

Veritas Storage Foundation™ and High Availability Solutions Read This First

AIX

5.0 Maintenance Pack 3 Rolling Patch
5

Veritas Storage Foundation and High Availability Solutions Read This First 5.0 Maintenance Pack 3 Rolling Patch 5

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: 5.0MP3RP5

Document version: 5.0MP3RP5.0

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

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

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

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

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

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

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

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

North America and Latin America supportsolutions@symantec.com

Documentation

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

docs@symantec.com

About Symantec Connect

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

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

Veritas Storage Foundation and High Availability Solutions: Read This First

This document includes the following topics:

- [About this document](#)
- [System requirements](#)
- [Storage Foundation High Availability fixed issues](#)
- [Storage Foundation and High Availability known issues](#)
- [Software limitations](#)
- [Changes in behavior for Storage Foundation High Availability](#)
- [Downloading the rolling patch archive](#)
- [Filesets included in this rolling patch](#)
- [Installing the Veritas software for the first time](#)
- [Prerequisites for upgrading to 5.0 MP3 RP5](#)
- [Supported upgrade paths](#)
- [Upgrading 5.0 MP3 to 5.0 MP3 RP5](#)
- [Verifying software versions](#)
- [Removing 5.0 MP3 RP5](#)
- [Documentation addendum](#)

- [Documentation errata](#)

About this document

This document provides release information about the products in the Veritas Storage Foundation and High Availability 5.0 Maintenance Pack 3 Rolling Patch 5 (SFHA 5.0 MP3 RP5) release.

For the latest information on updates, patches, and known issues regarding this release, see the following TechNote on the Symantec Technical Support website:

<http://www.symantec.com/docs/TECH46478>

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

See also the *Veritas Storage Foundation™ for Oracle® RAC 5.0 MP3 RP5 Application Note: Installing or upgrading to Oracle RAC 11g Release 2*.

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

- *Veritas Cluster Server™ 5.0 MP3 Release Notes*
- *Veritas Storage Foundation™ 5.0 MP3 Release Notes*

Note: The Veritas Storage Foundation Cluster File System 5.0 MP3 Release Notes information is located in the *Veritas Storage Foundation 5.0 MP3 Release Notes*.

- *Veritas Storage Foundation™ for Oracle® RAC 5.0 MP3 Release Notes*

System requirements

This section describes the system requirements for this release.

Supported operating systems

The 5.0 MP3 RP5 release operates on the architectures and operating systems shown below:

- AIX 5.3 TL7 with SP2 or any higher TLs
- AIX 6.1 TL0 with SP4 or any higher TLs

This release adds support for POWER7 systems to run in native POWER7 mode.

DB2 support

This release of Storage Foundation for DB2 offers support for DB2 9.7 in addition to the DB2 database versions that are supported in 5.0 MP3 release.

Oracle support

This release of Storage Foundation for Oracle and Veritas Cluster Server (VCS) offers support for the following in addition to the Oracle database versions that are supported in 5.0 MP3 release:

- Oracle 11g Release 2 (11.2.0.1.0)
- Oracle 11g Release 2 (11.2.0.2.0)

See the *Veritas Storage Foundation™ for Oracle® RAC 5.0 MP3 RP5 Application Note: Installing or upgrading to Oracle RAC 11g Release 2*.

This release of Storage Foundation for Oracle RAC offers support for the following in addition to the Oracle database versions that are supported in 5.0 MP3 release:

- Oracle 11g Release 1 (11.1.0.7)
- Oracle 11g Release 2 (11.2.0.1.0)

For latest information on support for Oracle RAC releases, visit the following link:

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

Storage Foundation High Availability fixed issues

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

- [Veritas Volume Manager fixed issues](#)
- [Veritas File System fixed issues](#)
- [Storage Foundation Cluster File System fixed issues](#)
- [Storage Foundation for Oracle fixed issues](#)
- [Storage Foundation for DB2 fixed issues](#)
- [Veritas Cluster Server fixed issues](#)
- [Storage Foundation for Oracle RAC fixed issues](#)
- [Veritas Cluster Server agents for Veritas Volume Replicator fixed issues](#)
- [Veritas Enterprise Administrator fixed issues](#)

Veritas Volume Manager fixed issues

Table 1-1 describes fixed issues in the Veritas Volume Manager 5.0 MP3 RP5 release.

Table 1-1 Veritas Volume Manager 5.0 MP3 RP5 fixed issues

Fixed issues	Description
2310284	Fixed the issue that for VxVM versions prior to 5.1SP1, don't allow CDS disk initialization for LUN size > 1TB
2226813	Fixed the issue with VVR: rlinks remain disconnected with UDP protocol if data ports are specified
2199496	Fixed the issue that Data Corruption issue with "site mirror" Campus Cluster feature
2197254	Fixed the issue that while creating volumes on thinrlm disks, the option "logtype=none" does not work with vxassist command
2188590	Fixed the issue that interlock acquired by SLAVE node for a read on a DCL object leads to IO hang when the node becomes MASTER before completion of the read
2183984	Fixed the system panic in dmp_update_stats routine
2181631	Fixed the issue that striped-mirror volume cannot be grown across sites with -oallowspansites with DRL
2176601	SRDF-R2 devices are seen in error state when devices are in write-protected mode
2160199	Fixed the issue that master takeover fails as the upcoming Master could not import shared diskgroup
2158438	Fixed the issue that vxsnap restore operation for 500 volumes spits garbage strings and sometime dumps core
2152830	Fixed the issue that in multilevel clone disks environment, regular diskgroup import should be handled properly and in case of diskgroup import failure, it should report correct error message
2141176	Fixed the panic: BAD TRAP: type=31 in dmp_decode_modmap_dmpnode routine
2139179	Fixed the issue that SSB check invalid when lun copy
2133503	Fixed the issue that renaming enclosure results in dmpevents.log reporting Mode for Enclosure has changed from Private to Private

Table 1-1 Veritas Volume Manager 5.0 MP3 RP5 fixed issues (*continued*)

Fixed issues	Description
2131814	Fixed the issue with VVR:System panic due to corrupt sio in _VOLRPO_REMOVE
2122009	vxddladm list shows incorrect hba information after running vxconfigd -k
2112568	Fixed the issue that system panicked by vxconfigd daemon in voldco_or_acmbuf_to_pvmbuf routine
2105722	Fixed the issue with VVR:I/O hang on Primary with link-breakoff snapshot
2105547	Fixed the issue that tag meta info records are not cleaned-up during DGSJ operation and leading to huge delay in DGSJ operation after few iterations
2094685	Fixed the issue that diskgroup corruption following an import of a cloned BCV image of a SRDF-R2 device
2081515	vxresize command hung due to insufficient memory
2076700	Fixed the issue with VVR:Primary panic due to NULL pointer dereference
2070531	Fixed the issue with campus cluster: Couldn't enable siteconsistency on a DCL volume, when trying to make the disk group and its volumes siteconsistent
2060785	Fixed the issue with VIS: Panic in vol_rv_merge_config
2055609	Fixed the issue that allocation specification should be propagated for DCO during grow operation
2052459	Fixed the issue that DMP failed to register one path resulted in CFS mount failure on slave node
2040150	Fixed the issue that existence of 32 or more keys per LUN leads to loss of SCSI3 PGR keys during cluster reconfiguration
2038735	Fixed the issue that incorrect handling of duplicate objects resulting in node join failure and subsequent panic
2038137	Fixed the issue that system panics if the I/O breakup routine is called recursively
2036929	Fixed the issue that renaming a volume with an attached link object can break linked snapshots
2034564	Fixed the issue that I/Os hung in serialization after one of the disk which formed the raid5 volume was pulled out

Table 1-1 Veritas Volume Manager 5.0 MP3 RP5 fixed issues (*continued*)

Fixed issues	Description
2029735	Fixed the system panic in volobject_iogen routine due to null gio_object
2029480	Fixed the diskgroup join failure with error "Configuration too large for configuration copies" renders source diskgroup into inconsistent state
2027831	Fixed the issue that 'vxdg free' not reporting free space correctly on CVM master. vxprint not printing DEVICE column for SDs
2025593	Fixed the issue that vxdg join hang/failure due to presence of non-allocator inforecords and when tagmeta=on
2020373	Fixed the issue that system crashes (due to page fault) in gendmpopen while running the test suite
2016129	Fixed the issue that need to disable vxesd when PP is present on the system
2015577	Fixed the issue that VVR init scripts need to exit gracefully if VVR license not installed
2010426	Fixed the issue that settag and rmtag do not handle wrong enclosure name
2009439	Fixed the issue with CVR:Panic in vol_ru_check_limits routine while running the stress test suite
1982715	Fixed the issue with vxclustadm dumping core in realloc
1974393	Fixed the issue that avoiding cluster hang when the transaction client timed out
1960341	Fixed the issue that toggling of naming scheme is not properly updating the DM records in incore database of vxconfigd
1954062	Fixed the issue that vxrecover results in system crash
1946460	Fixed the issue that reattchsite is not doing da-dm assosiation on slave after diskgroup deport-import, when disks are in detached state
1933375	Changed value of "voliomem_chunk_size" should be align with page-size granularity
1831634	Fixed the issue with CVR:Wrong sending sibling count causing replication hang, which can result in I/O hang
1589715	Fixed the issue that vxconfigd dumps core, after vxdmpadm getportids ctrl=ctrl_name on a disabled ctrl

Table 1-1 Veritas Volume Manager 5.0 MP3 RP5 fixed issues (*continued*)

Fixed issues	Description
1586095	Fixed the issue with DMP:'vxdumpadm include' failed for including excluded dmpnodename
1513385	Fixed the issue with VVR:Primary panic during autosync or dcm replay
1426480	Fixed the issue with VOLCVM_CLEAR_PR ioctl does not propagate the error returned by DMP to the caller
1946941	vxsnap print shows incorrect year
1441406	'vxdisk -x list' displays wrong DGID
1482555	vxconfigd hung when unloading apm modules vxdump_unload_all_apms()
1315575	On linux, If SNIA HBA API is not available, then hba information should be gathered from proc or sys fs.
2022148	Hang on VxFS fs which are part of RVG.
1933528	During Dynamic reconfiguration vxvm disk ends up in error state after replacing physical LUN.
1669719	After removing LUN from array side , listenclosure all output doesn't list that array enclosure
1998447	Vxconfigd dumped core due to incorrect handling of signal
2016129	DDL: Tunable to disable OS event monitoring by vxesd
2052203	Master vold restart can lead to DG disabled and abort of pending transactions.
2012016	Slave node panics while vxrecovery is in progress on master
2232789	supporting NetApp Metro Cluster
1829285	vxconfigd coredumps while assigning unique native name to a disk
2201149	DMP should try all possibilities to service I/O upon receipt of a SCSI illegal request following HBA fault
2015570	File System read failure seen on space optimized snapshot after cache recovery

[Table 1-2](#) describes fixed issues in the Veritas Volume Manager 5.0 MP3 RP4 release, which are included in this release.

Table 1-2 Veritas Volume Manager 5.0 MP3 RP4 fixed issues

Fixed issues	Description
2040150	Loss of SCSI-3 PGR Keys during cluster reconfiguration if there are 32 or more keys per LUN.
2029735	System panic in volobject_iogen due to null gio_object
2029480	DG join failure with error "Configuration too large for configuration copies" renders source DG into inconsistent state (can't re-imported).
2021737	HDS TrueCopy S-VOL type devices displayed as "error" in vxdisk list whilst read-only
2002703	Warning message for the open failure is displayed as read-only device is opened using default RW flags
1999004	I/Os hang in drl_busyq in VxVM on linked-based snapshot
1998606	vxconfigd memory allocation failure occurs when lots of DGs and Volumes are present into the system.
1993953	CVM Node unable to join in Sun Cluster environment selecting wrong coordinator
1992872	Multiple issues with vxdisk resize code
1992537	6 node 50mp1rp5hf1 cluster had a DiskGroup Agent panic. Can see vxsnap threads started at same time vxdg/vxprint stopped working
1989057	Cannot restore root file system that was backed up by mksysb
1982715	vxclustadm dumping core in realloc
1972851	DG import failed after VxVM upgrade with DG version doesn't support feature error
1960341	Toggling of naming scheme is not properly updating the DM records in incore database of vxconfigd
1956777	VVR:Stale references in prpendingq causes random corruption
1952197	Running vxtrace against a volume shows response times as negative
1950328	vxfmrshowmap: dumping core while freeing the allocated memory
1947089	vxdisk resize cannot handle over 1TB gpt labelled disk as expected
1946939	CVM: Panic during master takeover, when there are cache object I/Os being started on the new master

Table 1-2 Veritas Volume Manager 5.0 MP3 RP4 fixed issues (*continued*)

Fixed issues	Description
1946936	CVM: IO hangs during master takeover waiting for a cache object to quiesce
1939432	2 TP:reclaim with raid5 volumes caused system panic
1938907	5.0MP4:Linux:"vxdmpadm getportids" is not showing pWWN no.
1936611	vxconfigd core dump while splitting a dg
1920761	I/O is not going through the disk after connecting the storage back to master node in local detach policy
1915356	I/O stuck in vxvm caused cluster node panic
1911546	vxrecover hung after hitting e1909954 and reboot of all nodes
1911137	Possible cache object corruption due to CREC and shadow tree updates conflicting
1907796	Corrupted Blocks in Oracle after Dynamic LUN expansion and vxconfigd core dump
1901827	vxdg move failed silently and drops disks.
1899688	[VVR]Every I/O on smartsync enabled volume under VVR leaks memory
1897007	vxesd dumps core when it starts
1884070	[VVR]: when running iotest on volume, primary node runs out of memory
1881639	Node panicked while testing recovery of space optimized snapshot based thin provisioning volume after unplug / plug of FC cable
1881336	VVR: Primary Panic in vol_ru_replica_sent()
1874034	Race between modunload and an incoming IO leading to panic.
1872743	Layered volume is not startable due to duplicate rid in vxrecover maintained volume_list
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
1853644	vxvm vxcdsconvert fails when a diskmedia name starts with a number
1846165	Data corruption seen on cdsdisks in several customer cases

Table 1-2 Veritas Volume Manager 5.0 MP3 RP4 fixed issues (*continued*)

Fixed issues	Description
1825516	Unable to initialize and use ramdisk for VxVM use
1825270	Need for dmp_revive_paths() in dmp reconfiguration/restore_demon code path.
1810655	When vxesd is enabled, dmp/dr procedure with PowerPath panics the system
1805669	TP : vxdisk reclaim should skip cache object update
1797508	vxconfigd level join hang when a slave rejoin is followed by a master-takeover
1797203	NULL pointer dereference panic :vxio:vol_cvol_dshadow1_done()
1766931	During boot time at VxVM start-up, paths keep getting enabled/disabled for BCV/BCV-NR devices in a loop resulting in system hang
1762344	vxconfigd hang on one of cluster node when performed some switch operations + vxdisk scandisks (dmpcert run)
1747275	SFRAC/CVR - panic in _VOLRPQ_APPEND: corrupted queue at 3002c12cc28
1720155	(Equallogic iSCSI) vxfcntlsthaw consistently failed while using dmp devices
1677149	vxesd core dump in strncmp() during boot
1671264	vxconfigrestore fails when only one disk is under the dg.
1668978	ASL Request for Hitachi USPV HAM
1665400	CCT:'vxsnap refresh' hung for long time, as vxconfigd response is slow due to GAB EAGAIN errors
1664952	Refreshing private region structures degrades performance during "vxdisk listtag" on a setup of more than 400 disks.
1662744	VVR:Hang in tli_send() for duration equivalent to "tcp_ip_abort_interval" time period (> 8 minutes)
1650663	Segmentation fault (core dumped) by vxsnap admir
1634547	I/O stops after manually changing secondary paths to the primary.
1603445	vxconfigd dumped core during "vxddladm -c assign names" after vxdiskunsetup on UDN name.
1591146	mirrored volume grow doesn't works well if "mirror=enclosure" option is used, it leads to data corruption issue.

Table 1-2 Veritas Volume Manager 5.0 MP3 RP4 fixed issues (*continued*)

Fixed issues	Description
1589715	vxconfigd dumps core, after vxdmpadm getportids ctrl=ctrl_name on a disabled ctrl
1558384	VxVM: checkin the fmrshowmap utility
1541662	System panicked in DRL code when running flashsnap
1529858	Site detach due to error falsely updates ssb on available disks on that site
1482555	AIX nightlytc: /ddl/dmpapm.tc#3 is getting hung at vxdmpadm -u cfgapm
1481493	panic seen at vxio:vol_cvol_bplus_walk3+1bec while running FMR stress test
1471741	vxdmpadm getdmpnode all option is not supported
1471602	AIX Rootability Enhancements
1471487	Critical Minimum Queue and Round-robin Improvements
1461717	`vxsnap make' command result in vxconfigd and IO sleep too long time
1457182	slab allocator in DMP can loop indefinitely with out getting additional memory from system.
1393756	vxvm commands hung on master and slave after FC-site link disconnected
1237675	vxdiskadm option 16-5 not working due to vxconfigd dumping core

Table 1-3 describes fixed issues in the Veritas Volume Manager 5.0 MP3 RP3 release, which are included in this release.

Table 1-3 Veritas Volume Manager 5.0 MP3 RP3 fixed issues

Fixed issues	Description
990338	FMR Refreshing a snapshot should keep the same name for the snap object.
963951	INSTSNAPTMP marked dco log not getting deleted during vxrecover or volume restart
339187	CVM activation tag in vxprint -m output breaks vxprint.
1923764	Starting all 128 way stripe volumes of stripe-mirror configuration fails.

Table 1-3 Veritas Volume Manager 5.0 MP3 RP3 fixed issues (*continued*)

Fixed issues	Description
1920288	64 bit VxMS plugins of VxVM are missing in the path /opt/VRTSvxms/lib/map/aix64.
1873491	3PAR ASL: MAXRCLM should be set to 2MB instead of 256MB.
1873390	Update HDS USP ASL with extended attribute information and TP/RCLM.
1871038	ASL Request Hitachi AMS 2000 with TP support.
1864547	VM commands getting hung on master node with 32-node cluster.
1863155	Unable to boot system restored with NIM mksysb image.
1850166	vxvm vxdisk error v-5-1-8643 device resize failed:
1850124	CVM TC errors for admin/fmr2 snapback.tc.
1842144	Update to 5.1 HP XP 24K ASL with extended attribute information and TP/RCLM.
1837189	DMP: System panic when perform excludearray operation with powerpath.
1835569	Incorrect dropping of messages when the messages arrive out of order during kernel-level join leading to hang/system crash.
1835139	CERT : pnate test hang I/O > 200 seconds during the filer giveback.
1834848	TP: reclamation causes data corruption
1826088	After pulling out FC cables of local site array, plex became DETACHED/ACTIVE.
1825637	VM cannot recognize over 2TB LUNs.
1824993	da_is_any_same_disk skipped disk, blank udid_asl "is same disk same as".
1822681	memory leak in vxio/voldr1_cleansio_start
1820179	vxdctl debug 0 core dumps after vxconfig -X #n -x logfile=file_name
1819777	Panic in voldiosio_start(as race window exists while handling duplicate DA records.
1798883	TP/RCLM and other extended attributes support in EMC CLARiiON ASL.
1797540	VxVM: vxdisk resize intermittently causes vxconfigd to dump core.

Table 1-3 Veritas Volume Manager 5.0 MP3 RP3 fixed issues (*continued*)

Fixed issues	Description
1792795	supportability feature/messages for plex state change, DCO map clearance, usage of fast re-sync by vxplex.
1787437	VXPLEX CPU USAGE IS very high for snapback operation.
1782111	TP : vxdisk -o thin list should say N/A instead of '0' for PHYSALLOC, if array don't support it.
1779257	VVR: Disable Secondary logging through a tunable.
1766452	VVR: VRAS: AIX: vradmind dumps core during collection of memory stats.
1762561	DMP: System panic when perform excludearray operation with powerpath.
1762534	vxdctl settz and vxconfigd core dump if TZ environment variable is not set.
1755830	kmsg: sender: the logic for resend of messages needs to be optimized.
1755810	kmsg: sender thread is woken up unnecessarily during flowcontrol.
1755788	for a broadcast message, sender thread may end up sending the same message multiple times (not resend).
1755735	recovery I/Os get broken down to voliomem_chunk_size.
1755707	vxtask list shows the same taskid for parent and child tasks.
1755689	During recovery, -o delayrecover option does not work as expected for value of 0.
1755628	kmsg layer: with heavy messaging in the cluster the receiver thread slows down processing.
1755519	kmsg layer: receiver side flowcontrol is not supported.
1745992	CVR:I/O hang in 4 node CVR cluster.
1744672	Primary slave hangs in volcvm_rvgrecovery_send_iocont(TC remote_write_reconfigure_2.tc.
1744224	FMR3: multiple vxplex attach cmds running in parallel on a volume lead to clearing DCO map and subsequently lead to corruption.
1742702	vxvmconvert fails, probably due to wrong disk capacity calculation.
1733811	System panic on voldco_isdirty code path while doing vxsnap make operation after upgrading from DCO version 10.

Table 1-3 Veritas Volume Manager 5.0 MP3 RP3 fixed issues (*continued*)

Fixed issues	Description
1732969	Data loss after SSB.
1732200	[DMP] [Usability] When NEW dmp_native_multipathing tunable is set to 'on' - unlabelled LUNs vanish from format until turned off.
1729558	multiple vxplex attach cmds running in parallel on a volume lead to clearing DCO map and subsequently lead to corruption in FMR2.
1728587	VVR: Replication started with a checkpoint remains inconsistent/cant_sync after SRL is drained if the replication is interrupted.
1728269	Incorrect cur_pri_path updation for A/PG arrays leading to dmp database inconsistency.
1726902	vxconfigd dumped core while trying to choose a path in dmp_dmpdevice_to_pathlist_ebn().
1722984	Memory leak in vold_dg_get_clone_disks(.
1715889	Unable to encapsulate an unmanaged EMC DMX PP LUN.
1713366	BCV: vxsnap restore fails for dg on bcv split devices
1678370	VM_VVR: RLINK disconnected and "vx" commands hung on Secondary while load in progress.
1677217	DMP does not autofailback to the Primary paths following LCC card restoration.
1676061	System panic'd after 2 out of 4 paths to disk were removed.
1673002	Need to remove thousands of empty /tmp/vx.* directories.
1668334	vxdisk -o list thin reports different results on master/slave nodes during dmp array failover tests.
1638494	VVR: vxnetd stop causing 100% CPU & vx commands hanging.
1638174	vxconfigd memory leak found.
1635815	/etc/vx/diag.d/kmsgdump panicked the system
1630572	Creating cdsdisk layout on GPT-labeled disks on Linux platform is defective.
1594928	Avoid unnecessary retries on error buffers when disk partition is nullified.
1594325	need to backout *unit_io and *pref_io changes after 5.0GA.

Table 1-3 Veritas Volume Manager 5.0 MP3 RP3 fixed issues (*continued*)

Fixed issues	Description
1589018	num_retries field is getting re-initialized to initial value leading to looping and delay in error handling time.
1545835	vxconfigd core dump during system boot after VxVM4.1RP4 applied.
1544478	cvr run hung in the tc remote_write_reconfigure_2.tc . primary slave panics on linux .tcrun hangs
1538053	CVM_MSG_REQ_GSLOCK repeatedly resent resulting in hang
1537027	SECURITY: ddl_change_naming_scheme(should set mode when creating .newnames.
1531330	Multiple node panic seen during master reconfig process in 4 node sfcs cluster.
1528160	An ioctl interrupted with EINTR causes frequent vxconfigd exit()'s on 4.1MP4RP3
1517719	vxconfigd core in ddb_creat_record via commit via dg_trans_commit via dasup_check after array disable
1508462	vxconfigd hung after cluster nodes split simulation - VxVM 5.0 MP3 RP1
1485075	vmtest/tc/scripts/admin/voldg/cds/set.tc hits DMP ted assert dmp_select_path:2a.
1479735	CVR: I/O hang on slave if master (logowner crashes with DCM active.
1471784	[5.0MP3RP1 x64] vm can not create stripe-mirror/mirror- stripe/mirror volume with maxsize.
1471263	machine has panicked when added the disk from dg as a foreign device using "vxdmpadm addforeign".
1470251	volslabd utilized 100% cpu time.
1468647	vxdmpdebug fails to find ugettxt
1459000	Fail over cmd on a bad LUN can cause an infinite loop in dmpCLARiiON_issue_failover.
1457182	slab allocator in DMP can loop indefinitely with out getting additional memory from system.
1442139	CVM join hang in VOLDJOIN, when all 4 nodes rebooted.

