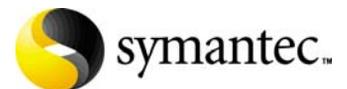


Veritas Storage Foundation™ and High Availability Solutions Read This First

AIX

5.0 Maintenance Pack 3 Rolling Patch 2



Veritas Storage Foundation and High Availability Solutions Read This First

Copyright © 2009 Symantec Corporation. All rights reserved.

Storage Foundation and High Availability Solutions 5.0 Maintenance Pack 3
Rolling Patch 2

Document version: 5.0MP3RP2.0

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

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

THIS 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" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the *Veritas Storage Foundation 5.0 Release Notes*.

The *Veritas Storage Foundation 5.0 Release Notes* can be viewed at the following URL:

<http://entsupport.symantec.com/docs/283990>

The *Veritas Cluster Server 5.0 Release Notes* can be viewed at the following URL:
<http://entsupport.symantec.com/docs/283978>

AIX is a registered trademark of IBM Corporation.

Licensing and registration

Veritas Storage Foundation is a licensed product. See the *Veritas Storage Foundation Installation Guide* for license installation instructions.

Veritas Cluster Server is a licensed product. See the *Veritas Cluster Server Installation Guide* for license installation instructions.

Technical support

For technical assistance, visit

http://www.symantec.com/enterprise/support/assistance_care.jsp and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Chapter 1	Veritas Storage Foundation and High Availability Solutions Read This First	
	System requirements	8
	Supported operating systems	8
	DB2 support	8
	Storage Foundation High Availability fixed issues	9
	Veritas Volume Manager fixed issues	9
	Veritas File System fixed issues	17
	Storage Foundation Cluster File System fixed issues	19
	Storage Foundation for Oracle fixed issues	20
	Storage Foundation for DB2 fixed issues	22
	Storage Foundation for Oracle RAC fixed issues	22
	Veritas Cluster Server fixed issues	23
	Veritas Cluster Server agents for Veritas Volume Replicator fixed issues	30
	Storage Foundation and High Availability known issues	31
	Storage Foundation and High Availability known issues	31
	Veritas Volume Manager known issues	32
	Veritas File System known issues	33
	Storage Foundation Cluster File System known issues	33
	Storage Foundation for Oracle known issues	34
	Storage Foundation for DB2 known issues	34
	Storage Foundation for Oracle RAC known issues	35
	Veritas Cluster Server known issues	35
	Software limitations	36
	Storage Foundation for Oracle software limitations	36
	Storage Foundation for DB2 software limitations	36
	Veritas Cluster Server software limitations	36
	Changes in behavior for Storage Foundation High Availability	37
	Changes in Veritas Cluster Server behavior	37
	Changes in Veritas File System behavior	40
	Downloading the rolling patch archive	40
	Filesets included in this rolling patch	41
	Veritas Cluster Server filesets	41
	File System filesets	41
	Storage Foundation Cluster File System filesets	42
	Storage Foundation for Oracle RAC filesets	43

Storage Foundation filesets	44
Storage Foundation for DB2 filesets	45
Storage Foundation for Oracle filesets	46
Volume Manager filesets	47
Veritas Volume Replicator filesets	47
Installing the Veritas software for the first time	48
Installing Storage Foundation or Storage Foundation Cluster File System and 5.0 MP3 RP2	48
Installing Storage Foundation for Oracle RAC and 5.0 MP3 RP2	51
Prerequisites for upgrading to 5.0 MP3 RP2	54
Upgrading 5.0 MP3 to 5.0 MP3 RP2	54
Performing a phased upgrade to 5.0 MP3 RP2 on a cluster	55
Performing a full upgrade to 5.0 MP3 RP2 on a cluster	69
Upgrading to 5.0 MP3 RP2 on a standalone system	86
Verifying software versions	89
Removing 5.0 MP3 RP2	90
Removing 5.0 MP3 RP2 from Veritas Cluster Server	90
Removing 5.0 MP3 RP2 on Storage Foundation or Storage Foundation Cluster File System	95
Removing 5.0 MP3 RP2 on Storage Foundation for Oracle RAC	97
Documentation addendum	98
MemCPUAllocator agent	98
Documentation errata	106
Manual pages errata	106
Veritas Cluster Server Installation Guide errata	107
Veritas Cluster Server database installation and configuration guides errata	

Veritas Storage Foundation and High Availability Solutions Read This First

This document provides release information about the products in the Veritas Storage Foundation and High Availability 5.0 Maintenance Pack 3 (MP3) Rolling Patch 2 (RP2) AIX 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://entsupport.symantec.com/docs/282024>

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

For further details, depending on the product for which you want to install this Rolling Patch, refer to one of the following Release Notes 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 RP2 release operates on the architectures and operating systems shown below:

- AIX 5.3 TL7 with SP2 or later
- AIX 6.1 TL0 with SP4 or later

DB2 support

This release supports DB2 9.5 FixPak 1, in addition to the DB2 database versions that are supported in the 5.0 MP3 release.

