluster:

```
# hastop -local
```

2 Stop cluster fencing, VCSMM, LMX, ODM, and GAB:

For Solaris 9:

```
# /etc/init.d/vxfen stop
# /etc/init.d/vcsmm stop
# /etc/init.d/lmx stop
# /etc/init.d/odm stop
# /etc/init.d/gab stop
```

For Solaris 10:

```
# svcadm disable -t vxfen
# svcadm disable -t vcsmm
# svcadm disable -t lmx
# svcadm disable -t vxodm
# svcadm disable -t gab
```

3 On each node, unload the vxfen, LMX, GAB, LTT, VCSMM, GMS, and GLM kernel modules if they are still loaded.

- Verify if the vxfen kernel module is loaded. For example:

```
# modinfo|grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.1 RP1)
```

If the vxfen kernel module is loaded then unload it. For example:

```
# modunload -i 210
```

- Verify if the LMX kernel module is loaded. For example:

```
# modinfo | grep lmx
257 ffffffff0444000 13f48 257 1 lmx (LLT Mux 5.1 RP1)
```

If the LMX kernel module is loaded then unload it. For example:

```
# modunload -i 257
```

- Verify if the VCSMM kernel module is loaded. For example:

```
# modinfo | grep vcsmm
312 78bc0000 43ae8 293 1 vcsmm (VRTSvcsmm 5.1 RP1)
```

If the VCSMM kernel module is loaded then unload it. For example:

```
# modunload -i 312
```

- Verify if the GMS kernel module is loaded. For example:

```
# modinfo | grep gms
253 ffffffff040c000 4550 244 1 vxgms
(VxGMS 5.1.0.0,REV=13Sep2009 (So))
```

If the GMS kernel module is loaded then unload it. For example:

```
# modunload -i 253
```

- 4 Verify if the GLM kernel module is loaded. For example:

```
# modinfo | grep glm
247 fffffefb2a000 27390 238 1 vxglm (VxGLM 5.1,REV=13Sep2009
SunOS 5)
```

If the GLM kernel module is loaded then unload it. For example:

```
# modunload -i 247
```

- 5 Verify if the GAB kernel module is loaded. For example:

```
# modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.1 RP1)
```

If the GAB kernel module is loaded then unload it. For example:

```
# modunload -i 149
```

- 6 Stop LLT:

For Solaris 9:

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

For Solaris 10:

```
# svcadm disable -t llt
```

- Verify if the LLT kernel module is loaded. For example:

```
# modinfo|grep llt
147 50ca4000 d6bc 110 1 llt (LLT 5.1 RP1)
```

If the LLT kernel module is loaded then unload it. For example:

```
# modunload -i 147
```

- 7 To uninstall only 5.1 RP1 patches from SF Oracle RAC cluster, execute following commands on each node of the cluster:

- For Solaris 9:

```
# patchrm 141270-02
# patchrm 142629-02
# patchrm 142631-02
# patchrm 142633-02
# patchrm 142634-02
# patchrm 143260-02
# patchrm 143261-02
# patchrm 143262-02
# patchrm 143263-02
# patchrm 143264-02
# patchrm 143265-02
# patchrm 143270-02
# patchrm 143271-02
# patchrm 143273-02
# patchrm 143274-02
# patchrm 143276-02
# patchrm 143279-02
# patchrm 143687-01
# patchrm 143696-01
# patchrm 143697-01
# patchrm 143706-02
# patchrm 143707-02
```

- For Solaris 10:

```
# patchrm 141752-02
# patchrm 142630-02
# patchrm 142632-02
# patchrm 142635-02
# patchrm 143266-02
# patchrm 143267-02
# patchrm 143268-02
# patchrm 143269-02
# patchrm 143272-02
# patchrm 143275-02
# patchrm 143277-02
# patchrm 143280-02
# patchrm 143693-01
# patchrm 143698-01
# patchrm 143708-02
```

- 8 To uninstall only 5.1 RP1 patches from SF Oracle RAC cluster, execute following commands on each node of the cluster:
- 9 After uninstalling the packages, refer to the Storage Foundation for Oracle RAC 5.1 Installation and Configuration Guide to reinstall the 5.1 software.
- 10 After removing the patches, reboot the nodes:

```
# /usr/sbin/shutdown -g0 -y -i6
```

Documentation addendum

The following sections contain additions to current documents.