Table 1-3 Veritas Volume Manager 5.0 MP3 RP3 fixed issues (*continued*)

Fixed issues	Description
1060336	vxresize should not roll back if fsadm failed but disabled vxfs.

Table 1-4 describes fixed issues in the Veritas Volume Manager 5.0 MP3 RP2 release, which are included in this release.

Table 1-4 Veritas Volume Manager 5.0 MP3 RP2 fixed issues

Fixed issues	Description
850816	You can now delete snap objects from a mounted volume.
1714619	Fixed the cause of a panic due to DMP calling iodone() twice.
1711339	VVR tunables can now be modified using the vxtune command.
1677416	Fixed some CVM join and takeover issues that occurred in shared Active/Passive storage configurations that were due CVM sending messages that were more than 64k in size.
1675221	Fixed an issue in which the vxdmadm setattr enclosure command created identical disk attribute names.
1663338	Fixed an issue in which I/O load balancing to LUNs between two VIO servers not working.
1661938	Enhanced the vxconfigd daemon to take advantage of the AIX very large address space model.
1598706	Fixed the cause of a system crash that occurred while mirroring the rootdisk.
1598145	Fixed the cause of a crash in the vxdmproot install command during startup on an AIX 5.3 SAN boot disk.
1597868	Fixed an issue in which, on a secondary node, rlink paused and generated the "Incorrect magic number or unexpected upid" error message, and the secondary_log_err flag got set.
1596811	Improved performance when re-enabling individual paths.
1594325	Backed out *unit_io and *pref_io changes that were made in the 5.0 release.
1590314	The vxdmadm getsubpaths dmpnodename command now validates the dmpnodename value before getting the subpath information.

Table 1-4 Veritas Volume Manager 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1589881	Fixed an issue in which the dump device was changed to none (dumps disabled) after encapsulating a boot disk.
1589172	Fixed an issue in which the vxdisksetup and vxdiskunsetup commands sometimes failed for EFI disks.
1589022	Fixed the cause of an infinite loop in the DMP error handling code path with a CLARIION array, which led to an I/O hang.
1586879	Improved performance of the vxdisk online command when used on large configurations.
1544051	Fixed an issue in which the incorrect bit was being checked for an EMC Symmetrix thin device.
1543803	Fixed an issue in which the vxconvert command failed to convert a volume group.
1534379	Fixed an issue in which the vxdg split command failed with the following error: Internal configuration daemon error
1534038	Fixed an issue in which DMP stats sometimes used invalid I/O stats entries, which led to a panic on the host.
1532363	Fixed an issue in which the vxdisk updateudid command is corrupted the UDID of a disk, which made diskgroup imports fail.
1529157	The vxdisk -e list command output now includes a wider range of values for the TYPE and STATUS columns.
1528368	Fixed the cause of an I/O hang during the data change map transition after performing vxresize operations on the primary node.
1527247	Fixed an issue in which the vxstat command showed twice the I/O activity on a mirror volume compared to the source volume.
1525819	Fixed an issue in which the vxconfigbackup command failed to work on a diskgroup that had 2 TB LUNs.
1525121	Fixed an issue in which EFI disks were in an error state after installing the Storage Foundation 5.0 MP3 RP1 patches.
1515581	Fixed an issue in which recreating a shared diskgroup put CVMVolDg in an empty KSTATE and offlined clustered file systems.

Table 1-4 Veritas Volume Manager 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1512805	Fixed an issue in which two disks from different enclosures with the same array volume ID were treated as duplicate array volume IDs.
1512352	Fixed an issue in which the vxconfigstore command failed with the following error: VxVM vxconfigstore ERROR V-5-2-3706 Diskgroup configuration
1508462	Fixed the cause of a vxconfigd hang that occurred due to a split brain condition on a cluster.
1507935	Fixed an issue in which the vxconfigd daemon dumped core when the settag keyword was used to specify a long site name in a campus cluster.
1507291	Fixed an issue in which setting the dmp_monitor_fabric value to ON triggered unexpected offlining of paths on a DMX4 array.
1503168	Fixed an issue in which the diskgroup for disks without a private region (nopriv disks) could not be imported.
1502842	Fixed an issue in which the dmpolicy.info file did not get updated after upgrading the packages from Storage Foundation (SF) 5.0 MP3 RP1 to SF 5.1.
1501165	Changed the V-5-1-2140 message from an error to a warning.
1500389	The vxrootadm command now automatically enables the use-nvramrc? variable.
1488084	Fixed an issue in which the vxdmpadm iostat command reported different amounts of read/write blocks than the vxstat, iostat, and sar -d commands.
1485379	Fixed an issue in which the vxtask -l list command displayed incorrect progress of the vxsnap addmir command, which was used to link a snapshot volume to the source volume.
1484919	Fixed an issue in which a system that was upgraded to the 5.0 MP3 release could not be booted.
1483643	Fixed an issue in which a raid 5 volume would not start on 3PAR Thin Provisioning LUNs.
1483201	Fixed an issue in which the Device Discovery Layer (DDL) sometimes set the unique disk identifier (UDID) value to INVALID. Multiple disks set to INVALID resulted in the following error: VxVM vxio V-5-0-1056 new disk disk_id has a non-unique UDID

Table 1-4 Veritas Volume Manager 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1483164	Fixed an issue in which disks with the NOLABEL state were usable via the CLI.
1482687	Fixed the cause of an I/O hang during and after CVM reconfiguration following a reboot of the master node.
1480315	Fixed an issue in which VxVM performed a full re-sync of a volume that was created in the background when the volume's diskgroup was imported.
1479735	Fixed the cause of an I/O hang on a slave if the master (logowner) crashed with a data change map active.
1479729	Fixed the cause of an I/O hang on the primary node after a secondary node crashed.
1477143	The cluster volume manager failback protocol is now triggered when cur_pri is null and at least one DMP node of the same LUN group is DMPNODE_SHARED.
1475707	Added an error message for attempting to import unwritable disks.
1473638	Fixed the cause of a failover in the IOCTL context for coordinator disks.
1472736	Fixed the cause of a system panic in the vxdmp module that was due to a NULL pointer dereference.
1471813	The vxdmpadm start restore command now indicates that the restore daemon cannot be started when the dmp_enable_restore tunable is set to off.
1471763	Fixed the cause of the following error: build_devlink_list: readlink failed for /dev/vx/rdisk/ludg: Invalid argument
1471754	Fixed the cause of a vxconfigd daemon core dump that occurred in the ssb_increment_dg_abort() call.
1471658	Fixed the cause of a vxconfigd daemon core dump that occurred in the priv_get_all_udid_entry() call.
1471003	Fixed an issue in which the vxdg -s import oradg command failed with following error on a campus cluster: V-5-1-10978 Disk group oradg: import failed: required lock not held in transaction
1470548	Fixed an issue in which a campus cluster detached site could not be reattached.

Table 1-4 Veritas Volume Manager 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1469487	The I/O buffer start time is no longer modified as part of error analysis.
1463547	Fixed the cause of a vxconfigd core dump that occurred when dynamically reconfiguring a LUN.
1461717	Fixed an issue in which the vxsnap make command caused the vxconfigd daemon to hang.
1461338	VxVM now correctly handles CPU affinity switching.
1461314	DMP no longer uses the SCSI bypass on single path disks for path-suppressing TPD.
1460582	Fixed an issue in which the vxdmpadm -f disable enclosure command failed with a Segmentation Fault and dumped core.
1459831	Fixed an issue in which replication hung due to a deadlock on a secondary that had a TCP multiconnection and was managed by nmcom.
1458792	Fixed in issue in which the *unit_io and *pref_io tunables became set to 32 MB after upgrading from the Storage Foundation 5.0 MP1 release to the 5.0 MP3 release.
1457132	Fixed the cause of data corruption when running the vxdmpadm disable path and vxdmpadm disable ctlr commands.
1450348	Fixed a potential hang/panic that was due to a race condition between an RU thread and a volume read completing during DCM replay.
1449264	Fixed the cause of a panic in dmp_get_path_deferq that occurred while creating a volume.
1449001	Enabled the auto-compiler flag selection, which is used for enabling and disabling storage keys support.
1446216	Fixed an issue in which LUN mapping changes on node 0 caused node 1 and node 2 to panic.
1446208	Changed message V-5-1-2140 from an error message to an informational message.
1445884	Fixed an issue in which the vxdisk -e list command dumped core.
1425338	Fixed an issue in which connect rlinks failed to be connected, followed by vxconfigd hanging on a secondary node.

Table 1-4 Veritas Volume Manager 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1421353	Fixed an issue in which I/O got stuck in the <code>drl_logbusy</code> queue due to corruption of the age node LRU list.
1416080	Fixed the cause of a system panic in the <code>vol_change_disk()</code> routine that was due to NULL pointer dereference.
1414469	Fixed an issue in which the <code>vxddladm listsupport all</code> did not display up-to-date information.
1414336	Fixed an issue in which some disk devices did not appear in the <code>vxdisk list</code> command output.
1411438	The dynamic multi-pathing tunable <code>dmp_enable_restore</code> value is now persistent when it is set to off.
1408367	Fixed the cause of a system panic when <code>mutex_panic()</code> was called from <code>vol_rwsleep_wrlock()</code> .
1393764	Fixed an issue in which the <code>vxconfigd</code> daemon hung on a node that tried to become the master on site 2 when the Fibre Channel and heartbeat link was disabled at same time.
1388883	Fixed an issue in which rebooting a controller caused the diskgroups to be disabled.
1380386	The appropriate number of I/O threads are now created for systems with more than 8 CPUs.
1374603	Fixed a cause of data corruption in the <code>dmp_bypass_iodone()</code> call.
1370927	Fixed an issue in which the VTOC of disks in a cluster became corrupted.
1321298	Fixed the cause of a <code>vxconfigd</code> daemon core dump that occurred after reconnecting the FC site link and heartbeat link.
1321272	Fixed the an issue in which some VxVM commands hung after disconnecting, then reconnecting to the FC site link.
1314961	Fixed an issue in which if a disk tag name contained long DBCS characters, the configuration daemon was not accessible.
1287975	The <code>vxclustadm</code> command has a segmentation fault when the <code>main.cf</code> file contains lines that are greater than 512 characters.
1286298	Fixed an issue in which proper locks were not taken in all necessary places while modifying <code>last_sent_seqno</code> .

Table 1-4 Veritas Volume Manager 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1224659	Fixed an issue in which the vxconfigbackup -p script sometimes created a zero-length .binconfig file.
1222125	Fixed an issue in which after converting a big volume group, you could not mount a file system on the volume.
1195591	Fixed the cause of a panic when a cluster had an empty RVG.
1184280	Added additional debug messages around the VE_BADPROTOV error message to improve debugging.
1108839	Turning on the dmp_cache_open tunable no longer slows down the vxconfigd daemon when run with 2048 dual path LUNs.
1097258	The vxconfigd daemon no longer hangs when an array is disconnected.

[Table 1-5](#) describes fixed issues in the Veritas Volume Manager 5.0 MP3 RP1 release, which are included in this release.

Table 1-5 Veritas Volume Manager 5.0 MP3 RP1 fixed issues

Fixed issues	Description
425273	Fixed an issue with VVR RU thread not starting nio after it is created from than waiting for all replicas to have NIO's created.
1710030	Fixed the cause of a panic the occurred with LUNs larger than 1 TB.
1512521	StorageKeys: Adding compilation flag for VxVM to make function pointer references to legacy modules work in AIX 6.1 key aware code.
1455062	Fixed an issue with the master node crashing if a node leaves before responding to the MV- serialization protocol.
1449266	Fixed a panic in dmp_get_path_deferq while creating a volume.
1443752	Fixed an issue in a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load.
1443706	Fixed an issue in FMR3, I/Os initiating DCO updates for clearing DRL async clear region may not wait for its completion.
1442369	Fixed a bug in vxconfigbackupd script leading to 0 byte binconfig file being created.

Table 1-5 Veritas Volume Manager 5.0 MP3 RP1 fixed issues (*continued*)

Fixed issues	Description
1441131	Fixed an issue with VxFS Corruption Detected when DCM log plex are attached with mirrored volume and VVR is not configured.
1441072	Fixed an issue with siteread policy is not honoured.
1441020	Fixed a secondary panic due to double free of message with TCP protocol and 16 connection.
1440837	Fixed a panic due to memory allocation failure in interrupt context.
1436917	Fixed an issue with after installing VM patch on AIX, install-db would be created if the vxio in Defined state.
1435471	Fixed an issue with the cluster nodes panicking in voldco_or_pvmbuf_to_pvmbuf code after installing 5.0 MP3.
1431279	Fixed an issue with vxconfigd core dumps.
1428106	Fixed a system panic in vxio:voldr1_trans_copy.
1427284	Fixed an issue with vxdmpadm dumped core when executing vxdmpadm list dmpnode command.
1427267	Fixed a CVR panic in VOLSIOQ_MORE due to corrupted volsioq_start queue.
1425919	Fixed an issue with vxesd looping using 100% of one CPU.
1425434	Fixed an issue with CVR fails to connect rlinks followed by vxconfigd hangs on secondary.
1414441	The vxsnap manual page includes mirror=enclosure parameter to avoid being mirrored on the same enclosure.
1414381	Fixed an issue with VVR I/O hanging due to the wrong generation number assignment after recovery.
1411636	Fixed a secondary log error causing rlink disconnect after IBC unfreeze.
1403123	Fixed an issue with vxconfigd sleeping and no vx commands were responding.
1397879	Enhanced the vxresize manual page to run from non-CVM master.
1397712	Fixed an issue with the vxsnap restore manual page is unable to properly freeze or thaw filesystems in a CVM environment.
1396427	Enhanced DMP to handle failing IO when it is not able to interpret sense data.

Table 1-5 Veritas Volume Manager 5.0 MP3 RP1 fixed issues (*continued*)

Fixed issues	Description
1389584	Fixed a system panic in vol_putdisk() code.
1389512	Able to force import diskgroup version 80 in VxVM 5.0.
1387033	Fixed a system panic in bcopy() due to null passed in from volioctl_copyin()
1385922	Fixed a system panic due to memory allocation.
1376656	Fixed an issue with vxcached never deletes old snaps when cache hits HWM.
1372340	Fixed an issue with vxplex core dumps during vxassist addlog due to DRL log length being less than 33 blocks.
1368752	Fixed an issue when there are no mirrors to read, VOL_READ_MIRRORS ioctl returns -1 instead of 1.
1364335	Created a command to verify whether the DMP is controlling the rootdisk.
1364332	Fixed an issue with the VM commands not working on DMP Enabled boot path.
1364324	Fixed an issue with VSCSI: A/P LBI/O policy not working with enabled DMP support on boot devices.
1364320	Fixed a issue with vxdmproot install causing machine to hang upon reboot.
1266730	Fixed the vxtask command to display the resync progress subtask for shared volumes with DRL.
1230360	Fixed a system panic in vol_klog_start() due to accessing freed mv read_sio.
1192105	Fixed the vxdg -n [newdg] deport [origdg] command causing a memory leak.
1182475	Fixed the vxdg split failing if the CVM master changes.

Veritas File System fixed issues

Table 1-6 describes fixed issues in Veritas File System 5.0 MP3 RP5, which are included in this release.

Table 1-6 Veritas File System 5.0 MP3 RP5 fixed issues

Incident	Description
2329889 (2316094)	vxfstat's "vxi_bcache_maxkbyte" counter shows maximum memory available for buffer allocation. Maximum memory available for Buffer allocation depends on total memory available for Buffer cache i.e. "vx_bc_bufhwm". vxfstat can incorrectly report "vxi_bcache_maxkbyte" greater than "vx_bc_bufhwm" after re-initialization of buffer cache globals. reinitialization can happen in case of dynamic reconfig operations.
2292376 (1392781)	GLM hang due to an EXCLUSIVE waiting lock request to be starved by giving priority to SHARED lock requests.
2292359 (2253938)	EAU delegation timeouts
2292355 (2283893)	Add functionality of free space defragmentation through fsadm.
2292333 (2272072)	GAB panics the box because VCS engine "had" did not respond. The lobolt wraps around.
2270699 (2209182)	Three file systems hung during during netbackup processes creating/removing checkpoints while file systems were frozen
2253494 (1296491)	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2244377 (2126233)	Real time pri threads looping in vx_ireuse() causing system hang during HP's hazard testing
2239263 (2198173)	node panicked while umounting a filesystem.
2230175 (2226762)	Large number of transactions (overflow) causes a lot of buffer flushes
2220460 (2074806)	a dmapi program using dm_punch_hole may result in corrupted data
219691 (2180476)	System running VxFS panic in vx_iupdat_clustblks()

Table 1-6 Veritas File System 5.0 MP3 RP5 fixed issues *(continued)*

Incident	Description
2196905 (2184528)	fsck fails to repair corrupt directory blocks having duplicate directory entries.
2194623 (2178147)	Link of IFSOC file does not call vx_dotdot_op resulting in a corrupted inode
2151075 (1709113)	System got panic unexpectedly.
2132244 (2106154)	Lots of xmalloc requests can be seen from vx_sched
2118517 (2081499)	vxfscd may loop uninterruptible on system using default values for max_thread_proc:3000 and nkthread :8416.
2114130 (2061177)	fsadm -de' command erroring with 'bad file number' on filesystem(s)
2098378 (2098371)	VxFS Performance issue based on IO size & mount option 'nodatainlog'.
2092078 (2029556)	Two panics in mutex_exit: not owner
2090270 (2074281)	fcladm dump dumps core if no savefile is specified
2081439 (2080469)	Setting soft limit more than hard limit by vxquota doesn't report any error. Same with setting quota more than 1TB.
2080412 (2079457)	storage quota fs userquota user hardlimit numspace XX FS does not work, customer can make files after reaching hardlimit.
2066174 (2060219)	Panic is caused due to a race between inode memory de-initialization and inactive list processing code
2036213 (2028811)	ncheck dumps core in printname() when lot of io is going on.

Table 1-6 Veritas File System 5.0 MP3 RP5 fixed issues (*continued*)

Incident	Description
2026620 (1978884)	Customer uses heavy workload to monitor for abuses of application API by users These abuses were missed twice due to high average load at time of worklist threads
2026568 (2009472)	System hang as threads hang during directory creation using transaction
2016062 (1999493)	Corruption found when using DirectIO write with buffered read
1995398 (1985626)	Multiple panics in VxFS when doing large numbers of amounts in parallel
1665461 (1634788)	fsadm core dumps intermittently

[Table 1-7](#) describes fixed issues in Veritas File System 5.0 MP3 RP4, which are included in this release.

Table 1-7 Veritas File System 5.0 MP3 RP4 fixed issues

Incident	Description
2036662	Fixed an internal assert f:xted_fslst_lock1:2.
2034181	Added Power PC 7 support for storage keys.
2013519	Enhance mkdstfs for explicitly selecting the purpose of volumes.
2012501	Fixed an issue in thin reclam algorithm to ensure at least 95 % reclamation.
2007755	Optimized full fsck for file system with many ilist holes.
1991584	Fixed a panic in vx_timeout_free.
1983199	Fixed fsadm(1M) issue returning EFAULT while doing reclamation.
1959373	Fixed a resize issue with corrupt IFDEV.
1954691	Fixed a panic due to null pointer dereference during reverse name lookup operation.
1917200	Fixed a panic due to a bad mutex in a CFS environment.
1913794	Fixed a panic during a resize operation.

Table 1-7 Veritas File System 5.0 MP3 RP4 fixed issues *(continued)*

Incident	Description
1665434	Fixed an issue in vxupgrade which earlier used to fail with ENFILE.

Table 1-8 describes fixed issues in Veritas File System 5.0 MP3 RP3, which are included in this release.

Table 1-8 Veritas File System 5.0 MP3 RP3 fixed issues

Incident	Description
1923881	Thin provisioning: fixed vxdisk reclaim operation which was not freeing space properly.
1923883	Thin provisioning: fixed a problem of space not getting reclaimed properly because of alignment issues.
1716183	Thin provisioning: storage is not reclaimed on mirrored disk after data removed in volume.
477923	Performance improvement of file relocation by doing asynchronous I/O.
1865182	Fixed a hang during transaction commit during large directory hash setup by returning error instead of keep trying forever for resources.
1865175	Fixed a file system shrink operation failure during transaction commit due to unavailability of resources by committing shorter transactions sooner.
1748740	Fixed post upgrade problems in a multi TB file system by fixing a wrong 32 bit cast.
1826727	Fixed a file system disabled problem by rearranging rename code.
1591095	Fix to make mount -o mntlock and unlock work properly.
1441713	Fixed high resource contention problem during multiple checkpoint removal by throttling the clone remove operation.
1799963	Fixed a Data Storage Interrupt by adding protection gates to the fdd/odm iodone entry points.
1763947	Fixed a fsck usage error by fixing options processing in fsck.
1672319	Fixed bad inode problem resulting in a disabled filesystem by refreshing properly the inode data during a trunc operation.
1763956	Fixed the problem of fclear beyond eof and not fs block aligned failing by updating the size between the reorg inode and main inode properly.

Table 1-8 Veritas File System 5.0 MP3 RP3 fixed issues (*continued*)

Incident	Description
1537737	Fixed panic in vx_write during resize by initializing a local pointer.
1665472	Fixed problems in vxresize by now relocating structural files too.
1635131	Fixed space reclaim problems by reverting back max_seqio_extent_size from 104857 to 2048.
1415997	Fixed the missing Oracle archive log problem by making the link, remove etc operations flush the transactions synchronously to disk even if mount option is delaylog.
1443687	Fixed problems during vxfs upgrade by improving inactive processing, attribute-illist reorg inodes and inode reuse logic.
1596358	VxFS quota hard limits were not getting properly enforced on cluster file systems.

[Table 1-9](#) describes fixed issues in Veritas File System 5.0 MP3 RP2, which are included in this release.

Table 1-9 Veritas File System 5.0 MP3 RP2 fixed issues

Incident	Description
847454	The VxFS cp command now has similar performance as the JFS2 cp command.
1298054	Fixed an issue in which ODM randomly displayed the following error: ORA-17500: ODM err:ODM ERROR V-41-4-2-42-11 Resource temporarily unavailable
1299313	Fixed the cause of a panic in vx_unlockmap() due to a null ml_tranp pointer.
1317905	Fixed the cause of a core dump when running the quotacheck command with more than 30 quota-enabled file systems in the /etc/fstab file.
1363639	Fixed the cause of an exception in the vx_vnmap_min() call that was due to the vx_fsext_info() call of a force umounted file system having a null fse_fsext_info pointer.
1387074	Fixed the cause of the fsclustadm cfsdeinit command failing with the "device busy" error.
1416745	Fixed an issue in which the dm_get_allocinfo() call failed with the EIO error for EXT4 inodes with indirects.