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 for Oracle fixed issues](#)
- [Storage Foundation for DB2 fixed issues](#)
- [Storage Foundation for Oracle fixed issues](#)
- [Veritas Cluster Server fixed issues](#)

Veritas Volume Manager fixed issues

[Table 1-1](#) describes fixed issues in the Veritas Volume Manager 5.0 MP3 RP2 release.

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

Incident	Description
850816	You can now delete snap objects from a mounted volume.
1097258	The <code>vxconfigd</code> daemon no longer hangs when an array is disconnected.
1108839	Turning on the <code>dmp_cache_open</code> tunable no longer slows down the <code>vxconfigd</code> daemon when run with 2048 dual path LUNs.
1184280	Added additional debug messages around the <code>VE_BADPROTOV</code> error message to improve debugging.
1195591	Fixed the cause of a panic when a cluster had an empty RVG.
1222125	Fixed an issue in which after converting a big volume group, you could not mount a file system on the volume.
1224659	Fixed an issue in which the <code>vxconfigbackup -p</code> script sometimes created a zero-length <code>.binconfig</code> file.
1286298	Fixed an issue in which proper locks were not taken in all necessary places while modifying <code>last_sent_seqno</code> .
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.
1314961	Fixed an issue in which if a disk tag name contained long DBCS characters, the configuration daemon was not accessible.

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

Incident	Description
1321272	Fixed the an issue in which some VxVM commands hung after disconnecting, then reconnecting to the FC site link.
1321298	Fixed the cause of a <code>vxconfigd</code> daemon core dump that occurred after reconnecting the FC site link and heartbeat link.
1370927	Fixed an issue in which the VTOC of disks in a cluster became corrupted.
1374603	Fixed a cause of data corruption in the <code>dmp_bypass_iodone()</code> call.
1380386	The appropriate number of I/O threads are now created for systems with more than 8 CPUs.
1388883	Fixed an issue in which rebooting a controller caused the diskgroups to be disabled.
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.
1408367	Fixed the cause of a system panic when <code>mutex_panic()</code> was called from <code>vol_rwsleep_wrlock()</code> .
1411438	The dynamic multi-pathing tunable <code>dmp_enable_restore</code> value is now persistent when it is set to OFF.
1414336	Fixed an issue in which some disk devices did not appear in the <code>vxdisk list</code> command output.
1414469	Fixed an issue in which the <code>vxddladm listsupport all</code> did not display up-to-date information.
1416080	Fixed the cause of a system panic in the <code>vol_change_disk()</code> routine that was due to NULL pointer dereference.
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.
1425338	Fixed an issue in which connect <code>rlinks</code> failed to be connected, followed by <code>vxconfigd</code> hanging on a secondary node.
1445884	Fixed an issue in which the <code>vxdisk -e list</code> command dumped core.
1446208	Changed message V-5-1-2140 from an error message to an informational message.
1446216	Fixed an issue in which LUN mapping changes on node 0 caused node 1 and node 2 to panic.

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

Incident	Description
1449001	Enabled the auto-compiler flag selection, which is used for enabling and disabling storage keys support.
1449264	Fixed the cause of a panic in <code>dmp_get_path_deferq</code> that occurred while creating a volume.
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.
1457132	Fixed the cause of data corruption when running the <code>vxdumpadm disable path</code> and <code>vxdumpadm disable ctlr</code> commands.
1458792	Fixed in issue in which the <code>*unit_io</code> and <code>*pref_io</code> tunables became set to 32 MB after upgrading from the Storage Foundation 5.0 MP1 release to the 5.0 MP3 release.
1459831	Fixed an issue in which replication hung due to a deadlock on a secondary that had a TCP multiconnection and was managed by <code>nmcom</code> .
1460582	Fixed an issue in which the <code>vxdumpadm -f disable enclosure</code> command failed with a Segmentation Fault and dumped core.
1461314	DMP no longer uses the SCSI bypass on single path disks for path-suppressing TPD.
1461338	VxVM now correctly handles CPU affinity switching.
1461717	Fixed an issue in which the <code>vxsnap make</code> command caused the <code>vxconfigd</code> daemon to hang.
1463547	Fixed the cause of a <code>vxconfigd</code> core dump that occurred when dynamically reconfiguring a LUN.
1469487	The I/O buffer start time is no longer modified as part of error analysis.
1470548	Fixed an issue in which a campus cluster detached site could not be reattached.
1471003	Fixed an issue in which the <code>vxdbg -s import oradg</code> command failed with following error on a campus cluster: V-5-1-10978 Disk group oradg: import failed: required lock not held in transaction
1471658	Fixed the cause of a <code>vxconfigd</code> daemon core dump that occurred in the <code>priv_get_all_udid_entry()</code> call.
1471754	Fixed the cause of a <code>vxconfigd</code> daemon core dump that occurred in the <code>ssb_increment_dg_abort()</code> call.