Disk agent

Monitors a physical disk or a partition. You can use the Disk agent to monitor a physical disk or a slice that is exported to LDomS (available using LDomS 1.2 or later). For LDomS with a physical disk or slice based boot image, a dependency must exist between the guest domain and primary domain. You configure the primary domain as the master of the guest domain.

Perform the following:


- Set the failure-policy of primary (control) domain to stop. For example, in the primary domain enter the following command to set the dependent domain to stop when the primary domain faults:

```
# ldm set-domain failure-policy=stop primary
```

- Set the primary domain as the master for the guest domain.

```
# ldm set-domain master=primary guest1dom
```

Figure 1-1 Sample service group that includes a Disk resource on Solaris



Agent functions

Monitor—Performs read I/O operations on the raw device to determine if a physical disk or a partition is accessible.

State definitions

ONLINE—Indicates that the disk is working normally

FAULTED—Indicates that the disk has stopped working or is inaccessible.

UNKNOWN—Indicates that a problem exists either with the configuration or the ability to determine the status of the resource.

Attribute

The Disk agent has one required attribute.

Partition—Indicates which partition to monitor. Specify the partition with the full path beginning with a slash (/).

If this path is not specified, the name is assumed to reside in `/dev/rdisk/`.

Example: `"/dev/rdisk/c2t0d0s2"`

Type and dimension: string-scalar

Resource type definition

The following is the agent's resource type definition.

```
type Disk (  
    static int OfflineMonitorInterval = 60  
    static str ArgList[] = { Partition }  
    static str Operations = None  
    str Partition  
)
```

Using the preonline_vvr trigger for RVGLogowner resources

For VCS configurations that use RVGLogowner resources, perform the following steps on each node of the cluster to enable VCS control of the RVGLogowner resources. For a service group that contains a RVGLogowner resource, change the value of its PreOnline trigger to 1 to enable it.

To enable the PreOnline trigger from the command line on a service group that has an RVGLogowner resource

- ◆ On each node in the cluster, perform the following command:

```
# hagrps -modify RVGLogowner_resource_sg PreOnline 1 -sys system
```

Where the service group is the service group that contains the RVGLogowner resource (*RVGLogowner_resource_sg*). The *system* is the name of the node where you want to enable the trigger.

On each node in the cluster, merge the preonline_vvr trigger into the default triggers directory.

To merge the preonline_vvr trigger

- ◆ On each node in the cluster, merge the preonline_vvr trigger to the /opt/VRTSvcs/bin/triggers directory.

```
# cp /opt/VRTSvcs/bin/sample_triggers/preonline_vvr \  
/opt/VRTSvcs/bin/triggers
```

Refer to the sample configurations directory for samples of how to enable these triggers (/opt/VRTSvcs/bin/sample_triggers.)

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.

Note: If you want to install Oracle RAC 10g or Oracle RAC 11g Release 1, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide (5.1)* for instructions.