Table 1-9 Veritas File System 5.0 MP3 RP2 fixed issues (*continued*)

Incident	Description
1432323	Fixed a loop issue in which the vx_do_putpage() call due to the page going beyond i_wsize.
1443034	Fixed the cause of hangs with the vxresize and fsadm -b commands on a VxFS file system.
1445008	Fixed the cause of a system panic that occurred when auditing files on a VxFS file system.
1445131	Fixed a bug in the vx_ifree_scan_list() call in vx_iflush.c.
1454783	The vx_multi_bufinval () call now releases the CPU for a local mount with large extents.
1459329	New VxFS tunables and new vxfsstat counters now increase the number of VMM buffers per PDT, which improves performance.
1500197	Scanning of pages in v_write during putpage no longer causes GAB/LLT failure.
1517415	Fixed the cause of a core dump when running the ncheck command.
1526568	The vx_tflush_map() call no longer disables the file system when a map is marked bad without having an actual I/O error.
1542572	Fixed a rounding-up of length issue in the vx_fsync_range() call.
1591301	Fixed an issue in which a smap was marked bad due to punching a hole in the allocated AU.
1633530	Fixed an issue in which calling vx_setext() while having the VX_GROWFILE flag set sometimes failed to grow the file or exposed uninitialized data.
1667628	Implemented a new tunable, flush_chunk_size, which is used by the vx_putpage_dirty() and vx_mm_invalidatep() calls.
1701842	Fixed a cause of Concurrent I/O returning the ENOTSUP error to DB2, which caused DB2 to crash.

Table 1-10 describes fixed issues in Veritas File System 5.0 MP3 RP1, which are included in this release.

Table 1-10 Veritas File System 5.0 MP3 RP1 fixed issues

Incident	Description
1366772	Fixed a performance issue.
1398904	Fixed an issue with VxFS filesystems temporarily hang in vx_delay().
1400046	Fixed an issue with the fsapadm enforceckpt mount_point command that resulted in a core dump.
1412160	Fixed a core dump caused by a VxFS function call while setting DST attributes.
1412465	Fixed a vxresize command failure to resize a volume, but the command could resize the file system.
1414178	Fixed an issue with VxFS using too much CPU while looking for odd-sized extents (vxi_alloc_fail).
1415188	Fixed a full fsck core dump that was due to running out of swap space and a malloc failure.
1423867	Fixed an issue in which the vx_convnodata_files() call could take more than 10 minutes to complete.
1487928	Fixed the build script to compile packages to enable the storage keys feature.
1517337	Fixed a panic issue with lockcount of an FS thread upon thread_terminate().

Storage Foundation Cluster File System fixed issues

[Table 1-11](#) describes fixed issues in Storage Foundation Cluster File System 5.0 MP3 RP5, which are included in this release.

Table 1-11 Storage Foundation Cluster File System 5.0 MP3 RP5 fixed issues

Fixed issues	Description
2329889 (2316094)	vxfsstat's "vxi_bcache_maxkbyte" counter shows maximum memory available for buffer allocation. Maximum memory available for Buffer allocation depends on total memory available for Buffer cache i.e. "vx_bc_bufhwm". vxfsstat can incorrectly report "vxi_bcache_maxkbyte" greater than "vx_bc_bufhwm" after re-initialization of buffer cache globals. reinitialization can happen in case of dynamic reconfig operations.
2292376 (1392781)	GLM hang due to an EXCLUSIVE waiting lock request to be starved by giving priority to SHARED lock requests.

Table 1-11 Storage Foundation Cluster File System 5.0 MP3 RP5 fixed issues
(continued)

Fixed issues	Description
2292359 (2253938)	EAU delegation timeouts
2292355 (2283893)	Add functionality of free space defragmentation through fsadm.
2292333 (2272072)	GAB panics the box because VCS engine "had" did not respond. The lobolt wraps around.
2277472 (2206065)	There is a mismanagement of the holds/releases of incore LDH-attribute-inodes(Large Directory Hash). Due to which cache of LDH-attribute-inodes is being expunged blindly in the function vx_inactive(), which is resulting in panic.
2270699 (2209182)	Three file systems hung during during netbackup processes creating/removing checkpoints while file systems were frozen
2253494 (1296491)	Panic occurs while doing nested mount when the base cluster mounted base fs gets force unmounted
2244377 (2126233)	Real time pri threads looping in vx_ireuse() causing system hang during HP's hazard testing
2239263 (2198173)	node panicked while umounting a filesystem.
2230294 (2220300)	'vx_sched' is hogging CPU resources.
2230175 (2226762)	Large number of transactions (overflow) causes a lot of buffer flushes
2220460 (2074806)	a dmapi program using dm_punch_hole may result in corrupted data
2196910 (2180476)	System running VxFS panic in vx_iupdat_clustblks()

Table 1-11 Storage Foundation Cluster File System 5.0 MP3 RP5 fixed issues
(continued)

Fixed issues	Description
2196905 (2184528)	fsck fails to repair corrupt directory blocks having duplicate directory entries.
2196883 (2184114)	Three node VCS/CFS cluster where the CVM/CFMount monitoring is getting timeout very frequently
2194637 (2161379)	One node of a 4-node CFS cluster repeatedly hangs due to deadlock between ILOCK and inode owner
2194623 (2178147)	Link of IFSOC file does not call vx_dotdot_op resulting in a corrupted inode
2151075 (1709113)	System got panic unexpectedly.
2132244 (2106154)	Lots of xmalloc requests can be seen from vx_sched
2118517 (2081499)	vxfsckd may loop uninterruptible on system using default values for max_thread_proc:3000 and nkthread :8416.
2114172 (2091103)	CFS hangs in cluster
2114130 (2061177)	fsadm -de' command erroring with 'bad file number' on filesystem(s)
2098378 (2098371)	VxFS Performance issue based on IO size & mount option 'nodatainlog'.
2092086 (2030289)	FS corruption issue
2092078 (2029556)	Two panics in mutex_exit: not owner
2090270 (2074281)	fcladm dump dumps core if no savefile is specified

Table 1-11 Storage Foundation Cluster File System 5.0 MP3 RP5 fixed issues
(continued)

Fixed issues	Description
2081439 (2080469)	Setting soft limit more than hard limit by vxedquota doesn't report any error. Same with setting quota more than 1TB.
2080412 (2079457)	storage quota fs userquota user hardlimit numspace XX FS does not work, customer can make files after reaching hardlimit.
2069082 (2069059)	CFS hang when setting LIBPATH to a CFS directory.
2066174 (2060219)	Panic is caused due to a race between inode memory de-initialization and inactive list processing code
2036213 (2028811)	ncheck dumps core in printname() when lot of io is going on.
2026620 (1978884)	Customer uses heavy workload to monitor for abuses of application API by users. These abuses were missed twice due to high average load at time of worklist threads.
2026568 (2009472)	System hang as threads hang during directory creation using transaction
2016062 (1999493)	Corruption found when using DirectIO write with buffered read
1995398 (1985626)	Multiple panics in VxFS when doing large numbers of umounts in parallel.
1665461 (1634788)	fsadm core dumps intermittently

Table 1-12 describes fixed issues in Storage Foundation Cluster File System 5.0 MP3 RP4, which are included in this release.

Table 1-12 Storage Foundation Cluster File System 5.0 MP3 RP4 fixed issues

Fixed issues	Description
2042468	Fixed an issue with CFSmount which earlier used to fail with the CFSMount:mwlogsnp-cfs:online:Mount Error : UX:vxfs mount: ERROR: V-3-21272: mount option(s) incompatible with file system.
2012504	Optimized the file attributes retrieval in a CFS environment.
2001745	Fixed an alignment issue during EMAP processing.
1959383	Fixed a performance issue during file removal.
1891145	Fixed an issue with updatation of the ias_list in CFS environment.
1702971	Fixed a stale cache data issue with CIO in a CFS environment.

[Table 1-13](#) describes fixed issues in Storage Foundation Cluster File System 5.0 MP3 RP3, which are included in this release.

Table 1-13 Storage Foundation Cluster File System 5.0 MP3 RP3 fixed issues

Fixed issues	Description
1880424	Fixed a hang during trunc by moving extent free operation to later time if the extent spans multiple AUs.
1841266	Fixed performance problem during write by replacing delay with sleep/wakeup mechanism.
1807557	Fix to make VX_FREEZE_ALL ioctl work with CFS file systems.
1799960	Fixed a CFS hang by changing lock mode of a lock from exclusive to shared thus allowing more threads to continue.
1793317	Fixed a problem of CFS hang after snap error by releasing the lock on snap after it has been marked bad.
1662270	Fixed a performance problem by making delegation processing of partial AUs more efficient.
1600552	LM & CFS-Conformance tests failed in fsync/fsync_new test in 'fclear' test cases, seeing data corruption.
1531028	VxFS quota hard limits were not getting properly enforced on cluster file systems.

Table 1-13 Storage Foundation Cluster File System 5.0 MP3 RP3 fixed issues
(continued)

Fixed issues	Description
1435654	Fixed a problem in CFSSMountAgent by making fsclustadm MT safe.

[Table 1-14](#) describes fixed issues in Storage Foundation Cluster File System 5.0 MP3 RP2, which are included in this release.

Table 1-14 Storage Foundation Cluster File System 5.0 MP3 RP2 fixed issues

Fixed issues	Description
1600241	Fixed the cause of a hang that occurred after another node in the cluster crashed.
1591313	Fixed the cause of a loop in the vx_dele_get_freespace() call due to a smap being marked bad.
1539892	Fixed an issue in which a clustered file system that was mounted on one node required fsck to be run.
1518713	The vxfsckd -n command now initializes the nthrs variable.
1286525	Fixed an issue in which Java threads hung on SFCFS functions.

[Table 1-15](#) describes fixed issues in Storage Foundation Cluster File System 5.0 MP3 RP1, which are included in this release.

Table 1-15 Storage Foundation Cluster File System 5.0 MP3 RP1 fixed issues

Fixed issues	Description
1487928	Fixed the build script to compile packages to enable the storage keys feature.

Storage Foundation for Oracle fixed issues

[Table 1-16](#) describes fixed issues in Storage Foundation for Oracle 5.0 MP3 RP3, which are included in this release.

Table 1-16 Storage Foundation for Oracle 5.0 MP3 RP3 fixed issues

Fixed issues	Description
--------------	-------------

[Table 1-17](#) describes fixed issues in Storage Foundation for Oracle 5.0 MP3 RP2, which are included in this release.

Table 1-17 Storage Foundation for Oracle 5.0 MP3 RP2 fixed issues

Fixed issues	Description
1651363	Fixed a security issue with the vxdbms server, in which an attacker could see the name and port of the server.
1533204	Fixed an issue in which the DBED GUI showed archive log mode as disabled when the archive log was actually enabled. Also, fixed an issue in which the number of file systems and the number of data files always showed as 0 (zero).
1530125	Fixed an issue in which the owner of the following directories was changed when installing VRTSdbms packages for the Storage Foundation for Oracle 5.0 or 5.0 MP3 releases: /etc /etc/default /etc/init.d /etc/rc2.d /opt
1526653	Fixed an issue in which the dbed_vmchecksnap script output an error if the dco object name was renamed from *_dco.
1511321	Fixed multiple issues with the dbed_checkconfig script. For example, the script can now distinguish if the control file is on a volume set and can identify if some of the Oracle files are not on a VxFS file system.
1508346	Added a date stamp to entries in the vxsnapadm_50.log file, which is used for trace vxsnapadm issues.
1481426	Fixed an issue in which the owner of the following directories was changed when installing patches or packages for the Storage Foundation for Oracle 5.0 or 5.0 MP3 releases: /etc /etc/default /etc/init.d /etc/rc2.d /opt

Table 1-17 Storage Foundation for Oracle 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1670804	Fixed an issue in which snapshot resyncs were performed serially and not in parallel as was true with the 5.0 MP3 release and earlier

[Table 1-18](#) describes fixed issues in Storage Foundation for Oracle 5.0 MP3 RP1, which are included in this release.

Table 1-18 Storage Foundation for Oracle 5.0 MP3 RP1 fixed issues

Fixed issues	Description
1435527	Improved boot time for DBEDAgent startup script.
1434688	Storage Foundation for Oracle is no longer creating world writable files under /tmp.
1433571	Sybase repository database server is no longer creating world writable files under /tmp.
1433244	Improved boot time for the DBED repository database server startup script.
1425261	Automatic truncation of the transaction log of the repository database. In addition incomplete recovery is automatically attempted in case the online transaction log was lost.
1425256	Support flashsnap CVM slave.

Storage Foundation for DB2 fixed issues

[Table 1-19](#) describes fixed issues in Storage Foundation for DB2 5.0 MP3 RP3, which are included in this release.

Table 1-19 Storage Foundation for DB2 5.0 MP3 RP3 fixed issues

Fixed issues	Description
1873755	Storage Foundation for DB2 no longer creates world writable log files in /var/vx/vxdba directory. The sfua_db_config command is modified to ask for DBA group information so we can set the correct group for various directories in /var/vx/vxdba.

Table 1-19 Storage Foundation for DB2 5.0 MP3 RP3 fixed issues (*continued*)

Fixed issues	Description
1854457	Fixed an issue with db2ed_clonedb Checkpoint clonedb fails for online, offline checkpoint on DB2 9.5 FixPak 2 and beyond.
1854456	Fixed an issue with db2ed_vmclonedb -o recoverdb fails for online snapshot mode on DB2 9.5 FixPak 2 and beyond.
1851299	Storage Foundation for DB2 no longer creates world writable directories under /var/vx/vxdba.

Table 1-20 describes fixed issues in Storage Foundation for DB2 5.0 MP3 RP2, which are included in this release.

Table 1-20 Storage Foundation for DB2 5.0 MP3 RP2 fixed issues

Fixed issues	Description
1651363	Fixed a security issue with the vxdbms server, in which an attacker could see the name and port of the server.
1508346	Added a date stamp to entries in the vxsnapadm_50.log file, which is used for trace vxsnapadm issues.

Table 1-21 describes fixed issues in Storage Foundation for DB2 5.0 MP3 RP1, which are included in this release.

Table 1-21 Storage Foundation for DB2 5.0 MP3 RP1 fixed issues

Fixed issues	Description
1435527	Improved boot time for DBEDAgent startup script.
1434688	Storage Foundation for DB2 is no longer creating world writable files under /tmp.
1433571	Sybase repository database server is no longer creating world writable files under /tmp.
1433244	Improved boot time for the DBED repository database server startup script.
1425261	Automatic truncation of the transaction log of the repository database. In addition incomplete recovery is automatically attempted in case the online transaction log was lost.

Veritas Cluster Server fixed issues

[Table 1-22](#) describes fixed issues in Veritas Cluster Server 5.0 MP3 RP5, which are included in this release.

Table 1-22 Veritas Cluster Server 5.0 MP3 RP5 fixed issues

Fixed issues	Description
2277517	Fixed the issue that cannot configure SCSI-3 fencing using RamSan DMP devices.
2296861	Fixed issue that VCS cannot bring the Sybase dataserver online. However, the Sybase dataserver can be started outside VCS control.
2278123	Fixed issue that MemCPUAllocator agent breaks if CPURequired attribute not declared in AIX 6.1 Virtualised environment
2278117	Fixed issue that Application agent clean script fails to work when using PidFiles due to bad use of array
2230371	Fixed the issue that <code>"/usr/sbin/_vxfsenwap -g fendg -a autoconfirm"</code> failed
2215888	Fixed issue that the monitor script of Db2udb do not handle the case when a parameter is undefined, which make an empty value being passed to next level.
2138472	Fixed issue that the agent dumps core when a Script based WAN heartbeat is added.
2133775	Fixed issue with LLT: increment the wrenable_clue if we can't dupmsg
2125454	Fixed issue that ASMInst is hard coded to perform health check monitoring which fails against Oracle 10.2.0.4
2109272	Fixed issue that unable to bring LVMVG online
2108826	Fixed issue that locator agent breaks if CPURequired attribute not declared in AIX 6.1 Virtualised environment
2106847	Fixed issue that GAB stuck with flow control due to race in handling LLT DISCONNECT and LLT CANPUT failure.
2104850	Fixed issue that while a failover service group is onlining on a node, Online operation is initiated on another node by an offline-local dependent service group
2098103	Fixed issue that If PreOnline and OnlineRetryLimit > 1, service group will not restart locally

Table 1-22 Veritas Cluster Server 5.0 MP3 RP5 fixed issues (*continued*)

Fixed issues	Description
2091444	Fixed issue that vxfcntlstdhw fails on unregistering keys in Japanese Locale
2078943	Fixed issue that DiskGroup agent forcing panic for no obvious reason
2077397	Fixed issue that VCS Dumps core under heavy load and the node goes for panic.
2077381	Improved heartbeating logic b/w engine and agfw during snapshotting and during steady cluster operation.
2077191	Fixed issue that nfs client I/O thrash got interrupted during cluster failover caused by split-brain.
2067297	Fixed issue that vxfcntlstdhw_common.sh: remove redundant KSH instances
2033683	Fixed issue that RemoteGroup resource doesn't currently support remote groups that are a parallel SG
1988109	Fixed issue that LVMVG Service Group switchover : Produces Invalid Type
1939135	Fixed issue that the hadsim core dumps when the systems are switching to running state for simulator.
2177206	Fixed the issue that split-brain occurs even when using VCS Steward
2318316	Fixed the issue that Oracle needs its database's \$Oracle_home/lib library to be first in LD_LIBRARY_PATH before /usr/lib
2117381	Fixed the issue that DB2 resourcec generating excessive logging in engine_A.log
2185545	Fixed the issue with IcmpAgent cpu consumption (e2005490) fix for 5.0MP3RP4
2194409	Fixed the issue that hacf dies dumping VCS configuration due to some attributes containing strings > 4Kb in length
2195571	Fixed the issue that RemoteGroup resource during network failure cannot be recovered without bouncing entire cluster.
2198773	Fixed the issue that hagr -switch of child group fails in 5.0MP3RP2 and later if 2 or more parent groups online on alternate
2204035	Fixed the issue that MonitorTimeStats incorrectly showing 303 secs Intermittently

[Table 1-23](#) describes fixed issues in Veritas Cluster Server 5.0 MP3 RP4, which are included in this release.

Table 1-23 Veritas Cluster Server 5.0 MP3 RP4 fixed issues

Fixed issues	Description
2033395	Notifier Agent is unable to get local IP address in linked-based IPMP.
2011417	LLT: Fix Memory leaks in ll_t_dmpi_setup and ll_t_closenode paths.
1994538	RemoteGroup resource detected offline even with patch for e1834858
1974589	LLT: Removing a link from LLT results in PANIC due to an invalid lower STREAMS queue.
1954778	Oracle ASM resource does not come to ONLINE state upon reboot in 11gR2 setup
1946365	Change in LLT delivery thread's priority limits
1915907	hares allows to create resources which has "." special character
1504121	Symantec SE - GCO failover does not work when user account has "!" in name.

[Table 1-24](#) describes fixed issues in Veritas Cluster Server 5.0 MP3 RP3, which are included in this release.

Table 1-24 Veritas Cluster Server 5.0 MP3 RP3 fixed issues

Fixed issues	Description
1924117	VRTSllt.rte 5.0.3.300 patch is unable to create correct link for drivers.
1906771	ASMAgent connecting as sysdba instead of sysasm for 11gR2. 11g and above the ASMInst and ASMDG agents use the role sysasm rather than sysdba in the offline, online, clean, and monitor entry points.
1898408	GAB: node panics when the master dies just before sending DLV for the RECONFIG msg
1898247	Netlsnr offline script does not kill listener process when ip is plumbed but the underlying MultiNICA resource is faulted.
1897072	MNICB resource failed to detect nic fault.
1891000	Switch parallel group to clus2 failed.

Table 1-24 Veritas Cluster Server 5.0 MP3 RP3 fixed issues (*continued*)

Fixed issues	Description
1884737	Fixed an issue in GAB's sequence recovery protocol where node may panic if it receives delayed response from the heavily loaded master node.
1880812	"hasys -display" shows the private link up even a node doesn't run llt
1860153	Send LLT heartbeats from various contexts so that even if the node is heavily loaded it doesn't miss to send heartbeats to other nodes.
1860148	After completion of `./installvcs -configure' process,'gabconfig -a' shows improper output for some time.
1860140	LLT should give error if an attempt is made to configure more than 8 links (LLT_MAX_LINK) under LLT.
1860123	Fixed an issue in vxfenswap command where it asks root password multiple times if the passwordless ssh/rsh is not configured.
1860115	"shutdown -r now" message to be removed or modified in vxfenclearpre command.
1860108	vxfenswap/vxfenclearpre prompt should be at the right place.
1860100	FENCING:Need of syntax check for /etc/vxfendg file.
1860090	Fixed an issue in vxfentsthdw command where it fails when an array has more than one paths to same I_T nexus.
1840045	Added a new check in the vxfentsthdw command to verify whether unregistering a Key from one node removes other node's key also.
1840037	vxfen symbol file for kdb not being built correctly.
1839349	kdb extension for GAB prints the same messages log (mlog) for all ports
1839321	LLT: We need to inform the senders immediately about the recv side flow control
1839318	LLT doesn't work over IPv4 UDP
1839299	LxRT5.1: VCS: GAB (fd 18) is continuously returning EAGAIN for vxconfigd port w
1838546	LLT: Cap the time for which we send out of context hbs
1836923	Application agent on Aix dump userinfo to engine log every monitor cycle when monitorProgram is configured with non root user.

Table 1-24 Veritas Cluster Server 5.0 MP3 RP3 fixed issues (*continued*)

Fixed issues	Description
1836633	hashadow core in restart_had /var/VRTSvc/lock/.hadargs parse resulted in attempt to deref null ptr
1836575	SMTP notification email should contain Entity name in subject line
1836512	`had' segv via notifier messages handler.
1834858	RemoteGroup faults when setup as monitoronly and local SG is taken offline
1834490	Application agent on Aix spits RC command not found message from monitor entry point
1829201	Application resource returns UNKNOWN on the failover node if owner's home directory not present
1807047	Issues found with SqlTest.pl script for Sybase agent
1780698	VCS Oracle agent not sending notification in case of an Oracle error defined in oraerror.dat
1875766	-parentprop option for hares was reported as Unknown option.
1839338	LLT: panic in gab_initllt() via gab_drv_init() via gab_config_drv() while installsfrac -configure
1744255	Agfw should not convert IntentionalOffline to Offline, (1) in first probe, (2) when probe is requested in Offline state
1739684	Case 320-192-581-- CCStor incorrect discovery as hasys output doesn't separate nodes by `#'

Table 1-25 describes fixed issues in Veritas Cluster Server 5.0 MP3 RP2, which are included in this release.

Table 1-25 Veritas Cluster Server 5.0 MP3 RP2 fixed issues

Fixed issues	Description
1513967/163872	[LLT/GAB] Fixed an issue in which one or more nodes would halt if LLT or GAB attempted to submit timer requests during a heavy I/O workload. This leads to an abend exception and a system panic.

Table 1-25 Veritas Cluster Server 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1709129	[GAB] Fixed an issue in which symbolic links that were created during the installation of the GAB package were not removed during the uninstallation of the GAB package.
1705098	[GAB] Fixed an issue in which the GAB module failed to load since existing configuration scripts used during patch installation did not remove pre-existing ODM entries. The fix involved adding a new kernel tunable (gab_timer_pri) to the PdAt class of the ODM database.
1703756	[VCS] Fixed an issue in which a warning message is displayed even when a parallel global group was brought online successfully. This happens because after a suitable target is determined, an internal variable is not incremented. This results in a re-visiting of the target selection algorithm, which causes error because the action is already initiated on the suitable target.
1678110	[GAB] Fixed an issue in which high priority processes running in the cluster may result in delayed response to the GAB timer function. the GAB timer function may not run as quickly as required if there are higher priority processes in the system. The priority for the gab timer is made into a tunable "gab_timer_pri". This tunable can have a value within a range.
1677412	[Agents] Fixed an issue so that when the SystemList of the service group is modified, you do not start all agents but only the required agents. The agent that was stopped by a user on a system gets restarted even if the group has no resource of that agent type, when the SystemList is modified to add that system. On SystemList modification to add new systems in SystemList, the engine starts all the agents without ensuring if the group has a resource of that type. Code changes so that only agents for which the group has resources are started whenever the SystemList is modified to add a new system.
1675815	[HAD] Fixed an issue so that the HostMonitor objects like VCShmg (Group), VCShm (Resource), and HostMonitor (Type) are not counted in each object's number.