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

Incident	Description
1471763	Fixed the cause of the following error: build_devlink_list: readlink failed for /dev/vx/rdisk/ludg: Invalid argument
1471813	The <code>vxddmpadm start restore</code> command now indicates that the restore daemon cannot be started when the <code>dmp_enable_restore</code> tunable is set to OFF.
1472736	Fixed the cause of a system panic in the <code>vxddmp</code> module that was due to a NULL pointer dereference.
1473638	Fixed the cause of a failover in the IOCTL context for coordinator disks.
1475707	Added an error message for attempting to import unwritable disks.
1477143	The cluster volume manager failback protocol is now triggered when <code>cur_pri</code> is null and at least one DMP node of the same LUN group is <code>DMPNODE_SHARED</code> .
1479729	Fixed the cause of an I/O hang on the primary node after a secondary node crashed.
1479735	Fixed the cause of an I/O hang on a slave if the master (logowner) crashed with a data change map active.
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.
1482687	Fixed the cause of an I/O hang during and after CVM reconfiguration following a reboot of the master node.
1483164	Fixed an issue in which disks with the <code>NOLABEL</code> state were usable via the CLI.
1483201	Fixed an issue in which the Device Discovery Layer (DDL) sometimes set the unique disk identifier (UDID) value to <code>INVALID</code> . Multiple disks set to <code>INVALID</code> resulted in the following error: VxVM vxio V-5-0-1056 new disk disk_id has a non-unique UDID
1483643	Fixed an issue in which a raid 5 volume would not start on 3PAR Thin Provisioning LUNs.
1484919	Fixed an issue in which a system that was upgraded to the 5.0 MP3 release could not be booted.
1485379	Fixed an issue in which the <code>vxtask -l list</code> command displayed incorrect progress of the <code>vxsnap admir</code> command, which was used to link a snapshot volume to the source volume.

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

Incident	Description
1488084	Fixed an issue in which the <code>vxmpadm iostat</code> command reported different amounts of read/write blocks than the <code>vxstat</code> , <code>iostat</code> , and <code>sar -d</code> commands.
1500389	The <code>vxrootadm</code> command now automatically enables the <code>use-nvramrc?</code> variable.
1501165	Changed the V-5-1-2140 message from an error to a warning.
1502842	Fixed an issue in which the <code>dmpolicy.info</code> file did not get updated after upgrading the packages from Storage Foundation (SF) 5.0 MP3 RP1 to SF 5.1.
1503168	Fixed an issue in which the diskgroup for disks without a private region (<code>nopriv</code> disks) could not be imported.
1507291	Fixed an issue in which setting the <code>dmp_monitor_fabric</code> value to ON triggered unexpected offlining of paths on a DMX4 array.
1507935	Fixed an issue in which the <code>vxconfigd</code> daemon dumped core when the <code>settag</code> keyword was used to specify a long site name in a campus cluster.
1508462	Fixed the cause of a <code>vxconfigd</code> hang that occurred due to a split brain condition on a cluster.
1512352	Fixed an issue in which the <code>vxconfigrestore</code> command failed with the following error: VxVM vxconfigrestore ERROR V-5-2-3706 Diskgroup configuration
1512805	Fixed an issue in which two disks from different enclosures with the same array volume ID were treated as duplicate array volume IDs.
1515581	Fixed an issue in which recreating a shared diskgroup put <code>CVMVolDg</code> in an empty <code>KSTATE</code> and offlined clustered file systems.
1525121	Fixed an issue in which EFI disks were in an error state after installing the Storage Foundation 5.0 MP3 RP1 patches.
1525819	Fixed an issue in which the <code>vxconfigbackup</code> command failed to work on a diskgroup that had 2 TB LUNs.
1527247	Fixed an issue in which the <code>vxstat</code> command showed twice the I/O activity on a mirror volume compared to the source volume.
1528368	Fixed the cause of an I/O hang during the data change map transition after performing <code>vxresize</code> operations on the primary node.
1529157	The <code>vxdisk -e list</code> command output now includes a wider range of values for the <code>TYPE</code> and <code>STATUS</code> columns.

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

Incident	Description
1532363	Fixed an issue in which the <code>vxdisk updateudid</code> command is corrupted the UDID of a disk, which made diskgroup imports fail.
1534038	Fixed an issue in which DMP stats sometimes used invalid I/O stats entries, which led to a panic on the host.
1534379	Fixed an issue in which the <code>vxdg split</code> command failed with the following error: Internal configuration daemon error
1543803	Fixed an issue in which the <code>vxconvert</code> command failed to convert a volume group.
1544051	Fixed an issue in which the incorrect bit was being checked for an EMC Symmetrix thin device.
1586879	Improved performance of the <code>vxdisk online</code> command when used on large configurations.
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.
1589172	Fixed an issue in which the <code>vxdisksetup</code> and <code>vxdiskunsetup</code> commands sometimes failed for EFI disks.
1589881	Fixed an issue in which the dump device was changed to none (dumps disabled) after encapsulating a boot disk.
1590314	The <code>vxdmadm getsubpaths dmpnodename</code> command now validates the <code>dmpnodename</code> value before getting the subpath information.
1594325	Backed out <code>*unit_io</code> and <code>*pref_io</code> changes that were made in the 5.0 release.
1596811	Improved performance when re-enabling individual paths.
1597868	Fixed an issue in which, on a secondary node, <code>rlink</code> paused and generated the “Incorrect magic number or unexpected upid” error message, and the <code>secondary_log_err</code> flag got set.
1598145	Fixed the cause of a crash in the <code>vxdmproot install</code> command during startup on an AIX 5.3 SAN boot disk.
1598706	Fixed the cause of a system crash that occurred while mirroring the rootdisk.
1661938	Enhanced the <code>vxconfigd</code> daemon to take advantage of the AIX very large address space model.