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 108.
2. Install Oracle RAC 11g Release 2.
See [“Installing Oracle Grid Infrastructure using the Oracle Universal Installer”](#) on page 118.
See [“Configuring LLT links in the GPnP profile”](#) on page 121.
See [“Installing the Oracle RAC database using the Oracle Universal Installer”](#) on page 122.
3. Complete the following post-installation tasks:

- Relink the SF Oracle RAC libraries with Oracle RAC:
See [“Linking the ODM library”](#) on page 123.
- Configure the CSSD resource.
See [“Configuring the CSSD resource manually”](#) on page 124.
- Modify the Oracle RAC configuration to prevent automatic startup of Oracle Clusterware.
See [“Preventing automatic startup of Oracle Clusterware”](#) on page 125.
- Create the Oracle RAC database.
See [“Creating the Oracle RAC database”](#) on page 126.
- Modify the Oracle RAC configuration to prevent automatic startup of the Oracle database.
See [“Preventing automatic database startup”](#) on page 126.
- Configure VCS service groups for Oracle RAC.
See [“Configuring VCS service groups for Oracle RAC”](#) on page 127.

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.

Some preparatory steps are similar to those documented in the version 5.1 guide and are therefore not described here. Only instructions specific to Oracle RAC 11g Release 2 are described in this section.

Complete the following preparatory tasks. For instructions, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide (5.1)*, Chapter “Before installing 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). The Oracle Universal Installer for Oracle RAC11g Release 2 does not

support OCR and voting disk on raw volumes. You can switch to raw volumes after the installation of the Grid Infrastructure software is complete.

Make sure that you perform the steps provided in the procedure for creating the OCR and voting disk directories on CFS.

- 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 109.
- Create Oracle Grid Infrastructure and Oracle database home directories manually.
See [“Creating Oracle Grid Infrastructure and Oracle database home directories manually”](#) on page 109.
- Set up user equivalence on cluster nodes.
See [“Setting up user equivalence”](#) on page 117.
- Synchronize time settings on cluster nodes.
See [“Synchronizing time settings on cluster nodes”](#) on page 117.
- Edit the user profiles.
See [“Editing the user profiles”](#) on page 117.
- Verify whether the Veritas Membership library is linked to the Oracle library.
See [“Verifying whether the Veritas Membership library is linked to the Oracle library”](#) on page 118.

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.

Creating Oracle Grid Infrastructure and Oracle database home directories manually

You can create the Oracle directories on the local file system, on Veritas file system, or on a 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
Oracle Grid Infrastructure Home Directory (GRID_HOME)	<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>
Oracle base directory (ORACLE_BASE)	<p>The base directory that contains all the Oracle installations. 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. The permissions on the Oracle base directory must be at least 750.</p>
Oracle home directory (ORACLE_HOME)	<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>

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):

```
# 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 /u01/grid
# chown -R grid:oinstall /u01/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 diskgroup `bindg_hostname` on each node:

```
# vxdg init bindg_galaxy Disk_1
```

2 Create a volume `binvol_hostname` on each node:

```
# vxassist -g bindg_galaxy make binvol_galaxy 12G
```

3 Create a filesystem with the volume, `binvol_hostname`, on each node.

```
# mkfs -F vxfs /dev/vx/rdisk/bindg_galaxy/binvol_galaxy
```

4 Mount the filesystem on each node:

```
# mount -F vxfs /dev/vx/dsk/bindg_galaxy/binvol_galaxy
/u01
```

5 Create the following directories for Oracle RAC on each node:

GRID_HOME directory:

```
# mkdir -p /u01/grid
```

ORACLE_BASE directory:

```
# mkdir -p /u01/app/oracle
```

ORACLE_HOME directory:

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

6 Change the ownership and permissions on each node:

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

7 Add an entry for the filesystem in the `/etc/vfstab` file on each node:

Edit the `/etc/vfstab` file, list the new file system, and specify "yes" for the mount at boot column for each node:

```
# device      device      mount  FS      fsck  mount      mount
# to mount to fsck  point  type   pass  at  boot options
#
.
/dev/vx/dsk/bindg_galaxy/binvol_galaxy
/dev/vx/rdisk/bindg_galaxy/binvol_galaxy \
/u01 vxfs 1 yes -
```

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`:

```
# vxvg -s init bindg Disk_1
```

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

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

- 3 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      /u01/app/oracle
                  /u01/app/oracle is on local storage.

$GRID_HOME        /u01/grid
                  /u01 is on local storage.
                  /u01/grid is on shared storage.

$ORACLE_HOME      /u01/app/oracle/product/11.2.0/dbhome_1
                  /u01/app/oracle/product/11.2.0 is on local storage.
                  /u01/app/oracle/product/11.2.0/dbhome_1 is on shared
                  storage.
```

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

- 4 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
```

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

```
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/crsbinvol \
/u01/grid
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/orabinvol \
/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 `/u01/app/oracle` must be owned by `oracle:oinstall`, otherwise `/u01/app/oracle/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

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

7 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 114.

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 "/u01/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 = "/u01/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.

Synchronizing time settings on cluster nodes

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

Editing the user profiles

Edit the Oracle user and grid user `.profile` file to set the paths to `ORACLE_BASE` and `ORACLE_HOME` on each node.

In the following sample procedure, the shell environment is `ksh`, 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 LD_LIBRARY_PATH=$ORACLE_HOME/lib
export PATH=$PATH:$ORACLE_HOME/bin
export CLASSPATH=$CLASSPATH:$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 ORACLE_BASE=/u01/app/oracle
export ORACLE_HOME=/u01/app/oracle/product/11.2.0/dbhome_1
export GRID_HOME=/u01/grid
export LD_LIBRARY_PATH=$GRID_HOME/lib
export PATH=$PATH:$GRID_HOME/bin:$ORACLE_HOME/bin
export CLASSPATH=$CLASSPATH:$ORACLE_HOME/jlib:\
$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib
```

- 2 Apply the profile changes:

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

Verifying whether the Veritas Membership library is linked to the Oracle library

The Veritas Membership library (VCSMM) must be linked with the Oracle library to enable coordinated exchange of cluster membership information and protection of data integrity. Oracle uses the linked skgxn library (libskgxn) to make ioctl calls to VCSMM, which in turn obtains membership information for clusters and instances.

Verify that the library `/opt/ORCLcluster/lib/libskgxn2.so` is linked to the `/opt/VRTSvcs/rac/lib/libskgxn2_64.so` library:

```
# ls -l /opt/ORCLcluster/lib
libskgxn2.so -> /opt/VRTSvcs/rac/lib/libskgxn2_64.so
```

If the link does not exist, create a symbolic link to the Oracle library as follows:

```
# ln -s /opt/VRTSvcs/rac/lib/libskgxn2_64.so \
/opt/ORCLcluster/lib/libskgxn2.so
```

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
```

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

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

- 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.
- 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: 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
bge0 10.2.156.0          global      public
bge1 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 \
bge2/192.168.12.0: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 \
bge2/192.168.12.0:cluster_interconnect
# $GRID_HOME/bin/oifcfg setif -global \
bge1/192.168.2.0:cluster_interconnect
```

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

```
# $GRID_HOME/bin/oifcfg getif
bge0 10.2.156.0      global      public
bge1 192.168.12.0   global      cluster_interconnect
bge2 192.168.12.0   global      cluster_interconnect
bge1 192.168.2.0    global      cluster_interconnect
bge2 192.168.2.0    global      cluster_interconnect
```

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

- 4 Restart the Oracle database on all nodes by running the following commands on one of the nodes in the cluster.

To stop the Oracle database:

```
$ srvctl stop database -d db_name
```

To start the Oracle database:

```
$ srvctl start database -d db_name
```

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 Install database software only .
Node selection	Select Real Application Clusters database installation . Select the nodes on which the Oracle RAC database software must be installed.
Select database edition	Select Enterprise Edition .
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
```

Linking the ODM library

Perform the steps in the procedure on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, copy the libraries on one node only. Use the mount command to check that the file system containing the Oracle libraries are mounted.