Table 1-25 Veritas Cluster Server 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1672405	<p>[VCS] Fixed an issue in which a switch operation on a child service group with an OLH (Online Local Hard) and OLF (Online Local Firm) parent results in a switch of the OLH parent and the child group even though the OLF parent was online. In a situation, where two service groups depend on one child and one parent has an online local hard dependency (OLH) while the other parent has an online local firm dependency (OLF):</p> <p>The command: <code>hagrp -switch Hard_ParentSG -any</code> switches both the parents. The command: <code>hagrp -switch Hard_ParentSG -to sysB</code> switches only the hard parent group along with the child group. When the <code>hargp -switch</code> command is executed with any of the following options:</p> <ul style="list-style-type: none"> i) <code>hagrp -switch SG_parent -any</code> ii) <code>hagrp -switch SG_parent -to sys</code> <p>The parent group switches (while the child group is online) only in the case of a hard dependency. The switch does not happen in the case of soft or firm dependency. The switch operation succeeds for an OLH parent, if only the parent group is online. The child group has no other parents online. The OLH parent and child group can have other parents. However, the OLH child group is always a leaf node.</p>
1672335	<p>[Fencing] The <code>vxfcntlpre</code> command now outputs the progress of its execution.</p>
1670337	<p>[LLT] Added a mechanism to track the operating system timeouts that are registered by LLT with the operating system.</p>
1668609	<p>[Agents] Fixed an issue in which the Proxy agent is updated to allow the target resource to be probed before scheduling the first probe of the Proxy resource.</p>
1665036	<p>[Fencing] The <code>vxfen</code> startup process on retry no longer displays the "RFSM GAB err 16" error when the cluster is fencing a node out of the cluster.</p>
1640292	<p>[HAD] Fixed an issue to increase the resilience of HAD to extreme load conditions.</p>

Table 1-25 Veritas Cluster Server 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1638725	[LLT] Fixed an issue in which the LLT timer function may not run as quickly as required if there are higher priority processes in the system. LLT uses the heartbeat mechanism between nodes to ensure and identify that the other node is alive. Any node in VCS/SFRAC sends heartbeat packets to all the other nodes in the cluster after every 50 millisecond. This heartbeat is sent with the LLT timer thread. Under a heavy load condition, LLT timer thread may not be scheduled to send heartbeat. If the LLT thread is on the busy node, it is not able to send a heartbeat for 16 seconds. The other node considers the busy node failed and this results in panic whenever the load of the busy node goes down and it starts communicating with the other node of cluster. The LLT heartbeat code has been moved from an llt thread context to a timer interrupt context. This ensures that the heartbeat is sent as soon as timer returns after 50 milliseconds. Interrupt handler will run real time and this removes scheduling delays.
1638240	[Agents] Fixed an issue in which the Sybase agent is unable to bring the Sybase resource online if the RUN_servername file is moved to some other (non default) location. The non default location for the Sybase dataserver RUN_servername file is not supported by the Sybase agent. Hence, if you move the RUN_servername file to some other location, the agent is unable to bring the Sybase resource online. A new attribute named Run_ServerFile of type string was introduced for the Sybase and SybaseBk agents. The value of this attribute can be set to the absolute path of the RUN_servername file.
1635044	[Fencing] The vxfen scripts now determine the location of ssh and scp from the environment variables if ssh or scp are not available in the standard path.
1634924	[VCS] Fixed an issue in which the engine logs indicated CPU usage even after the HostMonitor resource is deleted.
1633973	[VCS] Fixed an issue in which the node does not test the Authority attribute before bringing the faulted service group online, leading to concurrency violations and the service group being taken offline on the disaster recovery site.
1631012	[LLT] Fixed an issue in a configuration with LLT over UDP4, in which when an ip address is not plumbed on a link and an IP address is specified in llttab file, lltconfig would configure that link successfully when it should not.
1603120	[VCS] Fixed an issue where NFSRestart triggers were called despite no configured NFSRestart resources, which was detrimental to performance. See "Mandatory configuration change for the NFS and NFSRestart resources".

Table 1-25 Veritas Cluster Server 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1600786	<p>[Fencing] Fixed an issue in which I/O errors occur in case of a network partition at any point when the keys on the coordinator disks are being refreshed using the vxfsnwap command. If the keys on coordinator disks are accidentally cleared, they can be refreshed using the vxfsnwap command. However if there is a network partition at a particular point in the operation, it could result in I/O errors. If the keys that are registered on the coordinator disks are lost, the cluster may panic when a split-brain occurs. Using the vxfsnwap script to replace the coordinator disks with the same disks will register the missing keys again without any risk of data corruption. However there is a possibility of seeing I/O errors because the algorithm registers the keys in the modify phase and if there is a network partition then the register(s) could override preempt(s) without synchronization. If the vxfsnwap utility is run on existing coordinator disks, then the registrations are done in the commit phase instead of the modify phase.</p>
1600484	<p>[VCS] Fixed an issue so that user names are checked and validated while verifying the configuration and modifying the UserNames attribute. A user with a special character in the userid is accepted if it is the second or later user in the UserNames attribute within the main.cf file. Only the first user name is checked for valid names. If the attribute UserNames has more than one user defined in the main.cf file or the command haclus -modify UserNames u1 p1 u2 p2 is run, then even invalid user names were accepted.</p>
1598940	<p>[Agents] Fixed an issue in which the NIC and MultiNICA resources did not go online because the broadcast ping failed due to changes in the AIX operating system. Note that these agents will break if used with SP3+. Modify agent to correctly parse netstat output for both pre and post AIX 6.1 TL2 SP3 levels.</p>
1590726	<p>[VCS] Fixed an issue in which VCS generated notifications about high CPU/SWAP usage when notifications were configured. The HostMonitor feature is enhanced to give control to the user for enabling or (fully / partially) disabling the feature through the cluster object attribute - HostMonLogLvl. VCS has the HostMonitor feature enabled by default through the VCSmg group with a HostMonitor type resource VCSm. If notification is configured in VCS, you see the notifications whenever the CPU/SWAP usage is beyond critical levels. A new attribute HostMonLogLvl is added. The values can be 'ALL', 'HMAgentLog' or 'DisableHMAgent', with `ALL` as default.</p>
1589851	<p>[GAB] Fixed the cause of a system panic that was due to depleted memory reserves.</p>

Table 1-25 Veritas Cluster Server 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1545229	<p>[Agents] Fixed an issue to allow control of entry point scheduling priorities and scheduling class using the new attributes EPPriority, EPClass, OnlinePriority, and OnlineClass.</p> <p>See "Attributes to control the scheduling of class and priority of agent entry points".</p>
1545222	<p>[Agents] Fixed an issue to provide the ability to pass the entry point timeout value as a parameter to agent entry points in their argument list.</p> <p>See "New attribute EntryPointTimeout".</p>
1544263	<p>[Agents] Fixed an issue in which the Oracle agent performs an action corresponding to the last error even when it encounters multiple errors, thereby ignoring the previous error numbers. This happens because when the list of errors was parsed by the agent, it moved to the last error and got its state to perform the action corresponding to that error. The priority of actions are: FAILOVER, UNKNOWN, and IGNORE. If any error has FAILOVER/NOFAILOVER, the resource is FAULTED. If any error has UNKNOWN action, the resource is moved to UNKNOWN state. Else, we safely ignore the error and return the state as ONLINE.</p>
1542334	<p>[VCS] Fixed an issue in which the nfs_restart trigger was issuing too many hares -list commands, which impacted the response time of other HA commands invoked from the command line. The HA commands in nfs_postoffline trigger were replaced with more efficient HA commands.</p> <p>Also, the nfs_postoffline and nfs_preonline triggers were moved to the sample_triggers directory so that they are not invoked by default. Users are required to copy both the triggers from /opt/VRTSvcs/bin/sample_triggers to /opt/VRTSvcs/bin/triggers, if the configuration has the NFSRestart agent.</p>
1540807	<p>[GAB] Fixed an issue in which the error number returned by the gab_receive() function in the GAB library is wrong. The gab_receive() function returns -1, but the error number was set to 0.</p>
1539087	<p>[Agents] Fixed an issue in which the agent framework seems to be leaking memory during message logging.</p>
1538208	<p>[VCS] Fixed an issue in which the value of attribute HostUtilization is not 0 even after HostMonitor resource is deleted.</p>
1522568	<p>[Agents] Fixed an issue in which the agent framework crashed while setting the resource name for the dependent attribute.</p>

Table 1-25 Veritas Cluster Server 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1509742	<p>[GAB] Fixed an issue in which GAB membership to VCS (Port h) may not occur, if VCS is brought online before the Port a membership occurs. Clients of the GAB service may not get cluster membership. Symantec recommends that GAB must be configured to provide membership only after a minimum quorum number of nodes join the cluster. If a client of GAB comes up before GAB Port a forms membership on that node, then this client may not get cluster membership until it starts up on at least the configured quorum number of nodes, not even if Port a or any other GAB Ports receive cluster membership. Previously, seeding of Port a would trigger seeding on all other ports by seeding a CONNECTS message on all those ports. However, this led to a race which was fixed via e1424927. The fix disabled CONNECTS which used to propagate the SEED bit to other ports. SEED bit is now propagated to other ports after Port 'a' reconfigures. The master for each port just runs the reconfiguration calculation after Port a reconfigures there.</p>
1504693	<p>[GAB/LLT] Fixed an issue in which LLT cannot provide backenable to GAB. This resulted in an error being produced from the GAB module gabwrite() function.</p>
1478488	<p>[Agents] Fixed an issue in which file systems failed to mount on volume sets and required one or two failovers between nodes to work correctly. When there are no Volume resources and DiskGroup resource has StartVolumes=1, all the volumes of the disk group are started by the DiskGroup agent. If there are lots of volumes/volume sets in a disk group, the volume recovery for all the volumes can take time. The file system on them cannot be mounted before the volumes are started. The fix involved modifying the DiskGroup agent to start all the volumes before initiating recovery.</p>
1377324	<p>[Agents] Fixed a parsing error which caused an error message to appear in the /var/VRTSvcs/log/tmp/Oracle-0 file.</p>
1368385	<p>[Agents] Fixed an issue in which DiskGroupSnap does not work if layered volumes are used. VxVM creates layered volumes by default, in case of larger volumes spanning multiple disks. The agent expects each volume to have a plex at each site but VxVM does not assign a site tag to plex and there is only one top level plex. Thus, the agent reports that the configuration is invalid. This was a limitation in the original agent when no layered volumes were supported.</p>
1070177	<p>[Agents] Fixed an issue to include a new attribute to use the db2start command. There was no option to use the db2start command. Added optional attribute UseDB2start to allow users to start DB2 using the db2start command.</p>

Table 1-25 Veritas Cluster Server 5.0 MP3 RP2 fixed issues (*continued*)

Fixed issues	Description
1070177	[Agents] Fixed an issue to include a new attribute to use the db2start command. There was no option to use the db2start command. Added optional attribute UseDB2start to allow users to start DB2 using the db2start command.

[Table 1-26](#) describes fixed issues in Veritas Cluster Server 5.0 MP3 RP1, which are included in this release.

Table 1-26 Veritas Cluster Server 5.0 MP3 RP1 fixed issues

Fixed issues	Description
1510002	Fixed an issue where the HostMonitor tool misreported CPU as busy when it is idle in an AIX 5.0 MP3 Micropartition environment.
1469381	<p>Fixed an issue where the Share agent was 10x slower on 5.0 MP1 with 300+ Share resources in a service group.</p> <p>Note: This fix changes basic VCS functionality, it is critically important for you to implement these changes for all service groups that contain NFSRestart resources.</p> <p>You must set the value of the PreOnline attribute to 1 for all service groups that contain NFSRestart resources. Failure to set the service group's PreOnline attribute to a value of 1 results in broken NFSRestart resource configurations.</p> <p>The ha commands to change this attribute are:</p> <pre># haconf -makerw # hagr -modify servicegroup_name PreOnline 1 # haconf -dump -makero</pre>
1457429	Removed the VCS NOTICE V-16-1-53021 message after the hastart command is run.
1424927	Optimized GAB connect messages.
1414709	The hagr -offline command and hares -offline command now behave similarly when you bring the last resource in a service group offline.
1404384	Global groups can switch over to a node where WAC is not running, when PreSwitch is set to 1 and HAD runs properly.

Table 1-26 Veritas Cluster Server 5.0 MP3 RP1 fixed issues (*continued*)

Fixed issues	Description
1403471	Reduced time for global cluster fault detection.
1398750	Added the MemCPUAllocator agent. See “ MemCPUAllocator agent ” on page 146.
1397692	Removed a condition where VCS engine clients hung in connect when the target system was down.
1396639	Return code for SCSI commands are now logged before re-using a variable.
1395905	Changes implemented to close device file for device vxdmconfig.
1394624	LLT: fixed an issue where the lltdlv thread spun indefinitely.
1379299	LLT: fixed llt_recordmac() messages.

Storage Foundation for Oracle RAC fixed issues

[Table 1-27](#) describes fixed issues in Storage Foundation for Oracle RAC 5.0 MP3 RP5, which are included in this release.

Table 1-27 Storage Foundation for Oracle RAC 5.0 MP3 RP5 fixed issues

Fixed issues	Description
2276515	Fixed the issue that VCSMMDEBUG shows garbage entries after Oracle CRS has unregistered from VCSMM.
2266848	Fixed the issue that Oracle instance crashed, failure occurred at: vcsipc_dosnd
2053302	Fixed the issue with working of MTU attribute in PrivNIC & MultiPrivNIC resource

[Table 1-28](#) describes fixed issues in Storage Foundation for Oracle RAC 5.0 MP3 RP4, which are included in this release.

Table 1-28 Storage Foundation for Oracle RAC 5.0 MP3 RP4 fixed issues

Fixed issues	Description
1976654	The PrivNIC and MultiPrivNIC agents do not support VIO interfaces with the default MTU size of 64 Kb.
1945042	In WAIT, Don't call lmx poll when vcsipc done queue is not empty and remove light weight tracing.
1935473	LMX should register with NULL canput for performance
1633841	hacf -verify doesn't prompt for incorrect entries for Multiprivnic resource

[Table 1-29](#) describes fixed issues in Storage Foundation for Oracle RAC 5.0 MP3 RP3, which are included in this release.

Table 1-29 Storage Foundation for Oracle RAC 5.0 MP3 RP3 fixed issues

Fixed issues	Description
1877596	Fixed an issue in LMX where it may cause panic due to reuse of a freed buffer.
1847614	Improved the cssd script's performance to use ps command effectively.
1840213	Fixed an issue in Privnic agent where it was not handling the case where resource is configured and after that if netmask is changed manually, it will not update the netmask.

[Table 1-30](#) describes fixed issues in Storage Foundation for Oracle RAC 5.0 MP3 RP2, which are included in this release.

Table 1-30 Storage Foundation for Oracle RAC 5.0 MP3 RP2 fixed issues

Fixed issues	Description
1593859	Reduced the time it takes for Oracle to start with VCSIPC.
1525117	Fixed an issue in which the MultiPrivNIC agent was not able to plumb the IP address on the configured devices.
1382034	Fixed an issue in which the MultiPrivNIC agent failed over the IP address even when it was not required.
1597480	The LMX code now calls the tstop() function calling the tstart() function to avoid a race condition.

[Table 1-31](#) describes fixed issues in Storage Foundation for Oracle RAC 5.0 MP3 RP1, which are included in this release.

Table 1-31 Storage Foundation for Oracle RAC 5.0 MP3 RP1 fixed issues

Fixed issues	Description
1379299	LLT: fixed llt_recordmac() messages.

Veritas Cluster Server agents for Veritas Volume Replicator fixed issues

There are no fixed issues for Veritas Cluster Server agents for Veritas Volume Replicator in 5.0 MP3 RP5 release.

There are no fixed issues for Veritas Cluster Server agents for Veritas Volume Replicator in 5.0 MP3 RP3 release.

[Table 1-32](#) describes fixed issues in Veritas Cluster Server agents for the Veritas Volume Replicator 5.0 MP3 RP2 release.

Table 1-32 Veritas Cluster Server agents for Veritas Volume Replicator 5.0 MP3 RP2 fixed issues

Fixed issues	Description
1671357	Enabled the RVGPrimary agent to migrate a VVR primary to secondary in the case of having multiple secondaries.
1433149	Fixed issues related to the OnlineTimeout attribute with RVGPrimary and RVGSharedPri agents.
1295115	Enabled the fdsetup wizard to set up a firedrill SG in a secured VVR-GCO environment.

Veritas Enterprise Administrator fixed issues

There are no fixed issues for Veritas Enterprise Administrator in 5.0 MP3 RP5 release.

[Table 1-33](#) describes fixed issues in Veritas Enterprise Administrator 5.0 MP3 RP4 release.

Table 1-33 Veritas Enterprise Administrator fixed issues

Fixed issues	Description
1965998	vxsvc coredumping upon startup (VRTSobc)
1925371	vxpal.StorageAgent startup script need to be put into /etc/init.d instead of rc2.d (VRTSdsa)
1914596	isisd wont start, core file generated. (VRTSob)
1840724	Installation of patch VRTSaa 5.0mp3rp2 fails on AIX (VRTSaa)

Storage Foundation and High Availability known issues

The following sections describe the Veritas Storage Foundation High Availability (HA) known issues in this release.

- [Storage Foundation and High Availability known issues](#)
- [Veritas Volume Manager known issues](#)
- [Veritas File System known issues](#)
- [Storage Foundation Cluster File System known issues](#)
- [Storage Foundation for Oracle known issues](#)
- [Storage Foundation for DB2 known issues](#)
- [Storage Foundation for Oracle RAC known issues](#)
- [Veritas Cluster Server known issues](#)

Storage Foundation and High Availability known issues

The following are the Storage Foundation and High Availability issues that are known in this release.

Storage Foundation Manager 1.1.1 Central Server

The procedure to centrally manage Storage Foundation 5.0 MP3 RP5 hosts on Storage Foundation Manager 1.1.1 can be viewed at the following URL:

<http://entsupport.symantec.com/docs/315384>

Including and excluding ASLs on CLARiiON arrays can cause failed states (2008336)

When a CLARiiON array is configured to the given host, then the system administrator is not supposed to exclude and then include the corresponding ASL (CLARiiON). The LUNS can go to failed state if such operations are attempted.

Veritas Volume Manager known issues

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

VIOC: system is not coming up after enabling DMP rootability support on an alternative disk (2068475)

This issue exists in 5.0MP3 RP4 and onward. This situation can happen due to multiple odm entries present. You need to verify and remove additional entries.

To verify and remove additional entries

1 Find the attributes of the disk.

```
# lsattr -El hdisk18
max_transfer 0x40000 Maximum TRANSFER Size      True
pvid 00c9849012936dfd0000000000000000 Physical volume identifier False
queue_depth  3 Queue DEPTH                          True
reserve_policy no_reserve Reserve Policy            True
```

Multiple copies of the attributes exist in the output.

2 Query the ODM:

```
# odmget -q "attribute = pvid and uniquetype = disk/vscsi/\  
dmpvdisk" PdAt
```

```
PdAt:  
    uniquetype = "disk/vscsi/dmpvdisk"  
    attribute = "pvid"  
    deflt = "none"  
    values = ""  
    width = ""  
    type = "R"  
    generic = "D"  
    rep = "s"  
    nls_index = 11
```

```
PdAt:  
    uniquetype = "disk/vscsi/dmpvdisk"  
    attribute = "pvid"  
    deflt = "none"  
    values = ""  
    width = ""  
    type = "R"  
    generic = "D"  
    rep = "s"  
    nls_index = 11
```

```
PdAt:  
    uniquetype = "disk/vscsi/dmpvdisk"  
    attribute = "pvid"  
    deflt = "none"  
    values = ""  
    width = ""  
    type = "R"  
    generic = "D"  
    rep = "s"  
    nls_index = 11
```

```
PdAt:  
    uniquetype = "disk/vscsi/dmpvdisk"  
    attribute = "pvid"  
    deflt = "none"  
    values = ""  
    width = ""  
    type = "R"  
    generic = "D"
```

```
rep = "s"  
nls_index = 11
```

- 3 Keep only one entry per PdAt. Remove additional entries using the `odmdelete` command.

```
# odmdelete -o -PdAt -q "attribute = pvid and uniquetype = disk/\vscsi/dmpvdisk"
```

Perform this for the `max_transfer`, the `queue_depth`, the `reserve_policy` attributes too.

- 4 In the previous step, the command deletes all occurrences of PdAt entry. You now need to manually add the PdAt entry for `pvid`, `max_transfer`, `queue_depth`, `reserve_policy`. For example, copy the following to a file:

PdAt:

```
uniquetype = "disk/vscsi/dmpvdisk"  
  
attribute = "pvid"  
  
deflt = "none"  
  
values = ""  
  
width = ""  
  
type = "R"  
  
generic = "D"  
  
rep = "s"  
  
nls_index = 11
```

Once done, call `odmadd`. Perform this for the `max_transfer`, `queue_depth`, `reserve_policy` attributes.

Note: Verify that all the attributes appear only once.

If this procedure does not resolve the issue, or if multiple entries are not seen, contact Symantec support.

The vxvmconvert command may fail with PowerPath (1884301)

There are some known issues when using the vxvmconvert command with PowerPath.

Workaround:

Use the `vxconvert` command.

DMP attributes do not remain persistent (1887335)

After upgrading from 5.0 to 5.0 MP3 RP5, DMP attributes such as iopolicy and recovery option does not remain persistent.

Workaround:

You need to either set the policies again or edit the `/etc/vx/dmppolicy.info` file to have all enclosure names as lower case.

Slowness in disk group shared disk group creation operation on IBM DS4800 Arrays (2068466)

On IBM DS4800 arrays creation of a shared disk group is slow. VxVM does not hang but the operation takes some time to complete. This happens only with the DS4800 array.

Evaluate the need for intelligence in vxattachd to clear stale keys on failover/shared dg's in CVM and non CVM environment (1880279)

The sites/plexes might not be automatically reattached when the faulted disks are reconnected. Some of the dmp events might be ignored by the auto reattach daemon (`vxattachd`) due to the overflow of the events buffer. The auto reattach daemon might fail to reattach the site/disk in the presence of stale PGR keys in the disks.

Workaround:

Manually remove the stale keys that are still present on the reconnected faulted disks and then fire the `vxreattach` command to initiate the reattach.

vxsnap failed to prepare the volume set (2365466)

The `vxsnap` preparation on a volume set failed to create `dco` volume because it is unable to allocate the required space.

Workaround:

Whenever `vxsnap` prepare fails with space allocation issue, use the following the command to resolve it:

```
# vxsnap -g diskgroup prepare vset alloc=device_name
```

Where `device_name` can be picked up from the list of devices available in that diskgroup.

CDS (Cross-Platform Data Sharing) disks of size greater than 1 terabyte can experience data corruption (2291226)

Any disk of size greater than 1 terabyte, initialized as CDS (Cross-Platform Data Sharing) disk under VxVM, may experience block level data corruption at 1 terabyte boundary. The data corruption pattern would match with the CDS disk's backup label which is similar to Solaris SMI label.

Any CDS disk resize from less than 1 terabyte to more than 1 terabyte can also lead to this issue.

Workaround :

Move (relocate) the user data to CDS disk of size less than 1 terabyte and then migrate to other UNIX platforms if required. For more details refer to: <http://www.symantec.com/docs/TECH155921>.

Veritas File System known issues

The following are Veritas File System issues that are known in this release.

During file relocation file system can encounter a no space left on device condition (1888354)

File system can encounter a (possibly spurious) ENOSPC (no space left on device) condition during file relocation.