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

Incident	Description
1663338	Fixed an issue in which I/O load balancing to LUNs between two VIO servers not working.
1675221	Fixed an issue in which the <code>vxdmpadm setattr enclosure</code> command created identical disk attribute names.
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.
1711339	VVR tunables can now be modified using the <code>vxtune</code> command.
1714619	Fixed the cause of a panic due to DMP calling <code>iodone()</code> twice.

[Table 1-2](#) describes fixed issues in the Veritas Volume Manager 5.0 MP3 RP1 release, which are included in this release.

Table 1-2 Veritas Volume Manager 5.0 MP3 RP1 fixed issues

Incident	Description
425273	Fixed an issue with VVR RU thread not starting <code>nio</code> after it is created from than waiting for all replicas to have NIO's created.
1182475	Fixed the <code>vxdg split</code> failing if the CVM master changes.
1192105	Fixed the <code>vxdg -n [newdgm] deport [origdgm]</code> command causing a memory leak.
1230360	Fixed a system panic in <code>vol_klog_start()</code> due to accessing freed <code>mv read_sio</code> .
1266730	Fixed the <code>vxtask</code> command to display the resync progress subtask for shared volumes with DRL.
1364320	Fixed a issue with <code>vxdmproot install</code> causing machine to hang upon reboot.
1364324	Fixed an issue with V SCSI: A/P LB I/O policy not working with enabled DMP support on boot devices.
1364332	Fixed an issue with the VM commands not working on DMP Enabled boot path.
1364335	Created a command to verify whether the DMP is controlling the rootdisk.
1368752	Fixed an issue when there are no mirrors to read, <code>VOL_READ_MIRRORS ioctl</code> returns -1 instead of 1.
1372340	Fixed an issue with <code>vxplex core dumps</code> during <code>vxassist addlog</code> due to DRL log length being less than 33 blocks.

Table 1-2 Veritas Volume Manager 5.0 MP3 RP1 fixed issues

Incident	Description
1376656	Fixed an issue with vxcached never deletes old snaps when cache hits HWM.
1385922	Fixed a system panic due to memory allocation.
1387033	Fixed a system panic in bcopy() due to null passed in from voliocntl_copyin()
1389512	Able to force import diskgroup version 80 in VxVM 5.0.
1389584	Fixed a system panic in vol_putdisk() code.
1396427	Enhanced DMP to handle failing IO when it is not able to interpret sense data.
1397712	Fixed an issue with the vxsnap restore manual page is unable to properly freeze or thaw filesystems in a CVM environment.
1397879	Enhanced the vxresize manual page to run from non-CVM master.
1403123	Fixed an issue with vxconfigd sleeping and no vx commands were responding.
1411636	Fixed a secondary log error causing rlink disconnect after IBC unfreeze.
1414381	Fixed an issue with VVR I/O hanging due to the wrong generation number assignment after recovery.
1414441	The vxsnap manual page includes mirror=enclosure parameter to avoid being mirrored on the same enclosure.
1425434	Fixed an issue with CVR fails to connect rlinks followed by vxconfigd hangs on secondary.
1425919	Fixed an issue with vxesd looping using 100% of one CPU.
1427267	Fixed a CVR panic in VOLSIQ_MORE due to corrupted volsioq_start queue.
1427284	Fixed an issue with vxdmpadm dumped core when executing vxdmpadm list dmpnode command.
1428106	Fixed a system panic in vxio:voldr1_trans_copy.
1431279	Fixed an issue with vxconfigd core dumps.
1435471	Fixed an issue with the cluster nodes panicking in voldco_or_pvmbuf_to_pvmbuf code after installing 5.0 MP3.
1436917	Fixed an issue with after installing VM patch on AIX, install-db would be created if the vxio in Defined state.
1440837	Fixed a panic due to memory allocation failure in interrupt context.