To link the Veritas ODM library

- 1 Log in as the Oracle user.
- 2 Change to the \$ORACLE_HOME/lib directory:

```
$ cd $ORACLE_HOME/lib
```

- 3 Back up Oracle's ODM library:

```
$ mv libodm11.so libodm11.so.oracle-`date +%m_%d_%Y-%H_%M_%S`
```

- 4 Link the Veritas ODM library with Oracle's `libodm` library:

For Solaris SPARC:

```
$ ln -s /usr/lib/sparcv9/libodm.so libodm11.so
```

For Solaris x64:

```
$ ln -s /usr/lib/amd64/libodm.so libodm11.so
```

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 -modify cssd OnlineWaitLimit 5
```

4 Enable the CSSD resource:

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

5 Set the dependency of the CSSD resource on the CFSSMount 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 (5.1)*.

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

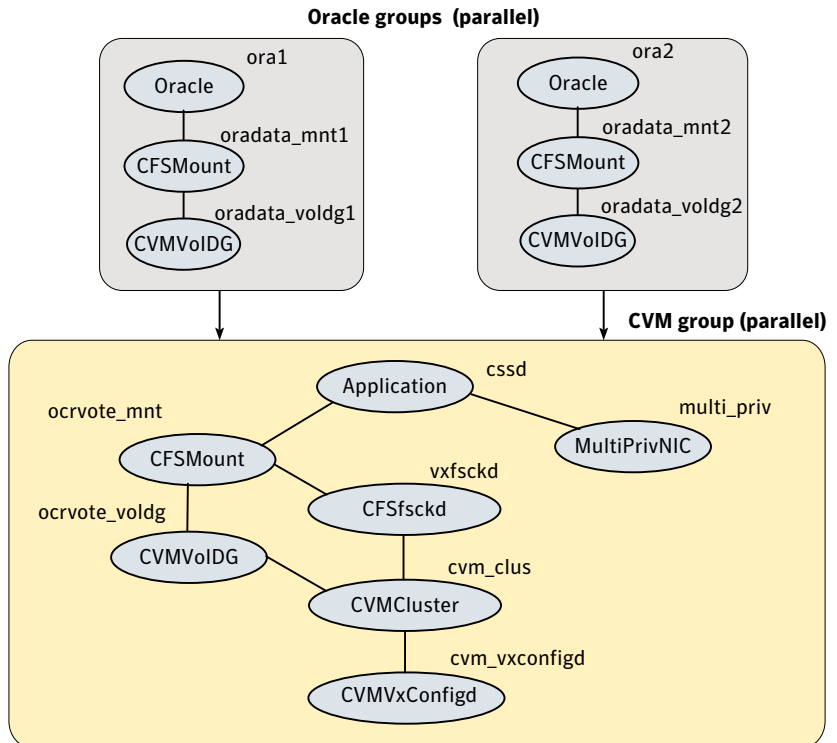



Figure A-2 illustrates a service group configuration without the VCS Oracle agent.

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 129.

- Use the CLI
See [“Using the CLI to configure VCS service groups for Oracle RAC”](#) on page 130.

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 permissions on the VCS configuration file: `# 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 129.
- Start VCS after modifying the configuration file
See [“To start VCS after modifying the configuration file”](#) on page 130.

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 one of the systems 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 CVMGroup 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 2t.

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 133.

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 134.
2. Upgrade Oracle RAC binaries.
See [“Upgrading Oracle RAC binaries”](#) on page 135.
3. Migrate the Oracle RAC database.
See [“Migrating the Oracle RAC database”](#) on page 135.
4. Complete the post-upgrade tasks.
See [“Performing post-upgrade tasks”](#) on page 135.

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 Clusterware 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 5.1 RP1
- 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
# hagrps -freeze oracle_group -persistent
```

- 4 Freeze the cvm service group to prevent VCS from reporting the resource as faulted when Oracle Clusterware is stopped and started during the upgrade:

```
# hagrps -freeze cvm_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.

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 [“Linking the ODM library”](#) on page 123.

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 135.

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 125.

- 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 126.

- 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
```