A failure message displays while rejecting VRTScavf (1744917)

While rejecting VRTScavf 5.0.3.200 with the `<command> installp -r</command>` command, the following failure message displays on the terminal:

```
/usr/lib/instl/reject[1065]: 258274 Segmentation fault (coredump)  
Finished processing all filesets. (Total time: time secs).
```

Workaround:

Run the following command manually after rejecting the patch:

```
# /usr/bin/ln -s /opt/VRTSvcs/bin/ScriptAgent \  
/opt/VRTSvcs/bin/CFSMount/CFSMountAgent
```

Storage Foundation Cluster File System known issues

The following are the Storage Foundation Cluster File System issues that are known in this release.

Mounting a filesystem as ‘seconly’ using cfsmount command may fail (2042351)

Mounting a filesystem as seconly using `cfsmount` command may fail with the following error:

```
Error: V-35-50: Could not mount [<volume name>] at <mount point>  
on <node name> Look at VCS engine_A.log on <node name> for  
possible errors for resource cfsmount1
```

Due to a timing issue in the `cfsmount` script, the seconly file system tried to mount before the primary mount operation is complete, which fails with the above mentioned error.

Storage Foundation for Oracle known issues

The following are the Storage Foundation for Oracle issues that are known in this release.

vxstorage_stats and dbed_analyzer might dump core (1899723)

`vxstorage_stats` and `dbed_analyzer` might dump core in some cases when the `gapsnapshot` plugin tries to claim the disk object.

Workaround:

To resolve this issue, enter the following commands:

```
■ # cd /opt/VRTSvxms/lib/map  
■ # mv libgapdisk.so libgapdisk.so_bak  
■ # mv libgapsnapshot.so libgapsnapshot.so_bak
```

The database fails over during Flashsnap operations (1469310)

In an SFHA environment, if the database fails over during Flashsnap operations such as the `dbed_vmsnap -o resync` command and various error messages appear. This issue occurs because Flashsnap commands do not create a VCS resource for

the SNAP disk group. As such, when the database fails over, only the primary disk group is moved to another node.

Workaround:

There is no workaround for this issue. The error messages depend on the timing of the database failover.

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

Storage Foundation for DB2 known issues

The following are the Storage Foundation for DB2 issues that are known in this release.

vxstorage_stats and dbed_analyzer might dump core (1899723)

vxstorage_stats and dbed_analyzer might dump core in some cases when the gapsnapshot plugin tries to claim the disk object.

Workaround:

To resolve this issue, enter the following commands:

- # cd /opt/VRTSvxms/lib/map
- # mv libgapdisk.so libgapdisk.so_bak
- # mv libgapsnapshot.so libgapsnapshot.so_bak

The database fails over during Flashsnap operations (1475719)

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

Workaround:

There is no workaround for this issue. The error messages depend on the timing of the database failover.

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

Storage Foundation for Oracle RAC known issues

The following are the Storage Foundation for Oracle RAC issues that are known in this release.

dbed_clonedb error when oracle pfile including DB_UNIQUE_NAME parameter(2317766)

A dbed_clonedb error occurs when the oracle pfile includes the DB_UNIQUE_NAME parameter.

Workaround:

- 1 Shut down the clone database using sqlplus.
- 2 Edit the clone database's pfile.
- 3 Modify DB_UNIQUE_NAME parameter and make it same as DB_NAME parameter.
- 4 Startup clone database using modified pfile.

Installing Oracle CRS 10R2 or 11gR1 on AIX 6.1 TL3 SP1 fails to start VIP

Installing Oracle CRS 10R2 or 11gR1 on AIX 6.1 TL3 SP1 fails to start VIP and displays the following error messages:

```
CRS-1006: No more members to consider  
CRS-0215: Could not start resource 'ora.system1.vip'.
```

Workaround:

See the Oracle MetaLink document ID 805536.1, "VIP cannot start on AIX 6.1 because netstat has a new column."

dbed_vmclonedb and dbed_clonedb fails on a RAC database when the clone SID name contains the primary SID name in the beginning (1743179)

If a clone SID name contains the primary SID name in the beginning, the dbed_vmclonedb and dbed_clonedb command fails with the following message:

```
ERROR V-81-4882 An error occurred while reconfiguring \  
Oracle instance 'clone_SID'
```

For example, the following commands, which have "Prod" as the primary SID and "Prod1" as the clone SID, produce this error message:

```
# dbed_vmclondb -S Prod -o recoverdb \  
new_sid=Prodl,server_name=srv_name -f snapplan -r relocate_path  
# dbed_vmclondb -S Prod -o mountdb \  
new_sid=Prodl,server_name=srv_name -f snapplan -r relocate_path  
# dbed_clonedb -S Prodl -m mount_point -c ckpt_name
```

Workaround:

Do not use a clone SID name that contains primary SID name in the beginning.

Issue with PrivNIC.cf and MultiPrivNIC.cf files (2053877)

When SF Oracle RAC 5.0 MP3 RP5 is installed, updated PrivNIC.cf and MultiPrivNIC.cf files do not get copied to /etc/VRTSvcs/conf/config directory.

Workaround:

You have to manually copy these files from /etc/VRTSvcs/conf directory to /etc/VRTSvcs/conf/config directory:

```
# cp -p /etc/VRTSvcs/conf/PrivNIC.cf \  
/etc/VRTSvcs/conf/config/PrivNIC.cf  
# cp -p /etc/VRTSvcs/conf/MultiPrivNIC.cf \  
/etc/VRTSvcs/conf/config/MultiPrivNIC.cf
```

The status of Veritas vxdbd cannot be retrieved on one of the cluster nodes after reboot (1802646, 1277640, 1848199)

Sometimes upon the reboot of the cluster, the status of Veritas vxdbd cannot be retrieved on one of the cluster nodes. This is caused by erroneous corruption, removal or truncation of the /var/VRTSat/.VRTSat/profile/VRTSatlocal.conf file on that node.

Workaround:

To resolve this issue,

- 1 Stop vxdbd on the node missing VRTSatlocal.conf file by executing the following command:

```
# /opt/VRTSdbcom/bin/vxdbdctrl stop
```

- 2 Copy the /var/VRTSat/.VRTSat/profile/VRTSatlocal.conf file from the other node, which is not experiencing this issue, to this node.

- 3 Start vxdbd with the new VRTSatlocal.conf file:

```
# /opt/VRTSdbcom/bin/vxdbdctrl start
```

Veritas Cluster Server known issues

There are no known issues for Veritas Cluster Server in this release.

Software limitations

The following sections describe the Veritas Storage Foundation High Availability (HA) software limitations in this release.

- [Veritas File System software limitations](#)
- [Storage Foundation Cluster File System software limitations](#)
- [Storage Foundation for Oracle software limitations](#)
- [Storage Foundation for DB2 software limitations](#)
- [Veritas Cluster Server software limitations](#)

Veritas File System software limitations

There are no software limitations for Veritas File System in this release.

Storage Foundation Cluster File System software limitations

There are no software limitations in this release of Storage Foundation Cluster File System.

Storage Foundation for Oracle software limitations

The following are the Storage Foundation for Oracle software limitations that are known in this release.

Older backups failing to be restored using the DBED scripts

If you are currently using backup and restore for the DBED repository, it is crucial to perform a full backup of the DBED repository database after installing 5.0 MP3 RP5. Otherwise, prior backups cannot be restored using the 5.0 MP3 RP5 restore script.

See the Veritas Storage Foundation for Oracle Administrator's Guide for the `sfa_rept_adm` command.

See "[Storage Foundation for Oracle fixed issues](#)" on page 42. for incident 1425261.

Storage Foundation for DB2 software limitations

The following are the Storage Foundation for DB2 software limitations that are known in this release.

No support for running DBED commands on Cluster File System

Storage Foundation for DB2 does not support running DBED commands on Cluster File System.

Older backups failing to be restored using the DBED scripts

If you are currently using backup and restore for the DBED repository, it is crucial to perform a full backup of the DBED repository database after installing 5.0 MP3 RP5. Otherwise, prior backups cannot be restored using the 5.0 MP3 RP5 restore script.

See the *Veritas Storage Foundation for DB2 Administrator's Guide* for the `sfua_rept_adm` command.

See "[Storage Foundation for DB2 fixed issues](#)" on page 44. for incident 1425261.

Veritas Cluster Server software limitations

The following are the Veritas Cluster Server software limitations that are known in this release.

Unable to bring Oracle ASM online on the second node of the cluster (2056688)

Symptom:

- 1 When a node reboots the ohasd process does not automatically start the ocspd.bin process.
- 2 When the ocspd.bin process starts, it automatically starts the ASM instance leading to concurrency violation if the ASMInst resource is configured in failover service group.

Description: When a node reboots the ohasd process does not automatically start the ocspd.bin process because the AUTO_START attribute of ora.cssd resource under CRS control is set to 'never'. The ocspd.bin process triggers startup of ASM instance automatically due to the ENABLED attribute of ora.asm resource under CRS control is set to 1.

Workaround:

To set the AUTO_START to 'always' and to disable the ENABLED attribute run the following commands:

```
$GRID_HOME/bin/crsctl modify resource ora.cssd -attr AUTO_START=always  
$GRID_HOME/bin/crsctl modify resource ora.asm -attr ENABLED=0
```

Note: This is applicable only for 11gR2.

Changes in behavior for Storage Foundation High Availability

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

About the installrp script

From version 5.0 MP3 RP5, Veritas Storage Foundation and High Availability Solutions provides a new upgrade script.

To upgrade from Veritas Storage Foundation and High Availability Solutions version 5.0 MP3 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. After using the installrp script you will need to reboot your system.

Note: If you have the VRTSvcsy fileset installed on your system, you need to update it manually. The installrp script does not update VRTSvcsy fileset. Refer to the operating system manuals for updating the fileset.

installrp script options

[Table 1-34](#) shows command line options for the product upgrade script.

Table 1-34 Available command line options

Command line options	Function
[<system1> <system2>...]	Specifies the systems on which to run the upgrade options. If not specified, the command prompts for a system name.
[-precheck]	The -precheck option is used to confirm that systems meet the products install requirements before installing.

Table 1-34 Available command line options (*continued*)

Command line options	Function
[-responsefile <response_file>]	The -responsefile option is used to perform automated installations or uninstalls using information stored in a file rather than prompting for information. <response_file> is the full path of the file that contains configuration definitions. The -enckeyfile option is required with the -responsefile option when the response file contains encrypted passwords.
[-patchpath <patch_path>]	The -patchpath 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 installrp.
[-tmppath <tmp_path>]	The -tmppath option is used to select a directory other than /var/tmp as the working directory for installrp. This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
[-logpath <log_path>]	The -logpath option is used to select a directory other than /opt/VRTS/install/logs as the location where installrp log files, summary file, and response file are saved.
[-rootpath <root_path>]	The -rootpath option is used to re-root the install of all packages to the given path. On Solaris, -rootpath passes -R <root_path> to pkgadd. The -rootpath option is not supported on AIX.
[-keyfile <ssh_key_file>]	The -keyfile option specifies a key file for SSH. When this option is used, -i <ssh_key_file> is passed to every SSH invocation.

Table 1-34 Available command line options (*continued*)

Command line options	Function
[-rsh]	The -rsh option is used when rsh and rcp are to be used for communication between systems instead of ssh and scp. When the -rsh option is not used, systems must be pre-configured such that ssh commands between systems execute without prompting for passwords or confirmations.
[-listpatches]	The -listpatches option is used to display product patches in correct installation order.

Changes in Veritas Cluster Server behavior

The following sections describe changes in Veritas Cluster Server behavior for this release.

Removing deprecated Cluster Manager (Java Console) directories

VCS Cluster Manager (Java Console) versions prior to 5.0 MP3 RP5 created the following directories under /opt/VRTSvcs/gui, on the systems where Java Console was installed:

- /opt/VRTSvcs/gui/attrpool
- /opt/VRTSvcs/gui/messages

Any users who could login to these systems had read-write permissions on these directories.

In the 5.0 MP3 RP5 release, VCS Cluster Manager (Java Console) no longer requires these directories. If you upgrade to 5.0 MP3 RP5 delete these directories. Deleting these directories does not affect VCS Java Console functionality.

To delete the directories, perform the following commands:

- # rm -rf /opt/VRTSvcs/gui/attrpool
- # rm -rf /opt/VRTSvcs/gui/messages

Mandatory configuration change for the NFS and NFSRestart resources

You must perform the following instructions for VCS configurations that have NFSRestart resources. Failure to perform these instructions can result in NFS/NFSRestart resources not functioning correctly.

Symantec implemented this change to prevent the invocation of NFSRestart-related triggers when no NFSRestart resources in the VCS configuration.

To copy the `nfs_preonline` and `nfs_postoffline` files

- ◆ Copy the `nfs_preonline` and `nfs_postoffline` files to the `/opt/VRTSvcs/bin/triggers` directory.

```
# cp /opt/VRTSvcs/bin/sample_triggers/nfs_preonline \  
/opt/VRTSvcs/bin/triggers  
# cp /opt/VRTSvcs/bin/sample_triggers/nfs_postoffline \  
/opt/VRTSvcs/bin/triggers
```

Attributes to control the scheduling of class and priority of agent entry points

Symantec has introduced four new attributes—`EPPriority`, `EPClass`, `OnlinePriority`, and `OnlineClass`—to enable you to control the scheduling of class and priority of the agent functions or entry points. The new attributes `OnlineClass` and `OnlinePriority` are used to set the scheduling class and priority for the online entry point. The new attributes `EPClass` and `EPPriority` are used to set the scheduling class and priority for all entry points, except the online entry point.

These attributes provide a single interface to tune the scheduling parameters for all entry points (except the online entry point). It does not matter if they are implemented as C-based or script-based entry points. The `OnlineClass` and `OnlinePriority` attributes provide the same functionality for only the online entry point.

It is usually required that the monitor, clean, offline and the other entry points running on an application have a higher scheduling class or priority without which they would compete with the application for system resources. However, running the online entry point with a higher scheduling class or priority may create problems because applications inherit the scheduling parameters from the application vendors, who specify that the applications are run using the default operating system scheduling parameters. Also, the online entry point is usually invoked before you start the application and the system is not very busy.

Hence, you must usually set the values of EPPriority and EPClass attributes to a higher value than the default value. You must usually set the value of the OnlinePriority and OnlineClass attribute to the default operating system scheduling values.

Note: You must either use all four new attributes or set them to -1 to go back to using the older Agent* and Script* attributes. A combination of the two different sets of attributes is not supported.

Table 1-35 indicates the values that apply to these new attributes.

Table 1-35 Attribute values to schedule class and priority of agent entry points

Attributes	Values
OnlineClass / EPClass	The default value for the attribute is -1. This indicates that this attribute is not in use and hence VCS will use the older AgentClass / AgentPriority and ScriptClass / ScriptPriority attributes.
OnlinePriority / EPPriority	<p>The default value for this attribute is -1. This indicates that this attribute is not in use and hence, VCS will use the older AgentClass/Priority and ScriptClass/Priority attributes.</p> <p>If the value of this attribute is 0, it indicates the base operating system priority for the configured scheduling class.</p> <p>For TS (TimeShare) class, the default priority value in VCS is 0 which translates to a process-priority of 60. If you set a higher or lower value for priority, it does not take effect because on AIX, a TS class process can only take a priority value of 60.</p> <p>For RT (RealTime) class, the default priority value in VCS of 0 translates to a process-priority of 0 on AIX. If you set a higher value than 0 in VCS, the same value will directly reflect for the process. For example: if you set the EPPriority to 10 and EPClass to RT, then the entry point processes run with a priority of 10. However, in RT scheduling class, you cannot set a value for priority that is lower than 0.</p>

New attribute AEPTIMEOUT

The new attribute EntryPointTimeout is used to pass the entry point timeout value as a parameter to agent entry points in their argument list. This is an internal attribute and you are not required to change the value of this attribute. This attribute is strictly for the use of agent developers.

New attribute WaitForRecovery in Sybase agent

Added a new attribute WaitForRecovery in Sybase agent. If this attribute is enabled, during the online function, the agent waits till recovery has been completed and all databases that can be made online are brought online. Sybase agent type definition should be updated in order to use the WaitForRecovery attribute.

Changes in Storage Foundation for Oracle RAC behavior

The following sections describe changes in Storage Foundation for Oracle RAC behavior for this release.

Storage key support for Storage Foundation for Oracle RAC

Storage key support is qualified in the SF Oracle RAC environment with file system components that have the full storage keys implementation and the VCS and SFRAC components that work in the legacy mode.

PrivNIC and MultiPrivNIC agents support for MTU size

The PrivNIC and MultiPrivNIC agents now support VIO interfaces with the default MTU size of 64 Kb.

Downloading the rolling patch archive

The patches included in the 5.0 MP3 RP5 release are available for download from the Symantec website. After downloading the 5.0 MP3 RP5 file, use gunzip and tar to uncompress and extract.

For the 5.0 MP3 RP5 download archive and instructions, see the following TechNote on the Symantec Technical Support Web site:

<http://www.symantec.com/docs/TECH46478>

Filesets included in this rolling patch

This section describes the AIX filesets included in this rolling patch:

[Veritas Cluster Server filesets](#)

[Storage Foundation Cluster File System filesets](#)

[Storage Foundation for Oracle RAC filesets](#)

[Storage Foundation filesets](#)

[Storage Foundation for DB2 filesets](#)

[Storage Foundation for Oracle filesets](#)

[Veritas Volume Replicator filesets](#)

Veritas Cluster Server filesets

[Table 1-36](#) describes the VCS filesets that are included in this rolling patch:

Table 1-36 VCS filesets

Filesets
VRTScscm.rte
VRTSgab.rte
VRTSllt.rte
VRTSvcsc.rte
VRTSvcscag.rte
VRTSvcscdb.rte
VRTSvcscor.rte
VRTSvcscsy.rte
VRTSvcscfen.rte

Storage Foundation Cluster File System filesets

[Table 1-37](#) describes the SFCFS filesets that are included in this rolling patch:

Table 1-37 SFCFS filesets

Filesets
VRTSobc33
VRTSob
VRTSobgui
VRTSccg
VRTSaa
VRTSllt.rte
VRTSgab.rte
VRTSvxfen.rte
VRTSvcs.rte
VRTSvcsag.rte
VRTSweb.rte
VRTSscsm.rte
VRTSvxvm
VRTSfspro
VRTSvmman
VRTSvmpro
VRTSdcli
VRTSalloc
VRTSvcsvr
VRTSvxfs
VRTSglm
VRTScavf

Storage Foundation for Oracle RAC filesets

[Table 1-38](#) describes the SF for Oracle RAC filesets that are included in this rolling patch:

Table 1-38 SF for Oracle RAC filesets

Filesets
VRTSaa
VRTSalloc
VRTScavf
VRTSccg
VRTScscm.rte
VRTSdbac.rte
VRTSdbcom
VRTSdbed
VRTSdbms3
VRTSdcli
VRTSfspro
VRTSgab.rte
VRTSglm
VRTSllt.rte
VRTSob
VRTSobc33
VRTSobgui
VRTSodm
VRTSvcsvcs.rte
VRTSvcsvcsag.rte
VRTSvcsvcsor.rte
VRTSvcsvr
VRTSvdid.rte
VRTSvmman
VRTSvmpro
VRTSvxfen.rte
VRTSvxfs
VRTSvxvm
VRTSweb.rte
VRTSdsa

Storage Foundation filesets

[Table 1-39](#) describes the Storage Foundation filesets that are included in this rolling patch:

Table 1-39 Storage Foundation filesets

Filesets
VRTSobc33
VRTSob
VRTSobgui
VRTSccg
VRTSaa
VRTSvxvm
VRTSfspro
VRTSvmman
VRTSvmpro
VRTSdcli
VRTSalloc
VRTSvcsvr
VRTSweb.rte
VRTSvxfs
VRTSllt.rte
VRTSgab.rte
VRTSvxfen.rte
VRTSvcs.rte
VRTSvcsag.rte
VRTScscm.rte

Storage Foundation for DB2 filesets

[Table 1-40](#) describes the Storage Foundation for DB2 filesets that are included in this rolling patch:

Table 1-40 Storage Foundation for DB2 filesets

Filesets
VRTSobc33
VRTSob
VRTSobgui
VRTSccg
VRTSaa
VRTSvxvm
VRTSfspro
VRTSvmman
VRTSvmpro
VRTSdcli
VRTSalloc
VRTSvcsvr
VRTSweb.rte
VRTSvxfs
VRTSdbms3
VRTSllt.rte
VRTSgab.rte
VRTSvxfen.rte
VRTSvcs.rte
VRTSvcsag.rte
VRTScscm.rte
VRTSdbcom
VRTSdb2ed
VRTSd2gui
VRTSvcsdb.rte

Storage Foundation for Oracle filesets

[Table 1-41](#) describes the SF for Oracle filesets that are included in this rolling patch:

Table 1-41 Storage Foundation for Oracle filesets

Filesets
VRTSobc33
VRTSob
VRTSobgui
VRTSccg
VRTSaa
VRTSvxvm
VRTSfspro
VRTSvmman
VRTSvmpro
VRTSdcli
VRTSalloc
VRTSvcsvr
VRTSweb.rte
VRTSvxfs
VRTSdbms3
VRTSllt.rte
VRTSgab.rte
VRTSvxfen.rte
VRTSvcs.rte
VRTSvcsag.rte
VRTScscm.rte
VRTSdbcom
VRTSdbed
VRTSorgui
VRTSvcsor.rte

Veritas Volume Replicator filesets

[Table 1-42](#) describes the Veritas Volume Replicator filesets that are included in this rolling patch:

Table 1-42 Veritas Volume Replicator filesets

Filesets
VRTSobc33
VRTSob
VRTSobgui
VRTSccg
VRTSaa
VRTSvxxvm
VRTSfspro
VRTSvman
VRTSvmpro
VRTSdcli
VRTSalloc
VRTSvcsvr
VRTSweb.rte

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 5.0 MP3 RP5. Review the 5.0 MP3 Installation Guide and Release Notes for your product.

See the *Veritas Storage Foundation™ for Oracle® RAC 5.0 MP3 RP5 Application Note: Installing or upgrading to Oracle RAC 11g Release 2* for installing or upgrading Oracle RAC 11gR2.

To install the Veritas software for the first time:

- 1 Mount the 5.0 MP3 product image or disc and navigate to the folder that contains the installation program. Choose one of the following to start the installation:

- For Storage Foundation:

```
# ./installsf -installonly [-rsh] node1 node2 ... nodeN
```

- For Storage Foundation Cluster File System:

```
# ./installsfdfs -installonly [-rsh] node1 node2 ... nodeN
```

- For Storage Foundation for Oracle:

```
# ./installsfora -installonly [-rsh] node1 node2 ... nodeN
```

- For Storage Foundation for Oracle RAC:

```
# ./installsfrac -installonly [-rsh] node1 node2 ... nodeN
```

- For Veritas Cluster Server:

```
# ./installvcs -installonly [-rsh] node1 node2 ... nodeN
```

The `-installonly` option is required to perform the installation without configuring the software.

- 2 Review the installation prerequisites for upgrading to 5.0 MP3 RP5. See [“Prerequisites for upgrading to 5.0 MP3 RP5”](#) on page 88.
- 3 Mount the 5.0 MP3 RP5 product image or disc and navigate to the folder that contains the installation program.

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

See [“About the installrp script”](#) on page 74. and See [“installrp script options”](#) on page 74. for more information.