Table 1-2 Veritas Volume Manager 5.0 MP3 RP1 fixed issues

Incident	Description
1441020	Fixed a secondary panic due to double free of message with TCP protocol and 16 connection.
1441072	Fixed an issue with siteread policy is not honoured.
1441131	Fixed an issue with VxFS Corruption Detected when DCM log plex are attached with mirrored volume and VVR is not configured.
1442369	Fixed a bug in vxconfigbackupd script leading to 0 byte binconfig file being created.
1443706	Fixed an issue in FMR3, I/Os initiating DCO updates for clearing DRL async clear region may not wait for its completion.
1443752	Fixed an issue in a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load.
1449266	Fixed a panic in dmp_get_path_deferq while creating a volume.
1455062	Fixed an issue with the master node crashing if a node leaves before responding to the MV- serialization protocol.
1512521	StorageKeys: Adding compilation flag for VxVM to make function pointer references to legacy modules work in AIX 6.1 key aware code.
1710030	Fixed the cause of a panic the occurred with LUNs larger than 1 TB.

Veritas File System fixed issues

[Table 1-3](#) describes fixed issues in the Veritas File System 5.0 MP3 RP2 release.

Table 1-3 Veritas File System 5.0 MP3 RP2 fixed issues

Incident	Description
847454	The VxFS <code>cp</code> command now has similar performance as the JFS2 <code>cp</code> 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 <code>quotacheck</code> command with more than 30 quota-enabled file systems in the <code>/etc/fstab</code> file.

Table 1-3 Veritas File System 5.0 MP3 RP2 fixed issues

Incident	Description
1363639	Fixed the cause of an exception in the <code>vx_vnmap_min()</code> call that was due to the <code>vx_fsext_info()</code> call of a force umounted file system having a null <code>fse_fsext_info</code> pointer.
1387074	Fixed the cause of the <code>fsclustadm cfsdeinit</code> command failing with the "device busy" error.
1416745	Fixed an issue in which the <code>dm_get_allocinfo()</code> call failed with the EIO error for EXT4 inodes with indirects.
1432323	Fixed a loop issue in which the <code>vx_do_putpage()</code> call due to the page going beyond <code>i_wsiz</code> .
1443034	Fixed the cause of hangs with the <code>vxresize</code> and <code>fsadm -b</code> 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 <code>vx_ifree_scan_list()</code> call in <code>vx_iflush.c</code> .
1454783	The <code>vx_multi_bufinval()</code> call now releases the CPU for a local mount with large extents.
1459329	New VxFS tunables and new <code>vxfsstat</code> counters now increase the number of VMM buffers per PDT, which improves performance.
1500197	Scanning of pages in <code>v_write</code> during <code>putpage</code> no longer causes GAB/LLT failure.
1517415	Fixed the cause of a core dump when running the <code>ncheck</code> command.
1526568	The <code>vx_tflush_map()</code> 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 <code>vx_fsync_range()</code> 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 <code>vx_setext()</code> while having the <code>VX_GROWFILE</code> flag set sometimes failed to grow the file or exposed uninitialized data.
1667628	Implemented a new tunable, <code>flush_chunk_size</code> , which is used by the <code>vx_putpage_dirty()</code> and <code>vx_mm_invalidatep()</code> calls.
1701842	Fixed a cause of Concurrent I/O returning the ENOTSUP error to DB2, which caused DB2 to crash.

Table 1-4 describes fixed issues in the Veritas File System 5.0 MP3 RP1 release, which are included in this release.

Table 1-4 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 <code>vx_delay()</code> .
1400046	Fixed an issue with the <code>fsapadm enforceckpt mount_point</code> 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 <code>vxresize</code> 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 (<code>vxi_alloc_fail</code>).
1415188	Fixed a full <code>fsck</code> core dump that was due to running out of swap space and a <code>malloc</code> failure.
1423867	Fixed an issue in which the <code>vx_convnode_data_files()</code> 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 <code>lockcount</code> of an FS thread upon <code>thread_terminate()</code> .

Storage Foundation Cluster File System fixed issues

Table 1-5 describes fixed issues in the Storage Foundation Cluster File System 5.0 MP3 RP2 release.

Table 1-5 Storage Foundation Cluster File System 5.0 MP3 RP2 fixed issues

Incident	Description
1286525	Fixed an issue in which Java threads hung on SFCFS functions.
1518713	The <code>vxfsckd -n</code> command now initializes the <code>nthrs</code> variable.
1539892	Fixed an issue in which a clustered file system that was mounted on one node required <code>fsck</code> to be run.
1591313	Fixed the cause of a loop in the <code>vx_delete_get_freespace()</code> call due to a <code>smap</code> being marked bad.

Table 1-5 Storage Foundation Cluster File System 5.0 MP3 RP2 fixed issues

Incident	Description
1600241	Fixed the cause of a hang that occurred after another node in the cluster crashed.

[Table 1-6](#) describes fixed issues in the Storage Foundation Cluster File System 5.0 MP3 RP1 release, which are included in this release.

Table 1-6 Storage Foundation Cluster File System 5.0 MP3 RP1 fixed issues

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

Storage Foundation for Oracle fixed issues

[Table 1-7](#) describes fixed issues in the Storage Foundation for Oracle 5.0 MP3 RP2 release.

Table 1-7 Storage Foundation for Oracle 5.0 MP3 RP2 fixed issues

Incidents	Description
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: <ul style="list-style-type: none"> ■ /etc ■ /etc/default ■ /etc/init.d ■ /etc/rc2.d ■ /opt
1508346	Added a date stamp to entries in the <code>vxsnapadm_50.log</code> file, which is used for <code>trace vxsnapadm</code> issues.
1511321	Fixed multiple issues with the <code>dbed_checkconfig</code> 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.
1526653	Fixed an issue in which the <code>dbed_vmchecksnap</code> script output an error if the dco object name was renamed from <code>*_dco</code> .

Table 1-7 Storage Foundation for Oracle 5.0 MP3 RP2 fixed issues

Incidents	Description
1530125	Fixed an issue in which the owner of the following directories was changed when installing VRTSGbms packages for the Storage Foundation for Oracle 5.0 or 5.0 MP3 releases: <ul style="list-style-type: none"> ■ /etc ■ /etc/default ■ /etc/init.d ■ /etc/rc2.d ■ /opt
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).
1651363	Fixed a security issue with the vxdbms server, in which an attacker could see the name and port of the server.
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 1-8](#) describes fixed issues in the Storage Foundation for Oracle 5.0 MP3 RP1 release, which are included in this release.

Table 1-8 Storage Foundation for Oracle 5.0 MP3 RP1 fixed issues

Incidents	Description
1425256	Support flashsnap CVM slave.
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.
1433244	Improved boot time for the DBED repository database server startup script.
1433571	Sybase repository database server is no longer creating world writable files under /tmp.
1434688	Storage Foundation for Oracle is no longer creating world writable files under /tmp.
1435527	Improved boot time for DBEDAgent startup script.

Storage Foundation for DB2 fixed issues

[Table 1-9](#) describes fixed issues in the Storage Foundation for DB2 5.0 MP3 RP2 release.

Table 1-9 Storage Foundation for DB2 5.0 MP3 RP2 fixed issues

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

[Table 1-10](#) describes fixed issues in the Storage Foundation for DB2 5.0 MP3 RP1 release, which are included in this release.

Table 1-10 Storage Foundation for DB2 5.0 MP3 RP1 fixed issues

Incidents	Description
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.
1433244	Improved boot time for the DBED repository database server startup script.
1433571	Sybase repository database server is no longer creating world writable files under <code>/tmp</code> .
1434688	Storage Foundation for DB2 is no longer creating world writable files under <code>/tmp</code> .
1435527	Improved boot time for DBEDAgent startup script.

Storage Foundation for Oracle RAC fixed issues

[Table 1-7](#) describes fixed issues in the Storage Foundation for Oracle RAC 5.0 MP3 RP2 release.

Table 1-11 Storage Foundation for Oracle RAC 5.0 MP3 RP2 fixed issues

Incidents	Description
1382034	Fixed an issue in which the MultiPrivNIC agent failed over the IP address even when it was not required.
1525117	Fixed an issue in which the MultiPrivNIC agent was not able to plumb the IP address on the configured devices.

Table 1-11 Storage Foundation for Oracle RAC 5.0 MP3 RP2 fixed issues

Incidents	Description
1593859	Reduced the time it takes for Oracle to start with VCSIPC.
1597480	The LMX code now calls the <code>tstop()</code> function calling the <code>tstart()</code> function to avoid a race condition.

Veritas Cluster Server fixed issues

[Table 1-12](#) describes fixed issues in the Veritas Cluster Server 5.0 MP3 RP2 release.

Table 1-12 Veritas Cluster Server 5.0 MP3 RP2 fixed issues

Incidents	Description
1070177	[Agents] Fixed an issue to include a new attribute to use the <code>db2start</code> command. There was no option to use the <code>db2start</code> command. Added optional attribute <code>UseDB2start</code> to allow users to start DB2 using the <code>db2start</code> command.
1368385	[Agents] Fixed an issue in which <code>DiskGroupSnap</code> does not work if layered volumes are used. <code>VxVM</code> 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 <code>VxVM</code> 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.
1377324	[Agents] Fixed a parsing error which caused an error message to appear in the <code>/var/VRTSvcs/log/tmp/Oracle-0</code> file.
1478488	[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 <code>DiskGroup</code> resource has <code>StartVolumes=1</code> , all the volumes of the disk group are started by the <code>DiskGroup</code> 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 <code>DiskGroup</code> agent to start all the volumes before initiating recovery.
1504693	[GAB/LLT] Fixed an issue in which LLT cannot provide backenable to GAB. This resulted in an error being produced from the GAB module <code>gabwrite()</code> function.