- 4 Reboot the nodes.
- 5 Navigate to the folder that contains the installation program. Run the same 5.0 MP3 installation script that you used in step 1, this time specifying the `-configure` option to configure the software.

```
# cd /opt/VRTS/install
```

- For Storage Foundation:

```
# ./installsf -configure [-rsh] node1 node2 ... nodeN
```

- For Storage Foundation Cluster File System:

```
# ./installsfcfs -configure [-rsh] node1 node2 ... nodeN
```

- For Storage Foundation for Oracle:

```
# ./installsfora -configure [-rsh] node1 node2 ... nodeN
```

- For Storage Foundation for Oracle RAC:

```
# ./installsfrac -configure [-rsh] node1 node2 ... nodeN
```

- For Veritas Cluster Server:

```
# ./installvcs -configure [-rsh] node1 node2 ... nodeN
```

Prerequisites for upgrading to 5.0 MP3 RP5

The following list describes prerequisites for upgrading to the 5.0 MP3 RP5 release:

- For any product in the Storage Foundation stack, regardless of your operating system, you must have the 5.0 MP3 release installed before you can upgrade that product to the 5.0 MP3 RP5 release.
- Each system must have sufficient free space to accommodate patches.
- Before you begin the upgrade, check the readiness of the nodes that you plan to upgrade. The command to start the pre-upgrade check is:

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

Supported upgrade paths

- 5.0 MP3 to 5.0 MP3 RP5
- 5.0 MP3 RP1 to 5.0 MP3 RP5
- 5.0 MP3 RP2 to 5.0 MP3 RP5
- 5.0 MP3 RP3 to 5.0 MP3 RP5
- 5.0 MP3 RP4 to 5.0 MP3 RP5

Upgrading 5.0 MP3 to 5.0 MP3 RP5

This section describes how to upgrade from 5.0 MP3 to 5.0 MP3 RP5 on a cluster or a standalone system.

- Installing or upgrading to Oracle RAC 11gR2.
See the Veritas Storage Foundation™ for Oracle® RAC 5.0 MP3 RP5 Application Note: Installing or upgrading to Oracle RAC 11g Release 2.
- [Performing a phased upgrade to 5.0 MP3 RP5 on a cluster](#)

Use the procedures to perform a phased upgrade to 5.0 MP3 RP5 on a cluster that has VCS, SFHA, SF for Oracle HA or SF for DB2 HA, SFCFS, or SF for Oracle RAC installed and configured.

- [Performing a full upgrade to 5.0 MP3 RP5 on a cluster](#)

Use the procedures to perform a full upgrade to 5.0 MP3 RP5 on a cluster that has VCS, SFHA, SF for Oracle HA or SF for DB2 HA, SFCFS, or SF for Oracle RAC installed and configured.

- [Upgrading to 5.0 MP3 RP5 on a standalone system](#)

Use the procedure to upgrade to 5.0 MP3 RP5 on a system that has Storage Foundation, SF for Oracle, or SF for DB2 installed.

Note: For Veritas Volume Manager before you upgrade to 5.0 MP3 RP5, you must disable DMP control on the root disk if it is enabled.

Performing a phased upgrade to 5.0 MP3 RP5 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.

In the following procedures a "subcluster" is a rough division of the cluster into two parts. Split the cluster so that your high priority service groups can remain online during the upgrade of the first subcluster.

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

- 1 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).
- 2 Move all the service groups from the group A to group B.
- 3 For the nodes in group A, start the upgrade using the `installrp` script.
- 4 Get the nodes in group B ready.
- 5 Activate the nodes in group A, then bring the service groups online.
- 6 Upgrade the nodes in group B.

Depending on your cluster's configuration, select one of the following procedures to upgrade to 5.0 MP3 RP5:

- [Performing a phased upgrade to 5.0 MP3 RP5 for VCS](#)

- [Performing a phased upgrade to 5.0 MP3 RP5 on a Storage Foundation HA cluster](#)

- [Performing a phased upgrade to 5.0 MP3 RP5 on a Storage Foundation Cluster File System cluster](#)
- [Performing a phased upgrade to 5.0 MP3 RP5 on a Storage Foundation for Oracle RAC cluster](#)

Performing a phased upgrade to 5.0 MP3 RP5 for VCS

The following procedure describes performing a phased upgrade for VCS.

Moving the service groups to the second subcluster

Perform the following steps to establish the service group's status and to switch the service groups.

To move service groups to the second subcluster

- 1 On the first subcluster, determine where the service groups are online.

```
# hagr -state
```

The output resembles:

#Group	Attribute	System	Value
parallel_sg1	State	node01	ONLINE
parallel_sg1	State	node02	ONLINE
parallel_sg1	State	node03	ONLINE
parallel_sg1	State	node04	ONLINE
parallel_sg2	State	node01	ONLINE
parallel_sg2	State	node02	ONLINE
parallel_sg2	State	node03	ONLINE
parallel_sg2	State	node04	ONLINE
failover_sg1	State	node01	ONLINE
failover_sg1	State	node02	OFFLINE
failover_sg1	State	node03	OFFLINE
failover_sg1	State	node04	OFFLINE
failover_sg2	State	node01	OFFLINE
failover_sg2	State	node02	ONLINE
failover_sg2	State	node03	OFFLINE
failover_sg2	State	node04	OFFLINE
VxSS	State	node01	ONLINE
VxSS	State	node02	ONLINE
VxSS	State	node03	ONLINE
VxSS	State	node04	ONLINE

- 2 Bring offline the parallel service groups (parallel_sg1 and parallel_sg2) and the VXSS group from the first subcluster. Switch the failover service groups (failover_sg1 and failover_sg2) from the first subcluster (node01 and node02) to the nodes on the second subcluster (node03 and node04).

```
# hagr -offline parallel_sg1 -sys node01
# hagr -offline parallel_sg2 -sys node01
# hagr -offline parallel_sg1 -sys node02
# hagr -offline parallel_sg2 -sys node02
# hagr -offline VxSS -sys node01
# hagr -offline VxSS -sys node02
# hagr -switch failover_sg1 -to node03
# hagr -switch failover_sg2 -to node04
```

- 3 On the nodes in the first subcluster, unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -k
```

Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on
/dev/hd4	20971520	8570080	60%	35736	2%	/
/dev/hd2	5242880	2284528	57%	55673	9%	/usr
/dev/hd9var	4194304	3562332	16%	5877	1%	/var
/dev/hd3	6291456	6283832	1%	146	1%	/tmp
/dev/hd1	262144	261408	1%	62	1%	/home
/dev/hd11admin	262144	184408	30%	6	1%	/admin
/proc	-	-	-	-	-	/proc
/dev/hd10opt	20971520	5799208	73%	65760	5%	/opt
/dev/vx/dsk/dg2/dg2vol1		10240	7600	26%	4	1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2		10240	7600	26%	4	1% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3		10240	7600	26%	4	1% /mnt/dg2/dg2vol3

```
# umount /mnt/dg2/dg2vol1
```

```
# umount /mnt/dg2/dg2vol2
```

```
# umount /mnt/dg2/dg2vol3
```

- 4 On the nodes in the first subcluster, stop all VxVM volumes (for each disk group) that VCS does not manage.
- 5 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

- 6 Freeze the nodes in the first subcluster.

```
# hasys -freeze -persistent node01
```

```
# hasys -freeze -persistent node02
```

- 7 Dump the configuration and make it read-only.

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

- 8 Verify that the service groups are offline on the first subcluster that you want to upgrade.

```
# hagrps -state
```

Output resembles:

```
#Group Attribute System Value
parallel_sg1 State node01 |OFFLINE|
parallel_sg1 State node02 |OFFLINE|
parallel_sg1 State node03 |ONLINE|
parallel_sg1 State node04 |ONLINE|
parallel_sg2 State node01 |OFFLINE|
parallel_sg2 State node02 |OFFLINE|
parallel_sg2 State node03 |ONLINE|
parallel_sg2 State node04 |ONLINE|
failover_sg1 State node01 |OFFLINE|
failover_sg1 State node02 |OFFLINE|
failover_sg1 State node03 |ONLINE|
failover_sg1 State node04 |OFFLINE|
failover_sg2 State node01 |OFFLINE|
failover_sg2 State node02 |OFFLINE|
failover_sg2 State node03 |OFFLINE|
failover_sg2 State node04 |ONLINE|
VxSS State node01 |OFFLINE|
VxSS State node02 |OFFLINE|
VxSS State node03 |ONLINE|
VxSS State node04 |ONLINE|
```

- 9 Perform this step on the nodes (node01 and node02) in the first subcluster if the cluster uses I/O Fencing. Use an editor of your choice and change the following:

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `scsi3` to `disabled`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=disabled
```

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `SCSI3` to `NONE`. You want the line in the `main.cf` file to resemble:

```
UseFence = NONE
```

- 10 Back up the llttab, llthosts, gabtab, types.cf, main.cf and AT configuration files on the first subcluster.

```
# cp /etc/llttab /etc/llttab.bkp
# cp /etc/llthosts /etc/llthosts.bkp
# cp /etc/gabtab /etc/gabtab.bkp
# cp /etc/VRTSvcs/conf/config/main.cf \
    /etc/VRTSvcs/conf/config/main.cf.bkp
# cp /etc/VRTSvcs/conf/config/types.cf \
    /etc/VRTSvcs/conf/config/types.cf.bkp
# /opt/VRTSat/bin/vssat showbackuplist
B|/var/VRTSat/.VRTSat/profile/VRTSatlocal.conf
B|/var/VRTSat/.VRTSat/profile/certstore
B|/var/VRTSat/ABAuthSource
B|/etc/vx/vss/VRTSat.conf
Quiescing ...
Snapshot Directory :/var/VRTSatSnapShot
```

Upgrading the first subcluster

You now navigate to the installer program and start it.

To start the installer for the phased upgrade

- 1 Confirm that you are logged on as the superuser and you mounted the product image or disc.
- 2 Make sure that you can ssh or from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 3 Navigate to the folder that contains installrp.
- 4 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
```

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

- 5 Review the output as the script displays the results of the check and saves the results of the check in a log file.

Make sure all your nodes meet the prerequisites, and all the issues reported by above pre-check process have been resolved. See [“Prerequisites for upgrading to 5.0 MP3 RP5”](#) on page 88.

- 6 Start the `installrp` program, specify the nodes in the first subcluster (`node1` and `node2`).

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

The program starts with a copyright message and specifies the directory where it creates the logs.

- 7 The installer performs a series of checks and tests to ensure communications and compatibility.
- 8 When you are prompted, reply **y** to continue with the upgrade.

```
Are you sure you want to install 5.0MP3RP5? [y,n,q] (y)
```

The upgrade is finished on the first subcluster. Do not reboot the nodes in the first subcluster until you complete the [Preparing the second subcluster](#) procedure.

Preparing the second subcluster

Perform the following steps on the second subcluster before rebooting nodes in the first subcluster.

2 Unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -k
```

Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on
/dev/hd4	20971520	8570080	60%	35736	2%	/
/dev/hd2	5242880	2284528	57%	55673	9%	/usr
/dev/hd9var	4194304	3562332	16%	5877	1%	/var
/dev/hd3	6291456	6283832	1%	146	1%	/tmp
/dev/hd1	262144	261408	1%	62	1%	/home
/dev/hd11admin	262144	184408	30%	6	1%	/admin
/proc	-	-	-	-	-	/proc
/dev/hd10opt	20971520	5799208	73%	65760	5%	/opt
/dev/vx/dsk/dg2/dg2vol1		10240	7600	26%		4 1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2		10240	7600	26%		4 1% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3		10240	7600	26%		4 1% /mnt/dg2/dg2vol3

```
# umount /mnt/dg2/dg2vol1
```

```
# umount /mnt/dg2/dg2vol2
```

```
# umount /mnt/dg2/dg2vol3
```

3 Stop all VxVM volumes (for each disk group) that VCS does not manage.**4** Make the configuration writable on the second subcluster.

```
# haconf -makerw
```

5 Unfreeze the service groups.

```
# hagrps -unfreeze parallel_sg1 -persistent
```

```
# hagrps -unfreeze parallel_sg2 -persistent
```

```
# hagrps -unfreeze failover_sg1 -persistent
```

```
# hagrps -unfreeze failover_sg2 -persistent
```

```
# hagrps -unfreeze VxSS -persistent
```

6 Dump the configuration and make it read-only.

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

7 Take the service groups offline on node03 and node04.

```
# hagrpl -offline parallel_sg1 -sys node03
# hagrpl -offline parallel_sg1 -sys node04
# hagrpl -offline parallel_sg2 -sys node03
# hagrpl -offline parallel_sg2 -sys node04
# hagrpl -offline failover_sg1 -sys node03
# hagrpl -offline failover_sg2 -sys node04
# hagrpl -offline VxSS -sys node03
# hagrpl -offline VxSS -sys node04
```

8 Verify the state of the service groups.

```
# hagrpl -state
#Group      Attribute  System  Value
parallel_sg1 State      node01  |OFFLINE|
parallel_sg1 State      node02  |OFFLINE|
parallel_sg1 State      node03  |OFFLINE|
parallel_sg1 State      node04  |OFFLINE|
parallel_sg2 State      node01  |OFFLINE|
parallel_sg2 State      node02  |OFFLINE|
parallel_sg2 State      node03  |OFFLINE|
parallel_sg2 State      node04  |OFFLINE|
failover_sg1 State      node01  |OFFLINE|
failover_sg1 State      node02  |OFFLINE|
failover_sg1 State      node03  |OFFLINE|
failover_sg1 State      node04  |OFFLINE|
VxSS        State      node01  |OFFLINE|
VxSS        State      node02  |OFFLINE|
VxSS        State      node03  |OFFLINE|
VxSS        State      node04  |OFFLINE|
```

9 Perform this step on node03 and node04 if the cluster uses I/O Fencing. Use an editor of your choice and change the following:

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `scsi3` to `disabled`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=disabled
```

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `SCSI3` to `NONE`. You want the line in the `main.cf` file to resemble:

```
UseFence = NONE
```

10 Stop VCS, I/O Fencing, GAB, and LLT on node03 and node04.

```
# /opt/VRTSvcs/bin/hastop -local
# /etc/init.d/vxfen.rc stop
# /etc/init.d/gab.rc stop
# /etc/init.d/llt.rc stop
```

11 Make sure that the VXFEN, GAB, and LLT modules on node03 and node04 not loaded.

```
# /sbin/vxfenconfig -l
VXFEN vxfenconfig ERROR V-11-2-1087 There are 0 active
coordination points for this node

# /sbin/gabconfig -l
GAB Driver Configuration
Driver state           : Unconfigured
Partition arbitration: Disabled
Control port seed     : Disabled
Halt on process death: Disabled
Missed heartbeat halt: Disabled
Halt on rejoin        : Disabled
Keep on killing        : Disabled
Quorum flag           : Disabled
Restart               : Disabled
Node count            : 0
Disk HB interval (ms): 1000
Disk HB miss count    : 4
IOFENCE timeout (ms) : 15000
Stable timeout (ms)  : 5000

# /usr/sbin/strload -q -d /usr/lib/drivers/pse/llt
/usr/lib/drivers/pse/llt: no
```

Activating the first subcluster

Get the first subcluster ready for the service groups.

To activate the first subcluster

- 1 Perform this step on node01 and node02 if the cluster uses I/O Fencing. Use an editor of your choice and revert the following to an enabled state before you reboot the first subcluster's nodes:

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `NONE` to `SCSI3`. You want the line in the `main.cf` file to resemble:

```
UseFence = SCSI3
```

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `disabled` to `scsi3`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=scsi3
```

- 2 Reboot node01 and node02 in the first subcluster.

```
# /usr/sbin/shutdown -r
```

Wait until the nodes come up.

- 3 Seed node01 and node02 in the first subcluster.

```
# gabconfig -xc
```

- 4 Start VCS on node01 and node02. On each node run:

```
# hstart
```

- 5 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

- 6 Unfreeze the nodes in the first subcluster.

```
# hasys -unfreeze -persistent node01
```

```
# hasys -unfreeze -persistent node02
```

- 7 Dump the configuration and make it read-only.

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

- 8 Bring the service groups online on node01 and node02.

```
# hagr -online parallel_sg1 -sys node01
# hagr -online parallel_sg1 -sys node02
# hagr -online parallel_sg2 -sys node01
# hagr -online parallel_sg2 -sys node02
# hagr -online failover_sg1 -sys node01
# hagr -online failover_sg2 -sys node02
# hagr -online VxSS -sys node01
# hagr -online VxSS -sys node02
```

Upgrading the second subcluster

Perform the following procedure to upgrade the second subcluster (node03 and node04).

To start the installer to upgrade the second subcluster

- 1 Confirm that you are logged on as the superuser and you mounted the product image or disc.
- 2 Navigate to the folder that contains installrp.
- 3 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
```

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

- 4 Review the output as the script displays the results of the check and saves the results of the check in a log file.

Make sure all your nodes meet the prerequisites, and all the issues reported by above pre-check process have been resolved. See [“Prerequisites for upgrading to 5.0 MP3 RP5”](#) on page 88.

- 5 Confirm that VCS is stopped on node03 and node04. Start the installrp program, specify the nodes in the second subcluster (node3 and node4).

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

The program starts with a copyright message and specifies the directory where it creates the logs.

- 6 The installer performs a series of checks and tests to ensure communications and compatibility.
- 7 When you are prompted, reply **y** to continue with the upgrade.

```
Are you sure you want to install 5.0MP3RP5? [y,n,q] (y)
```
- 8 Monitor the installer program answering questions as appropriate until the upgrade completes.

Finishing the phased upgrade

You now have to reboot the nodes in the second subcluster.

To finish the upgrade

- 1 Perform this step on node03 and node04 if the cluster uses I/O Fencing. Use an editor of your choice and revert the following to an enabled state before you reboot the second subcluster's nodes:
 - In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from disabled to `scsi3`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=scsi3
```

- 2 Reboot the node03 and node04 in the second subcluster.

```
# /usr/sbin/shutdown -r
```

The nodes in the second subcluster join the nodes in the first subcluster.

- 3 # **gabconfig -a**

```
GAB Port Memberships
```

```
=====  
Port a gen  nxxxxnn membership 0123  
Port b gen  nxxxxnn membership 0123  
Port h gen  nxxxxnn membership 0123
```

4 Run an `hastatus -sum` command to determine the status of the nodes, service groups, and cluster.

```
# hastatus -sum

-- SYSTEM STATE
-- System          State          Frozen

A node01          RUNNING        0
A node02          RUNNING        0
A node03          RUNNING        0
A node04          RUNNING        0

-- GROUP STATE
-- Group          System        Probed    AutoDisabled    State

B VxSS           node01        Y         N                ONLINE
B VxSS           node02        Y         N                ONLINE
B VxSS           node03        Y         N                ONLINE
B VxSS           node04        Y         N                ONLINE
B parallel_sg1   node01        Y         N                ONLINE
B parallel_sg1   node02        Y         N                ONLINE
B parallel_sg1   node03        Y         N                ONLINE
B parallel_sg1   node04        Y         N                ONLINE
B parallel_sg2   node01        Y         N                ONLINE
B parallel_sg2   node02        Y         N                ONLINE
B parallel_sg2   node03        Y         N                ONLINE
B parallel_sg2   node04        Y         N                ONLINE
B failover_sg1   node01        Y         N                ONLINE
B failover_sg1   node02        Y         N                OFFLINE
B failover_sg1   node03        Y         N                OFFLINE
B failover_sg1   node04        Y         N                OFFLINE
B failover_sg2   node01        Y         N                OFFLINE
B failover_sg2   node02        Y         N                ONLINE
B failover_sg2   node03        Y         N                OFFLINE
B failover_sg2   node04        Y         N                OFFLINE
```

5 After the upgrade is complete, mount the VxFS file systems and start the VxVM volumes (for each disk group) that VCS does not manage.

In this example, you have performed a phased upgrade of Veritas product. The service groups were down when you took them offline on node03 and node04, to the time Veritas product brought them online on node01 or node02.

Performing a phased upgrade to 5.0 MP3 RP5 on a Storage Foundation HA cluster

The following procedure describes performing a phased upgrade on a Storage Foundation HA, SF for Oracle HA or SF for DB2 HA cluster.

To perform a phased upgrade to 5.0 MP3 RP5 on a Storage Foundation HA 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 failover service groups from group A to the nodes in group B.

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

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

```
# haconf -makerw
```

- 5 Freeze the high availability 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 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

- 8 Unmount all Storage Checkpoints and file systems:

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

- 9 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 `vxlink status` command to verify that all RLINKs are up-to-date:

```
# vxlink -g diskgroup status rlink_name
```

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

- 10 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.
- 11 Stop all VxVM volumes. Enter the following command for each disk group:

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

Verify that no volumes remain open:

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

- 12 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
```

- 13 Close any instances of the VCS GUI that is running on the node.
- 14 If required, apply the operating system kernel patches on the nodes in the selected group.
- 15 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 script proceeds in a noninteractive mode to examine the systems for licenses, filesets, disk space, system-to-system communications, etc.

- 16** Review the output as the script displays the results of the check and saves the results of the check in a log file. Make sure all your nodes meet the prerequisites, and all the issues reported by above pre-check process have been resolved. See [“Prerequisites for upgrading to 5.0 MP3 RP5”](#) on page 88.

- 17** Start the upgrade using the installrp script.

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

- 18** Take the service groups offline on group B.

```
# hagrps -offline service_group -any
```

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

- Stop VCS on each node in the selected group:

```
# hastop -local
```

- Stop the VCS command server:

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

Where `pid_of_CmdServer` is the process ID of `CmdServer`.

- Stop cluster fencing, GAB, and LLT.

```
# /etc/rc.d/rc2.d/S97vxfen stop  
# /etc/rc.d/rc2.d/S92gab stop  
# /etc/rc.d/rc2.d/S70llt stop
```

- 20** Reboot the nodes you have upgraded, after the nodes come up, on the nodes that you have rebooted, seed the nodes.

```
# gabconfig -xc
```

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

```
# haconf -makerw
```

- 22** Unfreeze the nodes in the first sub-cluster:

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

23 Make the VCS configuration read-only:

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

24 Online service groups to the original node:

```
# hagrpr -online service_group -sys nodename
```

25 Repeat step 14 to step 17 to upgrade the nodes in group B, then use the following command to reboot them:

```
/usr/sbin/shutdown -r
```

The nodes in the group B join the nodes in the first subcluster.

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

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

27 If you stopped any RVGs in step 9, restart each RVG:

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

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

```
# mount /filesystem
```

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

```
# mount /checkpoint_name
```

Performing a phased upgrade to 5.0 MP3 RP5 on a Storage Foundation Cluster File System cluster

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

When you perform a phased upgrade, you upgrade some of the nodes (group A) in the cluster first. You then upgrade the rest of the nodes (group B) in the cluster.

To perform a phased upgrade to 5.0 MP3 RP5 on an SFCFS cluster:

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

- 3 If you have a failover service group on a node in group A, switch the service group to a node that is running in group B.

```
# 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 in group A:

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

- 6 Make the configuration read-only:

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

- 7 Upgrade the nodes in group A first. Perform step 8 through step 33 for these nodes.

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

```
# mount | grep vxfs
```

If any Storage Checkpoints are mounted, on each node in the 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:

```
# mount | 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.
- Use the vxrvrg stop command to stop each RVG individually:

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

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

```
# vxlink -g diskgroup status rlink_name
```

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

11 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.

12 Verify that no volumes remain open, enter the following command:

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

13 Stop VCS on each node in the selected group:

```
# hastop -local
```

14 Stop the VCS command server:

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

where pid_of_CmdServer is the process ID of CmdServer.

15 If ODM is installed (port 'd' is up), stop the ODM service:

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

16 Stop cluster fencing, GAB, and LLT:

```
# /etc/rc.d/rc2.d/S97vxfen stop  
# /etc/rc.d/rc2.d/S92gab stop  
# /etc/rc.d/rc2.d/S70llt stop
```

17 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
```

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

See “[System requirements](#)” on page 8.and IBM’s documentation for the procedures.

19 Mount the 5.0 MP3 RP5 product image or disc and navigate to the folder that contains the installation script. On the first node of the selected group, 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 in the selected group that you want to upgrade.

20 After all of the nodes in the selected group are upgraded, on the second group of nodes (group B), stop the failover service groups. Downtime starts for failover service groups.

```
# hagr -offline failover_service_group
```

21 Repeat step 8 through step 18 on each node in the second group (group B).

Note: If HA is not stopped then after rebooting the first set of upgraded nodes, those will try to join the cluster and it will fail to join.

22 Shut down and reboot each of the nodes in the selected group that you have upgraded. After the nodes come back up, application failover capability is available for that group.