Table 1-12 Veritas Cluster Server 5.0 MP3 RP2 fixed issues

Incidents	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>
1513967/ 1638722	<p>[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.</p>
1522568	<p>[Agents] Fixed an issue in which the agent framework crashed while setting the resource name for the dependent attribute.</p>
1538208	<p>[VCS] Fixed an issue in which the value of attribute HostUtilization is not 0 even after HostMonitor resource is deleted.</p>
1539087	<p>[Agents] Fixed an issue in which the agent framework seems to be leaking memory during message logging.</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>
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>
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. See “New attribute EntryPointTimeout” on page 39.</p>

Table 1-12 Veritas Cluster Server 5.0 MP3 RP2 fixed issues

Incidents	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” on page 37.</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>
1589851	<p>[GAB] Fixed the cause of a system panic that was due to depleted memory reserves.</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 VCSshm group with a HostMonitor type resource VCSshm. 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>
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>
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 <code>haclus -modify UserNames u1 p1 u2 p2</code> is run, then even invalid user names were accepted.</p>

Table 1-12 Veritas Cluster Server 5.0 MP3 RP2 fixed issues

Incidents	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 <code>vx fenceswap</code> command. If the keys on coordinator disks are accidentally cleared, they can be refreshed using the <code>vx fenceswap</code> 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 <code>vx fenceswap</code> 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 <code>vx fenceswap</code> utility is run on existing coordinator disks, then the registrations are done in the commit phase instead of the modify phase.</p>
1603120	<p>[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” on page 37.</p>
1631012	<p>[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 <code>llttab</code> file, <code>lltconfig</code> would configure that link successfully when it should not.</p>
1633973	<p>[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.</p>
1634924	<p>[VCS] Fixed an issue in which the engine logs indicated CPU usage even after the HostMonitor resource is deleted.</p>
1635044	<p>[Fencing] The <code>vx fen</code> scripts now determine the location of <code>ssh</code> and <code>scp</code> from the environment variables if <code>ssh</code> or <code>scp</code> are not available in the standard path.</p>

Table 1-12 Veritas Cluster Server 5.0 MP3 RP2 fixed issues

Incidents	Description
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 the Sybase and SybaseBk agents. The value of this attribute can be set to the absolute path of the RUN_<servername> file.
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.
1640292	[HAD] Fixed an issue to increase the resilience of HAD to extreme load conditions.
1665036	[Fencing] The vxfen startup process on retry no longer displays the "RFSM GAB err 16" error when the cluster is fencing a node out of the cluster.
1668609	[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.
1670337	[LLT] Added a mechanism to track the operating system timeouts that are registered by LLT with the operating system.
1672335	[Fencing] The vxfenclearpre command now outputs the progress of its execution.

Table 1-12 Veritas Cluster Server 5.0 MP3 RP2 fixed issues

Incidents	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>hagrp -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>
1675815	<p>[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.</p>
1677412	<p>[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.</p>
1678110	<p>[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.</p>

Table 1-12 Veritas Cluster Server 5.0 MP3 RP2 fixed issues

Incidents	Description
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.
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.
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.

Table 1-13 describes fixed issues in the Veritas Cluster Server 5.0 MP3 RP1 release, which are included in this release.

Table 1-13 Veritas Cluster Server 5.0 MP3 RP1 fixed issues

Incidents	Description
1379299	LLT: fixed llt_recordmac() messages.
1394624	LLT: fixed an issue where the lltdlv thread spun indefinitely.
1395905	Changes implemented to close device file for device vxdmconfig.
1396639	Return code for SCSI commands are now logged before re-using a variable.
1397692	Removed a condition where VCS engine clients hung in connect when the target system was down.
1398750	Added the MemCPUAllocator agent. See " MemCPUAllocator agent " on page 98.
1403471	Reduced time for global cluster fault detection.
1404384	Global groups can switch over to a node where WAC is not running, when PreSwitch is set to 1 and HAD runs properly.
1414709	The <code>hagr -offline</code> command and <code>hares -offline</code> command now behave similarly when you bring the last resource in a service group offline.
1424927	Optimized GAB connect messages.
1457429	Removed the VCS NOTICE V-16-1-53021 message after the <code>hastart</code> command is run.

Table 1-13 Veritas Cluster Server 5.0 MP3 RP1 fixed issues

Incidents	Description
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 <code>ha</code> commands to change this attribute are:</p> <pre data-bbox="458 656 1029 730"># haconf -makerw # hagr -modify servicegroup_name PreOnline 1 # haconf -dump -makero</pre>
1510002	<p>Fixed an issue where the HostMonitor tool misreported CPU as busy when it is idle in an AIX 5.0 MP3 Micropartition environment.</p>

Veritas Cluster Server agents for Veritas Volume Replicator fixed issues

[Table 1-14](#) describes fixed issues in Veritas Cluster Server agents for the Veritas Volume Replicator 5.0 MP3 RP2 release.

Table 1-14 Veritas Cluster Server agents for Veritas Volume Replicator 5.0 MP3 RP2 fixed issues

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

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 RP2 hosts on Storage Foundation Manager 1.1.1 can be viewed at the following URL:

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

Veritas Volume Manager known issues

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

AIX Volume Manager 5.0MP3 documentation describes incorrect default behavior for vSCSI devices (1460494)

The AIX Volume Manager 5.0MP3 documentation incorrectly stated the default behavior for vSCSI devices. The correct default behavior is as follows:

In this release, DMP is enabled on VIO clients by default. All VSCSI devices are under DMP control. Use the `vxddladm disablevscsi` command followed by a reboot to disable DMP and enable MPIO on these vSCSI devices.