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

- 24** Make the VCS configuration writable again from any node in the selected group:
- ```
haconf -makerw
```
- 25** Enter the following command on each node in the selected group to unfreeze HA service group operations:
- ```
# hasys -unfreeze -persistent nodename
```
- 26** Make the configuration read-only:
- ```
haconf -dump -makero
```
- 27** Run the hagrps autoenable command for the failover service groups on each node in the selected group:
- ```
# hagrps -autoenable service_group -sys nodename
```
- 28** Bring the failover service groups online on the nodes in the selected group:
- ```
hagrps -online service_group -sys nodename
```
- 29** Bring the CVM service group online on each node in the selected group A:
- ```
# hagrps -online cvm -sys nodename
```
- 30** If required start all the volumes by entering the following command for each disk group:
- ```
vxvol -g diskgroup startall
```
- 31** If you stopped any RVGs in step 10, restart each RVG:
- ```
# vxrvrg -g diskgroup start rvg_name
```
- 32** Remount all VxFS file systems on all nodes in the selected group:
- ```
mount /filesystem
```
- 33** Remount all Storage Checkpoints on all nodes in the selected group:
- ```
# mount /checkpoint_name
```
- 34** Repeat step 19 and step 22 through step 33 for the second group (group B) of nodes.

Performing a phased upgrade to 5.0 MP3 RP5 on a Storage Foundation for Oracle RAC cluster

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

The phased upgrade methodology involves upgrading half of the nodes (group A and group B) in the cluster at a time.

To upgrade to 5.0 MP3 RP5 on an SF Oracle RAC cluster

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so that you can execute all product commands.
- 3 Switch the failover service groups from group A to the nodes in group B.

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

- 4 On the nodes in group A, freeze HA service group operations:

```
# haconf -makerw  
# hasys -freeze -persistent nodename  
# haconf -dump -makero
```

- 5 On the nodes in group A, stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 6 On the nodes in group A, stop Oracle database:

- If VCS manages the Oracle RAC instance:

```
# hagrps -offline oracle_group -sys nodename
```

- If VCS does not manage the Oracle RAC instance, log in as the Oracle user on one of the nodes in the group A and shut down the instances on nodes in group A:

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

- 7 On nodes in Group A, disable autostart of Oracle database instances (you need to disable autostart of the database instances to relink Oracle after installing this RP).
 - If VCS manages the Oracle RAC instance:

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

- If VCS does not manage the Oracle RAC instance, log in as the Oracle user on one of the nodes in the group A and change the startup policy to manual:

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

8 On the nodes in group A, unmount the VxFS and CFS file systems that are not managed by VCS.

- Make sure that no processes are running which 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
```

- Unmount the VxFS or CFS file system:

```
# umount /mount_point
```

9 On the nodes in group A, 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
```

10 On all nodes in group A, stop VCS.

- If CRS is under VCS control:

```
# hastop -local
```

- If CRS is not under VCS control, you need to stop CRS on all nodes in group A before you stop VCS:

```
# $CRS_HOME/bin/crsctl stop crs
# hastop -local
```

- 11** On any of the node in group A, mount the product image or disc and navigate to the folder that contains the installation program.

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

If passwordless ssh is not configured then enter:

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

where node1 and node2 are nodes in group A.

- 12** After installrp is finished, move /etc/llttab to /etc/llttab.bak on all nodes in group A. LLT should not start after systems in group A reboot, to avoid these nodes to join Group B systems having old version of LLT.

```
# mv /etc/llttab /etc/llttab.bak
```

- 13** Restart the nodes in group A.

```
# shutdown -r now
```

- 14** After nodes in group A comes back online, on the nodes in group B, stop all applications that are not configured under VCS. Use native application commands to stop the application.

Downtime starts here.

- 15** On the nodes in group B, stop Oracle database:

- If VCS manages the Oracle RAC instance:

```
# hagrps -offline oracle_group -sys nodename
```

- If VCS does not manage the Oracle RAC instance, log in as the Oracle user on one of the nodes in the group B and shut down the instances on nodes in group B:

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

- 16** Perform step 8 to step 10 on the nodes in group B.

- 17** On the nodes in group B, stop the remaining SF Oracle RAC processes. Run these commands from every node in group B:

- Stop the VCS command server if it is running:

```
# ps -ef | grep CmdServer  
# /opt/VRTSvcs/bin/CmdServer -stop
```

- Stop ODM, vcsmm, fencing, GAB, LMX, and LLT.

```
# /etc/rc.d/rc2.d/S99odm stop
# /etc/rc.d/rc2.d/S98vcsmm stop
# /etc/rc.d/rc2.d/S97vxfen stop
# /etc/rc.d/rc2.d/S92gab stop
# /etc/rc.d/rc2.d/S71lmx stop
# /etc/rc.d/rc2.d/S70llt stop
```

- Sometimes LLT repeats an error.

```
# /etc/rc.d/rc2.d/S70llt stop
Unconfiguring LLT...
LLT lltconfig ERROR V-14-2-15121 LLT unconfigure aborted,
unregister 2 port(s)
Unloading LLT Driver...strload: cannot terminate /usr/lib
/drivers/pse/llt: Device
busy
```

You might need to stop the lmx kernel extensions, before you stop LLT.

```
# /etc/methods/lmxext -stop
```

- Engine logs in Group A nodes show the following errors, which you can ignore.

```
2010/05/24 07:46:38 VCS WARNING V-16-1-11030 HAD not ready
to receive this command. Message was: 0x90a
2010/05/24 07:46:38 VCS INFO V-16-1-10125 GAB timeout set to
30000 ms
2010/05/24 07:46:38 VCS ERROR V-16-1-10116 GabHandle::open
failed errno = 161
2010/05/24 07:46:38 VCS ERROR V-16-1-11033 GAB open failed.
Exiting
```

- 18** After all the ports in Group B are stopped, start the ports in Group A. Note that you need to stop ODM because it starts in single instance mode.

```
single instance mode.  
# mv /etc/llttab.bak /etc/llttab  
# /etc/rc.d/rc2.d/S7011t start  
# /etc/rc.d/rc2.d/S711mx start  
# /etc/rc.d/rc2.d/S92gab start  
# gabconfig -x  
# /etc/rc.d/rc2.d/S97vxfen start  
# /etc/rc.d/rc2.d/S98vcsmm start  
# /etc/rc.d/rc2.d/S99odm stop  
# /etc/rc.d/rc2.d/S99odm start  
# hastart
```

After you perform these steps, the GAB ports a, b, d, o and h are online.

```
# gabconfig -a  
GAB Port Memberships  
=====
```

Port a	gen	bb9401	membership	01
Port b	gen	bb9403	membership	01
Port d	gen	bb9407	membership	01
Port h	gen	bb9409	membership	01
Port o	gen	bb9405	membership	01

- 19** On the nodes in group A, unfreeze the HA service group operations:

```
# haconf -makerw  
# hasys -unfreeze -persistent nodename  
# haconf -dump -makero
```

- 20** On one of the nodes in group A, start the cvm service group:

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

After starting the cvm group, the GAB ports v, w, and f come online. CRS comes online on Group A nodes, but database instance remains offline.

- 21** On the nodes in group A, relink the SF Oracle RAC libraries with Oracle.

Refer to Veritas Storage Foundation for Oracle RAC 5.0 MP3 Installation and Configuration Guide for more information.

- 22** On the nodes in group A, start the remaining VCS service groups:

- For parallel groups:

```
# hagrps -online group_name -sys nodename
```

- For failover groups:

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

23 On the nodes in group A, start the Oracle database:

- If VCS manages the Oracle RAC instance, perform the following command:

```
# hagrps -online oracle_group -sys nodename
```

- If VCS does not manage the Oracle RAC instance, log in as the Oracle user on one of the nodes in the group A and start the instances on nodes in group A:

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

Downtime ends here.

- 24** On the nodes in group A, manually mount the VxFS or CFS file systems that VCS does not manage.
- 25** On the nodes in group A, start all applications that VCS does not manage. Use native application commands to start the applications.
- 26** On the nodes in group B, upgrade SF Oracle RAC:

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

If passwordless ssh is not configured then enter:

```
# ./installrp -rsh node3 node4 ...
```

where node3 and node4 are nodes in group B.

27 Restart the nodes in group B.

```
# shutdown -r now
```

When the nodes in group B come up, the GAB ports a, b, d, f, h, o, v, and w are online. All service groups should be online except the oracle_group service group.

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen    bb9402 membership 0123
Port b gen    bb9404 membership 0123
Port d gen    bb9412 membership 0123
Port f gen    bb9410 membership 0123
Port h gen    bb940a membership 0123
Port o gen    bb9406 membership 0123
Port v gen    bb940c membership 0123
Port w gen    bb940e membership 0123
```

28 On the nodes in group B, relink the SF Oracle RAC libraries with Oracle.

Refer to Veritas Storage Foundation for Oracle RAC 5.0 MP3 Installation and Configuration Guide for more information.

29 On the nodes in group B, start the Oracle database:

- If VCS manages the Oracle RAC instance:

```
# hagrpl -online oracle_group -sys nodename
```

- Return the startup policy to AutoStart:

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

- If VCS does not manage the Oracle RAC instance, log in as the Oracle user on one of the nodes in the group A and start the instances:

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

Change startup policy back to automatic:

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

30 Perform step 24 and step 25 on the nodes in group B.

- 31 If you are currently using backup and restore for the DBED repository. Perform a full backup of the DBED repository database after completing the 5.0 MP3 RP5 installation.

For more information about older backups failing to be restored using the DBED scripts, see “Software limitations” on page 98.

For the `sfua_rept_adm` command, refer to the Veritas Storage Foundation for Oracle Administrator's Guide.

For more information See “[Storage Foundation for Oracle fixed issues](#)” on page 42. or [Storage Foundation for DB2 fixed issues](#) for incident 1425261.

- 32 If you are going to use the DBED feature, run the `sfua_db_config -o setperm` command to set the correct permissions, owner, and group for the following directories:

```
/var/vx/vxdba  
/var/vx/vxdba/logs  
/var/vx/vxdba/locks
```

You must perform this step to ensure that DBED features work.

Performing a full upgrade to 5.0 MP3 RP5 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:

- 1 Freeze service group operations and stop VCS on the cluster.
- 2 Take the nodes offline and install the software patches.
- 3 Bring the nodes online to restart cluster failover services.

Depending on your cluster's configuration, select one of the following procedures to upgrade to 5.0 MP3 RP5:

- [Performing a full upgrade to 5.0 MP3 RP5 on a VCS cluster](#)
- [Performing a full upgrade to 5.0 MP3 RP5 on a Storage Foundation HA cluster](#)
- [Performing a full upgrade to 5.0 MP3 RP5 on a Storage Foundation Cluster File System cluster](#)
- [Performing a full upgrade to 5.0 MP3 RP5 on a Storage Foundation for Oracle RAC cluster](#)

Performing a full upgrade to 5.0 MP3 RP5 on a VCS cluster

Note: If you have any volumes under VxVM and VxFS you need to stop them before upgrading which means your application will require downtime.

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

To perform pre-upgrade tasks:

- 1 Log in as superuser.
- 2 Verify that /opt/VRTS/bin is in your PATH so that you can execute all product commands.
- 3 List the service groups in your cluster and their status. On any node, type:

```
# hagr -state
```

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

```
# hagr -offline -force ClusterService -any
```

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

```
# haconf -makerw
```

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

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

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

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

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

- 8 Make a backup copy of the current main.cf and all types.cf configuration files. For example, on one node in the cluster, type:

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

To perform the pre-check:

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

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

See “Prerequisites for upgrading to 5.0 MP3 RP5” on page 153.

To perform the upgrade:

- 1 Navigate to the folder that contains the installrp script and start the installrp script:

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

Review the output.

- 2 Reboot per the installer instructions.

To perform a service restart or cluster reboot:

- 1 Reboot all nodes in the cluster.
- 2 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 -list | sort -u +0b -1 | \
  while read grp sys ; do
    hagrps -unfreeze $grp -persistent
  done
# haconf -dump -makero
```

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

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

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

where system is the node name.

Performing a full upgrade to 5.0 MP3 RP5 on a Storage Foundation HA cluster

The following procedure describes performing a full upgrade on a Storage Foundation HA, SF for Oracle HA or SF for DB2 HA cluster.

To perform a full upgrade to 5.0 MP3 RP5 on a Storage Foundation HA 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** On each node, enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

- 7** Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# 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 vxrvrg stop command to stop each RVG individually:

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

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

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

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

- 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** 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 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

12 Navigate to the folder that contains the installrp script and check the readiness of the systems where you plan to upgrade. The command to start the pre-upgrade check is:

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

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

13 Review the output as the script displays the results of the check and saves the results of the check in a log file.

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

See “Prerequisites for upgrading to 5.0 MP3 RP5” on page 153.

15 Navigate to the folder that contains the installrp script and start the installrp script:

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

Review the output.

16 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.

17 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

18 Unfreeze the service group operations on each node:

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

19 Make the VCS configuration read-only:

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

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

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

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

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

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

```
# mount /checkpoint_name
```

- 23** If you are currently using backup and restore for the DBED repository. Perform a full backup of the DBED repository database after completing the 5.0 MP3 RP5 installation.

For more information See “[Software limitations](#)” on page 72. about older backups failing to be restored using the DBED scripts.

See the *Veritas Storage Foundation for Oracle Administrator's Guide* for the `sfua_rept_adm` command.

For more information See “[Storage Foundation for Oracle fixed issues](#)” on page 42. or See “[Storage Foundation for DB2 fixed issues](#)” on page 44. on page 52 for incident 1425261.

- 24** If you are going to use the DBED feature, run the `sfua_db_config -o setperm` command to set the correct permissions, owner, and group for the following directories:

```
/var/vx/vxdba  
/var/vx/vxdba/logs  
/var/vx/vxdba/locks
```

Note: If you do not perform this step the DBED features will not work.

Performing a full upgrade to 5.0 MP3 RP5 on a Storage Foundation Cluster File System cluster

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

To perform a full upgrade to 5.0 MP3 RP5 on an SFCFS cluster

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

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

```
# haconf -makerw
```

- 4** Enter the following command to freeze the 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:

```
# mount | grep vxfs
```

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

```
# umount /filesystem
```

- 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 If ODM is installed and port 'd' is up. Stop ODM service using the following command on each node:

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

14 On each node, stop ODM, cluster fencing, GAB, and LLT :

```
# /etc/rc.d/rc2.d/S97vxfen stop  
# /etc/rc.d/rc2.d/S92gab stop  
# /etc/rc.d/rc2.d/S70llt stop
```

15 If required, apply the OS kernel patches.

See [“System requirements”](#) on page 8. and Oracle's documentation for the procedures.

- 16** 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
```

- 17** Mount the 5.0 MP3 RP5 product image or disc and navigate to the folder that contains the installation script. 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.

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

- 19** If necessary, reinstate any missing mount points in the /etc/filesystems 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 5.0 MP3 RP5 on a Storage Foundation for Oracle RAC cluster

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

To upgrade to 5.0 MP3 RP5 on an SFRAC cluster

- 1** Log in as superuser.
- 2** Verify that `/opt/VRTS/bin` is in your `PATH` so that you can execute all product commands.
- 3** From a node in the cluster, freeze HA service group operations:

```
# haconf -makerw  
# hasys -freeze -persistent nodename  
# haconf -dump -makero
```

(Old step 6). Stop Oracle database on the cluster: If the Oracle RAC instance is managed by VCS:

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

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

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

- 4** Stop all applications on the cluster that are not configured under VCS. Use native application's commands to stop them.
- 5** Stop the Oracle database on all the nodes in the cluster:
 - If VCS manages the Oracle RAC instance:

```
# hagrps -offline oracle_group -any
```

- If VCS does not manage the Oracle RAC instance, log in as the Oracle user on one of the nodes in the cluster and shut down the database:

```
$ srvctl stop database -d database_name
```

- 6 Disable AutoStart for the Oracle database. You need to disable AutoStart to relink Oracle after you install this release patch.

- If VCS manages the Oracle RAC instance:

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

- If VCS does not manage the Oracle RAC instance, log in as the Oracle user on one of the nodes in the cluster and change the startup policy to manual:

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

- 7 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
```

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

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

Verify that no volumes remain open:

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

- 9 Stop VCS on the cluster.

- If CRS is under VCS control, stop VCS from one of the nodes in cluster:

```
# hastop -all
```

- If CRS is not under VCS control, stop CRS on all nodes in cluster, and then from one node stop VCS:

```
# $CRS_HOME/bin/crsctl stop crs
# hastop -all
```

- 10 Stop the VCS command server on the cluster:

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

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

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

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

- 12 Mount the 5.0 MP3 RP5 product image or disc and navigate to the folder that contains the installation script. Enter the `installrp` command to start the script:

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

If passwordless ssh is not configured then enter:

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

where `node1` and `node2` are nodes in the SFRAC cluster.

- 13 Restart the nodes:

```
# shutdown -r now
```

- 14 From a node in the cluster, unfreeze the HA service group operations:

```
# haconf -makerw
# hasys -unfreeze -persistent nodename
# haconf -dump -makero
```

- 15** From a node in the cluster, manually start the cvm service group:

```
# hagrps -online cvm -any
```

After starting the cvm group, the GAB ports v, w, and f are online. CRS comes online on the cluster nodes if it is under VCS control.

- 16** Relink the SF Oracle RAC libraries with Oracle. Refer to Veritas Storage Foundation for Oracle RAC 5.0 MP3 Installation and Configuration Guide for more information.

(old step 17) Start VCS on each of the nodes:

```
# hastart
```

- 17** Bring the VCS service groups online:

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

Downtime ends for the failover service groups

(Was step 20ish) If VCS does not control CRS, use the following command on each node to start CRS.

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

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

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

- 18** From a node in the cluster, start the Oracle database and return the startup policy to Autostart:

- If VCS manages the Oracle RAC instance:

```
# hagrps -online oracle_group -any  
# haconf -makerw  
# hagrps -modify oracle_group AutoStart 1  
# haconf -dump -makero
```

- If VCS does not manage the Oracle RAC databases, log in as the Oracle user on one of the nodes in the cluster and start the database:

```
$ srvctl start database -d database_name
```

Change the startup policy back to automatic:

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

- 19 Start all applications on the cluster that are not configured under VCS. Use native application commands to start the application.
- 20 If you are currently using backup and restore for the DBED repository. Perform a full backup of the DBED repository database after completing the 5.0 MP3 RP5 installation.

For more information See “[Software limitations](#)” on page 72. about older backups failing to be restored using the DBED scripts.

See the *Veritas Storage Foundation for Oracle Administrator's Guide* for the `sfua_rept_adm` command.

For more information See “[Storage Foundation for Oracle fixed issues](#)” on page 42. or See “[Storage Foundation for DB2 fixed issues](#)” on page 44. for incident 1425261.

- 21 If you are going to use the DBED feature, run the `sfua_db_config -o setperm` command to set the correct permissions, owner, and group for the following directories:

```
/var/vx/vxdba  
/var/vx/vxdba/logs  
/var/vx/vxdba/locks
```

Note: If you do not perform this step the DBED features will not work.

Upgrading to 5.0 MP3 RP5 on a standalone system

You can use this procedure to upgrade on a standalone system that runs Storage Foundation, SF for Oracle, or SF for DB2.

To upgrade to 5.0 MP3 RP5 on a standalone system

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so that you can execute all product commands.
- 3 If required, apply the OS kernel patches.
See “[System requirements](#)” on page 8. and IBM’s documentation for the procedures.
- 4 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

5 Unmount all Storage Checkpoints and file systems:

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

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

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

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

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

- On the Primary node, use the 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
```

Verify that no volumes remain open:

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

9 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

- 10 Mount the 5.0 MP3 RP5 product image or disc and navigate to the folder that contains the installation script. Run the installrp script.

If ssh key authentication is configured then enter:

```
# ./installrp node1
```

If ssh is not configured then enter:

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

where node1 is the node that you plan to upgrade.

- 11 Shut down and restart the system.
- 12 If necessary, reinstate any missing mount points in the/etc/filesystems file.
- 13 Restart all the volumes by entering the following command for each disk group:

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

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

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

- 15 Remount all VxFS file systems and Storage Checkpoints:

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

- 16 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

- 17** If you are currently using backup and restore for the DBED repository. Perform a full backup of the DBED repository database after completing the 5.0 MP3 RP5 installation.

For more information See [“Software limitations”](#) on page 72. about older backups failing to be restored using the DBED scripts.

See the *Veritas Storage Foundation for Oracle Administrator’s Guide* for the `sfua_rept_adm` command.

For more information See [“Storage Foundation for Oracle fixed issues”](#) on page 42. or See [“Storage Foundation for DB2 fixed issues”](#) on page 44. for incident 1425261.

- 18** If you are going to use the DBED feature, run the `sfua_db_config -o setperm` command to set the correct permissions, owner, and group for the following directories:

```
/var/vx/vxdba  
/var/vx/vxdba/logs  
/var/vx/vxdba/locks
```

Note: If you do not perform this step the DBED features will not work.

Verifying software versions

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

```
# ls1pp -L VRTS\*
```

Removing 5.0 MP3 RP5

Roll back of the 5.0 MP3 RP5 to the release 5.0 MP3 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 5.0 MP3 software.

You can roll back 5.0 MP3 RP5 to the release 5.0 MP3 version for Veritas Cluster Server.

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

- [Rolling back 5.0 MP3 RP5 to 5.0 MP3 for Veritas Cluster Server](#)
- [Removing 5.0 MP3 RP5 on Storage Foundation or Storage Foundation Cluster File System](#)
- [Removing 5.0 MP3 RP5 on Storage Foundation for Oracle RAC](#)

Rolling back 5.0 MP3 RP5 to 5.0 MP3 for Veritas Cluster Server

Use the following procedure to roll back VCS 5.0 MP3 RP5 to VCS 5.0 MP3 on your cluster manually. To uninstall VCS, see the Veritas Cluster Server Installation Guide.

Caution: Use this procedure only when rolling back VCS. Do not roll back VCS when it is part of other products that rely on VCS, for example Storage Foundation Clustered File System or Storage Foundation for Oracle RAC.

To roll back 5.0 MP3 RP5:

- 1 Verify that all of the VCS 5.0 MP3 RP5 patches are in the APPLIED state. Create a text file called `filesets.to.reject` that contains the name and version of each fileset, one per line, exactly as shown below.

```
VRTSgab.rte           5.0.3.500
VRTSllt.rte           5.0.3.500
VRTSvc.s.rte          5.0.3.500
VRTSvc.sag.rte        5.0.3.500
VRTSvc.sdb.rte        5.0.3.500
VRTSvc.sor.rte        5.0.3.500
VRTSvc.ssy.rte        5.0.3.500
VRTSvcxfen.rte        5.0.3.500
```

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

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

3 Review the output and confirm that all of the updated filesets are in the APPLIED state. Example output follows:

Fileset	Level	State	Description

Path: /usr/lib/objrepos			
VRTSgab.rte	5.0.3.500	APPLIED	Veritas Group Membership and Atomic Broadcast by Symantec Veritas-5.0MP3RP5-05/04/11-05:49:14
Path: /etc/objrepos			
VRTSgab.rte	5.0.3.500	APPLIED	Veritas Group Membership and Atomic Broadcast by Symantec Veritas-5.0MP3RP5-05/04/11-05:49:14

Fileset	Level	State	Description

Path: /usr/lib/objrepos			
VRTSl1t.rte	5.0.3.500	APPLIED	Veritas Low Latency Transport by Symantec Veritas-5.0MP3RP5-05/04/11-05:49:14
Path: /etc/objrepos			
VRTSl1t.rte	5.0.3.500	APPLIED	Veritas Low Latency Transport by Symantec Veritas-5.0MP3RP5-05/04/11-05:49:14

Fileset	Level	State	Description

Path: /usr/lib/objrepos			
VRTSvc.s.rte	5.0.3.500	APPLIED	Veritas Cluster Server by Symantec Veritas-5.0MP3RP5-05/04/11-05:49:14
Path: /etc/objrepos			
VRTSvc.s.rte	5.0.3.500	APPLIED	Veritas Cluster Server by Symantec Veritas-5.0MP3RP5-05/04/11-05:49:14

Fileset	Level	State	Description

```
-----
Path: /usr/lib/objrepos
VRTSvcstag.rte          5.0.3.500  APPLIED  Veritas Cluster Server Bundled
                        Agents by Symantec
                        Veritas-5.0MP3RP5-05/04/11-05:
                        49:14

Fileset                  Level  State  Description
-----
```

```
Path: /usr/lib/objrepos
VRTSvcstb.rte          5.0.3.500  APPLIED  Veritas High Availability
                        Agent for DB2 by Symantec
                        Veritas-5.0MP3RP5-05/04/11-05:
                        49:14
```

```
Path: /etc/objrepos
VRTSvcstb.rte          5.0.3.500  APPLIED  Veritas High Availability
                        Agent for DB2 by Symantec
                        Veritas-5.0MP3RP5-05/04/11-05:
                        49:14

Fileset                  Level  State  Description
-----
```

```
Path: /usr/lib/objrepos
VRTSvcstor.rte         5.0.3.500  APPLIED  Veritas High Availability
                        Agent for Oracle by Symantec
                        Veritas-5.0MP3RP5-05/04/11-05:
                        49:14
```

```
Path: /etc/objrepos
VRTSvcstor.rte         5.0.3.500  APPLIED  Veritas High Availability
                        Agent for Oracle by Symantec
                        Veritas-5.0MP3RP5-05/04/11-05:
                        49:14

Fileset                  Level  State  Description
-----
```

```
Path: /usr/lib/objrepos
VRTSvcssy.rte          5.0.3.500  APPLIED  Veritas High Availability
                        Agent for Sybase by Symantec
                        Veritas-5.0MP3RP5-05/04/11-05:
                        49:14
```

```
Path: /etc/objrepos
VRTSvcssy.rte          5.0.3.500  APPLIED  Veritas High Availability
                        Agent for Sybase by Symantec
```

```
Veritas-5.0MP3RP5-05/04/11-05:
49:14
Fileset                Level  State  Description
-----
Path: /usr/lib/objrepos
VRTSvxfen.rte          5.0.3.500  APPLIED  Veritas I/O Fencing by
                        Symantec
                        Veritas-5.0MP3RP5-05/04/11-05:
                        49:14

Path: /etc/objrepos
VRTSvxfen.rte          5.0.3.500  APPLIED  Veritas I/O Fencing by
                        Symantec
                        Veritas-5.0MP3RP5-05/04/11-05:
                        49:14
```

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

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

```
# hagr -state
```

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

```
# hagr -offline -force ClusterService -any
```

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

```
# haconf -makerw
```

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

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

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

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

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

- 10** 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
```

- 11** Shut down VCS. On any node, type:

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

- 12** Shut down CmdServer. On each node, type:

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

- 13** Verify that VCS has shut down.

- On any node, type:

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

The output resembles:

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

Output for membership for port h does not appear.

- On each node, run the command:

```
# ps -ef | egrep "had|hashadow|CmdServer"
```

Terminate any instances of had, hashadow, or CmdServer that still run after 60 seconds.

- 14** Stop fencing, GAB, and LLT.

```
# /etc/rc.d/rc2.d/S97vxfen stop
# /etc/rc.d/rc2.d/S92gab stop
# /etc/rc.d/rc2.d/S701lt stop
```

- 15** Preview the patch removal selection and validity tests. On each node, type:

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

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

- 16** Perform the patch removal. On each node, type:

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

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

- 17** Reboot all nodes in the cluster.

- 18** 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
# hagr -list | sort -u +0b -1 | \
while read grp sys ; do
    hagr -unfreeze $grp -persistent
done
# haconf -dump -makero
```

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

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

```
# hagr -online ClusterService -sys system
```

where `system` is the node name.

Removing 5.0 MP3 RP5 on Storage Foundation or Storage Foundation Cluster File System

You can use the following procedure to uninstall 5.0 MP3 RP5 on Storage Foundation or Storage Foundation Cluster File System (SFCFS).

To uninstall 5.0 MP3 RP5 on Storage Foundation or 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 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```

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

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

```
# vxrvvg -g diskgroup stop rvg_name
```

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

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

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

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

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

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

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

- 8** Stop VCS along with all the resources and then stop the remaining resources manually:

```
# /etc/rc.d/rc2.d/S99vcs stop
```

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

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

- 10** Unmount /dev/odm on each node:

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

- 11** Unload the ODM module on each node:

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

- 12** Check if the VEA service is running:

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

- 13** If the VEA service is running, stop it:

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

- 14** To uninstall Veritas Volume Manager, it is essential to disable DMP control on the root disk if enabled. Run the following command to disable DMP support for the root disk and reboot the system.

```
# vxdmproot uninstall
```

- 15** 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, enter the following commands:

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

- Or for example, to uninstall Veritas Storage Foundation Cluster File System, enter the following commands:

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

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

Note: Provided that the remote shell (RSH) or secure shell (SSH) has been configured correctly, this command can be run on a single node of the cluster to install the software on all the nodes of the sub-cluster. After uninstalling the Veritas software, refer to the appropriate product's 5.0 MP3 Installation Guide document to reinstall the 5.0 MP3 software.

Removing 5.0 MP3 RP5 on Storage Foundation for Oracle RAC

You can use the following procedure to uninstall the 5.0 MP3 RP5 on Storage Foundation for Oracle RAC systems.

To uninstall the 5.0 MP3 RP5 on SF Oracle RAC

1 Stop Oracle instances on each node.

- If Oracle is not under VCS control, stop its instances on each node.

```
$ $ORACLE_HOME/bin/srvctl stop database -d dbname
```

- If Oracle is under VCS control, take the Oracle resource offline. Use the following command for each node:

```
# hares -offline ora_res -sys node
```

2 Stop applications that use native commands and that use CFS and CVM and are not under VCS control.

For example, if CRS is not under VCS control, use the following command to stop CRS:

```
# $CRS_HOME/bin/crsctl stop crs
```

Unmount CFS file systems on each node which are not under VCS control.

- 3 Stop VCS on all nodes. Use the following command:

```
# hastop -local
```

- 4 Stop all the applications on each node that use VxFS file system and that are not under VCS control.
- 5 Unmount VxFS file systems on each node that are not under VCS control.
- 6 Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped. In the `gabconfig -a` command output, the VCS engine or high availability daemon (HAD) port h is not displayed. This indicates that you have stopped VCS.

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

Sample output:

```
GAB Port Memberships
=====
Port a gen 5c3d0b membership 01
Port b gen 5c3d10 membership 01
Port d gen 5c3d0c membership 01
Port o gen 5c3d0f membership 01
```

- 7 Uninstall Storage Foundation for Oracle RAC.

```
# cd /opt/VRTS/install
# ./uninstallsfrac [-rsh] node1 node2
```

See the *Veritas Storage Foundation for Oracle RAC 5.0 MP3 Installation and Configuration Guide* for more information. After uninstalling the packages, refer to the *Storage Foundation for Oracle RAC 5.0 MP3 Installation and Configuration Guide* to reinstall the 5.0 MP3 software.

Documentation addendum

The following is an addition to the *Veritas Cluster Server Bundled Agents Reference Guide*.

MemCPUAllocator agent

Use the MemCPUAllocator agent to allocate CPU and memory to an IBM AIX dynamic logical partition (DLPAR). Set this resource's attribute values to specify

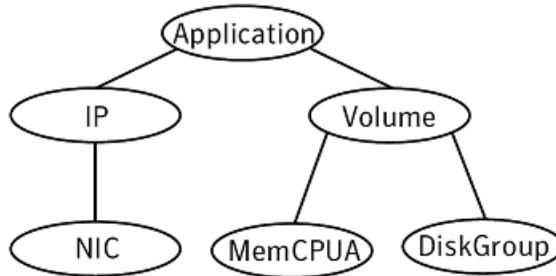
the amount of CPU and memory that you want to allocate to a service group on a DLPAR.

For prerequisites and other important information about this agent, See [“MemCPUAllocator agent notes”](#) on page 150.

Dependencies

Set the MemCPUAllocator resource as a leaf node in a resource dependency tree. Select the amount of CPU and memory that you want the DLPAR to have before it comes online.

Figure 1-1 Sample service group for a MemCPUAllocator resource, where the MemCPUA resource represents the MemCPUAllocator resource



Agent functions

Online	The MemCPUAllocator agent dynamically allocates the required amount of memory and CPU to the DLPAR from the Hardware Management Console (HMC).
Offline	The agent deallocates the amount of memory and CPU it acquired during the online agent function. It then returns the resources back to the pool.

Monitor

Checks that the online agent function succeeded. If it succeeded, then the monitor agent function reports the resource state as online. If it did not succeed, then the monitor agent function reports the resource state as offline.

If it did not succeed, then the monitor agent function reports the resource state as offline. If the agent is not able to allocate the required resources during the online agent function, the subsequent monitor reports offline and the resource faults. Because the resource is a leaf node, VCS engine stops bringing other resources online and marks the group as faulted. The VCS engine then tries to bring the group online on some other DLPAR. This check ensures that the agent can dynamically allocate the resources that the service group requires for the DLPAR.

Attributes

Table 1-43 Required Attributes

Required attribute	Description
ManagedSystem	<p>The name of the managed system that contains the partition.</p> <p>Type-dimension: string-scalar</p> <p>Example: mymachine</p>
HMC	<p>Name of the HMC</p> <p>The list of HMCs that control the managed systems.</p> <p>The agent tries to connect to any HMC on this list in the order that they are specified.</p> <p>Type-dimension: string-vector Example: HMC = { myhmc1, myhmc2 }</p>

Optional Attributes

Table 1-44 Optional Attributes

Optional attribute	Description
MemoryRequired	<p>Amount of RAM (in MB) that you want to allocate</p> <p>Type-dimension: string-scalar</p> <p>Default: 0</p> <p>Example: 256</p>
MemoryCritical	<p>Specifies whether the memory allocation is critical.</p> <p>A value of 0 indicates that the online agent function should go ahead even when the required memory was not successfully allocated.</p> <p>Type-dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: 1</p>
CPURequired	<p>The number of dedicated CPUs that you want to allocate.</p> <p>Type-dimension: string-scalar</p> <p>Example: 2</p>
CPUCritical	<p>Specifies whether the CPU allocation is critical.</p> <p>A value of 0 indicates that the online agent function should proceed even when the required CPU was not successfully allocated.</p> <p>Type-dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: 1</p>

Resource type definition

```
type MemCPUAllocator (
    static str ArgList[] = { ManagedSystem, HMC, MemoryRequired,
        MemoryCritical, CPUCritical, CPURequired }
```

```
str ManagedSystem
str HMC[]
str MemoryRequired
str CPURequired
boolean CPUCritical = 0
boolean MemoryCritical = 0
temp boolean IsOnline = 0
)
```

MemCPUAllocator agent notes

The MemCPUAllocator agent has the following notes:

- See [“Configuring password free SSH communication between VCS nodes and HMC”](#) on page 150.
- See [“Dynamic resource allocation scenarios”](#) on page 150.
- See [“Configuring MemCPUAllocator”](#) on page 153.

Configuring password free SSH communication between VCS nodes and HMC

To use remote command operations on the HMC, you must have SSH installed on the DLPAR nodes in the VCS cluster. You must configure the HMC to allow password free SSH access from these partitions. Refer to the appropriate IBM AIX documentation for information.

- ◆ To verify that you have password free SSH access:
 - From each DLPAR in the cluster, execute the following command to test if the password free access works.

```
Eagle> ssh -l hscroot hmc2.veritas.com
Last login:Thur Jun 16 22:46:51 2005 from 10.182.9.34
hscroot@hmc2:~>
```

Once each node can connect to the HMC using SSH without a password, you can start to use the MemCPUAllocator agent.

Dynamic resource allocation scenarios

This section describes different examples of the resource allocation scenarios that the MemCPUAllocator agent can handle. For ease of explanation, consider only the memory resource in these examples. CPU resource implementation is similar.

Consider two DLPARs named Eagle and Vulture. These DLPARs are configured with the following minimum and maximum values memory values.

Table 1-45 The minimum and maximum memory for the DLPARs Eagle and Vulture

DLPAR	Minimum	Maximum
Eagle	512 MB	2 GB
Vulture	512 MB	2 GB

Two service groups SG1 and SG2 have the following resource requirements.

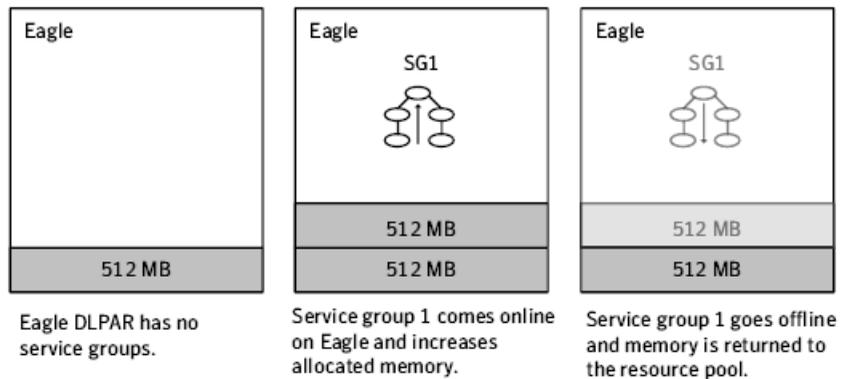
Table 1-46 The memory that is required for service group SG1 and SG2

Service group	Required memory
SG1	512 MB
SG2	512 MB

Scenario 1: A DLPAR node has minimum resources

Assume that the DLPARs start with the minimum values for memory. When SG1 is brought online on Eagle, the online agent function for the agent attempts to allocate 512 MB to Eagle from the free pool. The agent retains the minimum resources for the DLPAR’s overhead operations and allocates resources for the service group in addition to the existing memory. For SG1 to come online the agent allocates an additional 512 MB to Eagle. After this allocation the total current memory for eagle is 1 GB. If SG1 goes offline, the agent deallocates the 512 MB that it allocated when the service group came online. This deallocation brings back the current memory of Eagle to 512 MB.

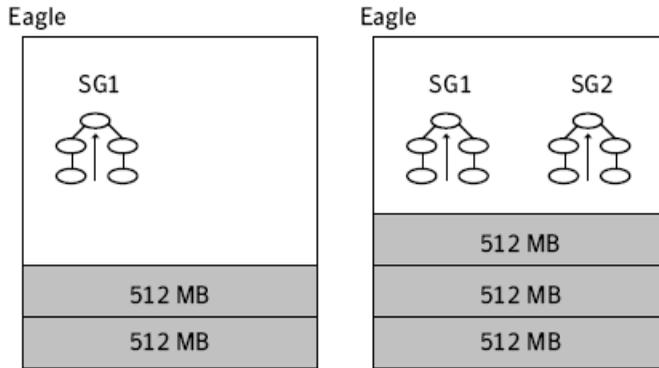
Figure 1-2 Bringing a service group online and taking it offline on a DLPAR



Scenario 2: Bringing another service group online

In this scenario, the Eagle DLPAR starts with 512 MB, and has SG1 online on it. It uses a total of 1 GB of memory. If SG2 is brought up on Eagle, the agent allocates an additional 512 MB of memory to Eagle. This reallocation brings the total memory to 1.5 GB.

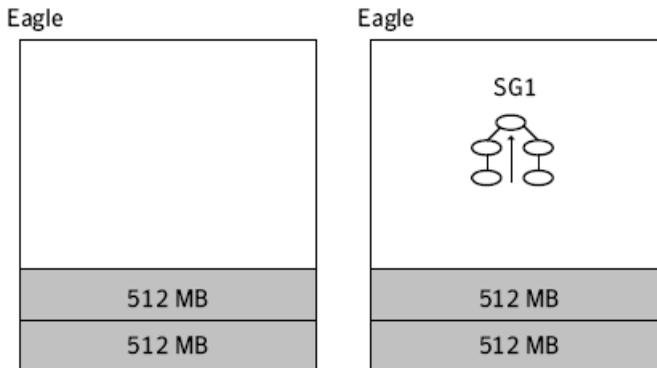
Figure 1-3 Bringing another service group online on a DLPAR



Scenario 3: DLPAR has required resources

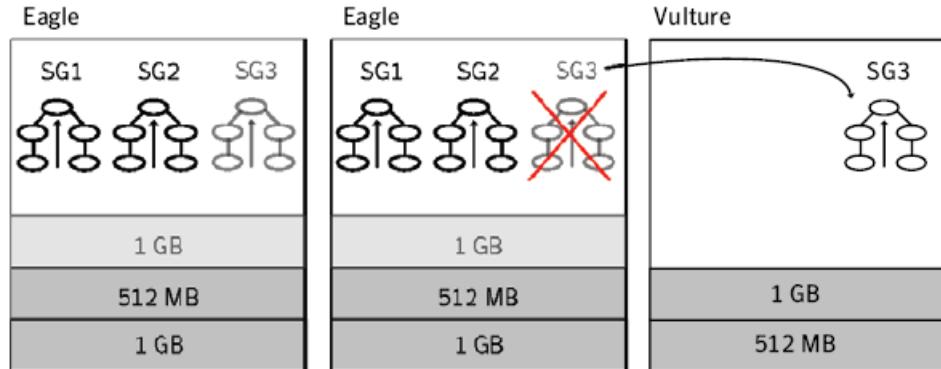
Instead of starting with 512 MB, Eagle starts with 1 GB of initial memory. Eagle has 512 MB more than its minimum amount. If SG1 is brought online on Eagle, the agent determines that Eagle has an extra 512 MB more than its minimum. No service groups use this extra 512 MB. The agent does not allocate any additional memory to Eagle. SG1 is brought online on Eagle and the current memory for Eagle stays 1 GB.

Figure 1-4 DLPAR Eagle starting with 1 GB of initial memmor



Scenario 4: Cannot allocate required resources

Figure 1-5 Exceeding the maximum amount of memory on a DLPAR



Consider the stage in Scenario 2, where SG1 and SG2 are both online on Eagle, which brings its current memory to 1.5 GB. An additional service group SG3 enters the picture and requires 1 GB memory. SG3 tries to come up on Eagle. The agent determines that allocating 1 GB more memory to Eagle exceeds its maximum limit of 2 GB. The agent therefore does not allocate the memory and the online agent function fails, which leads to a resource fault. This resource fault makes the VCS engine stop the online of SG3 on Eagle and try it on Vulture. If Vulture starts with 512 MB and the agent allocates an additional 1 GB to Vulture, its current memory is 1.5 GB. SG3 can fail over and come online on Vulture.

Scenario 5: Service group failover

As in Scenario 2, SG1 and SG2 are both online on Eagle, which brings its current memory to 1.5 GB. Vulture has a current memory configuration of 512 MB. If you switch the service groups from Eagle to Vulture:

- The MemCPUAllocator agent's offline agent function deallocates 1 GB from Eagle (512 MB for SG1 and 512 MB for SG2).
- The VCS engine migrates SG1 and SG2 to Vulture and the agent's online agent function allocates 1 GB to Vulture. This allocation brings Vulture's memory to 1.5 GB.

Configuring MemCPUAllocator

Before you can use the MemCPUAllocator agent, you need to set up SSH access between the HMC and the DLPAR nodes. You must also make sure to configure the MemCPUAllocator resource as a leaf node in the service group's dependency tree in the main.cf file.

See [Figure 1-1](#) on page 147.

Provide values to the MemCPUAllocator resource to specify the resource requirements for that service group. For example, if a service group needs 512 MB memory and two CPUs to start with, the MemCPUAllocator resource definition resembles:

```
MemCPUAllocator mymem (  
    ManagedSystem @eagle = eagle-server  
    ManagedSystem @vulture = vulture-server  
    HMC = { testhmc }  
    RequiredMemory = 512  
    RequiredCPU = 2  
    MemoryCritical = 1  
    CPUCritical = 1  
)
```

Documentation errata

The following sections describe documentation errata.

Veritas Cluster Server User Guide (2278673)

In the "VCS environment variables" section, read the default value of VCS_SERVICE environment variable as "vcs-app" instead of "vcs".

Veritas Volume Manager Administrator's Guide errata

In the Veritas Volume Manager 5.0 MP3 Administrator's Guide, Enabling and disabling vSCSI support section has been corrected.

Enabling and disabling vSCSI support

In this release, Dynamic Multipathing (DMP) is enabled on VIO clients by default. After you install or upgrade Veritas Volume Manager, any vSCSI devices are under DMP control. MPIO is disabled.

You can control the default behavior, so that the DMP support is not enabled when you install or upgrade VxVM. To do this, set the environment variable `__VXVM_DMP_VSCSI_ENABLE` to no.

Note: When you upgrade an existing VxVM installation that has DMP is enabled, then DMP remains enabled, regardless of whether or not the environment variable `__VXVM_DMP_VSCSI_ENABLE` is set to `no`.

To install or upgrade VxVM without enabling DMP support of vSCSI devices

- Before installing or upgrading VxVM, set the environment variable `__VXVM_DMP_VSCSI_ENABLE` to `no`.

```
# export __VXVM_DMP_VSCSI_ENABLE=no
```

Note: The environment variable name `__VXVM_DMP_VSCSI_ENABLE` begins with two underscore (`_`) characters.

After VxVM is installed, use the `vxddladm` command to switch vSCSI devices between MPIO control and DMP control.

To return control to MPIO, disable vSCSI support with DMP. After DMP support has been disabled, MPIO takes control of the devices. MPIO implements multipathing features such as failover and load balancing; DMP acts in pass-through mode.

To disable vSCSI support within DMP and enable MPIO

- 1 Disable vSCSI support.

```
# vxddladm disablevscsi
```

- 2 You are prompted to reboot the devices, if required.

You can add support for any arrays that do not have an ASL by using the `vxddladm addvscsi` operation.

To enable vSCSI support within DMP and disable MPIO

- 1 Enable vSCSI support.

```
# vxddladm enablevscsi
```

- 2 You are prompted to reboot the devices, if required.

To add support for using an array as a vSCSI device within DMP

- ◆ Add support for using an array as a vSCSI device within DMP:

```
# vxddladm addvscsi array_vid
```

To list all of the arrays that DMP supports for use with vSCSI devices

- ◆ List all of the arrays that DMP supports for use with vSCSI devices:

```
# vxddladm listvscsi
```

To remove DMP support on a vSCSI array

- ◆ Remove DMP support on a vSCSI array:

```
# vxddladm rmvscsi array_vid
```

Manual pages errata

One manual page has been updated in this Rolling Patch to include corrections for errors or omissions.

vxdisk(1M) (1528116)

The `rm` keyword description should be as follows:

```
rm Removes the specified disk access records, by disk access name.
```

The `scandisks` keyword description should be as follows:

```
scandisks  
Initiates the rescanning of devices in the operating system  
device tree by VxVM. If necessary, DMP reconfiguration is  
triggered. This allows VxVM to configure and multipath disks  
dynamically.
```

By default, VxVM performs ASL configuration for all of the devices when performing device discovery. To restrict ASL configuration for newly added disks that are not already known to VxVM, specify the `-f` option.

Veritas Cluster Server Installation Guide errata

The latest product documentation is available on the Symantec website:

<http://sort.symantec.com/documents>:

For the Veritas Cluster Server Installation Guide 5.0, the following procedures have updated instructions:

- To install VCS filesets on a node
- To remove VCS packages on a node manually

Veritas Cluster Server database installation and configuration guides errata

The latest product documentation is available on the Symantec website:

<http://sort.symantec.com/documents>:

- Veritas Cluster Server Agent for DB2 Installation and Configuration Guide
- Veritas Cluster Server Agent for Oracle Installation and Configuration Guide
- Veritas Cluster Server Agent for Sybase Installation and Configuration Guide

For these Installation and Configuring Guides 5.0, the following procedures have updated instructions:

- To install the agent
- To remove the agent