The `vxdmadm getsubpaths` command displays incorrect output in a vSCSI environment (1453073)

In a vSCSI environment, when the I/O policy is A/P, the output of the `vxdmadm getsubpaths` command displays all primary and secondary paths as active, even though only one path is used for I/O.

Workaround

There is no known workaround.

The `vxdmroot uninstall` command may not clean all stale ODM entries (1453103)

The `vxdmroot uninstall` command only cleans up ODM entries for the booting path. A few stale ODM entries could be seen for the remaining paths of the boot disk.

Workaround

There is no known workaround. The stale entries do not cause any harm to system functionality.

Possible vxconfigd crash with powerpath devices (1706503)

On a setup with powerpath devices, the `vxconfigd` daemon may crash after running following set of commands:

```
# vxddmpadm exclude ctlr=ctlr  
# vxddmpadm include ctlr=ctlr
```

Workaround

There is no known workaround.

Veritas File System known issues

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

A failure message displays while rejecting VRTScavf (1744917)

While rejecting VRTScavf 5.0.3.200 with the `installp -r` 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).
```

The message displays due to a bug in the `sysck` utility, which is maintained by IBM. The problem is reported and is being tracked via PMR #53405,756.

Workaround

Run the following command manually after rejecting the patch:

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

Storage is not reclaimed on a mirrored disk after data is removed from the volume (1741147)

A reclaim operation on a file system mounted on a mirrored volume fails to reclaim space from the mirror.

Workaround

None.

Storage Foundation Cluster File System known issues

There are no known issues in 5.0 MP3 RP2 release of Storage Foundation Cluster File System.

Storage Foundation for Oracle known issues

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

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.

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.

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" for the workaround.

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

If your clone SID name contains the primary SID name in the beginning, the `dbed_vmclondb` and `dbed_clondb` 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=Prod1,server_name=svr_name -f snapplan -r relocate_path
# dbed_vmclondb -S Prod -o mountdb \
new_sid=Prod1,server_name=svr_name -f snapplan -r relocate_path
# dbed_clondb -S Prod1 -m mount_point -c ckpt_name
```

Workaround

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

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.

- [Storage Foundation for Oracle software limitations](#)
- [Storage Foundation for DB2 software limitations](#)

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 RP2. Otherwise, prior backups cannot be restored using the 5.0 MP3 RP2 restore script.

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 20 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.

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 RP2. Otherwise, prior backups cannot be restored using the 5.0 MP3 RP2 restore script.

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

For more information see [“Storage Foundation for DB2 fixed issues”](#) on page 22 for incident 1425261.

Veritas Cluster Server software limitations

There are no known limitations in the 5.0 MP3 RP2 release of Veritas Cluster Server.

Changes in behavior for Storage Foundation High Availability

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

Changes in Veritas Cluster Server behavior

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

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 1-15 indicates the values that apply to these new attributes.

Table 1-15 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.

Table 1-15 Attribute values to schedule class and priority of agent entry points

Attributes	Values
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 EntryPointTimeout

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.

Changes in Veritas File System behavior

The following sections describe changes in Veritas File System behavior for this release.

Flushing in chunks

In this release, all the flushing and invalidation of client pages associated with VxFS files are performed in chunks. The default chunk size is 64 MB; however, the chunk size can be tuned to 128 MB or 256 MB, instead. Flushing in chunks might result in `vx_sched` consuming more CPU time. If the CPU utilization of `vx_sched` is too high, the chunk size can be increased using the `vxtunefs` command.

The tunable `flush_chunk_size` has been introduced as a precaution, only. The default chunk size of 64 MB should suffice for all workloads.

To tune the flush chunk size

- 1 Specify the chunk size:

```
# vxtunefs -D flush_chunk_size=value
```

value can be 1 for 256 MB, 2 for 128 MB, or 3 for 64 MB. 64 MB is the default size. For example, to tune the flush chunk size to 256 MB, enter:

```
# vxtunefs -D flush_chunk_size=1
```

The new chunk size takes effect immediately. However, if you instead manually edited the `/etc/vx/vxfssystem` file to add or change the `flush_chunk_size` tunable value, then the new chunk size only takes effect when mounting the first VxFS file system on the server.

- 2 Verify the flush chunk size value:

```
# cat /etc/vx/vxfssystem | grep flush_chunk_size  
flush_chunk_size
```

If the `/etc/vx/vxfssystem` file does not exist, the `vxtunefs` command creates the file.

Downloading the rolling patch archive

The patches included in the 5.0 MP3 RP2 release are available for download from the Symantec website. After downloading the 5.0 MP3 RP2 file, use the `tar -z` command to uncompress and extract the archive.

For the 5.0 MP3 RP2 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/282024>

Filesets included in this rolling patch

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

- [Veritas Cluster Server filesets](#)
- [File System 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](#)
- [Volume Manager filesets](#)
- [Veritas Volume Replicator filesets](#)

Veritas Cluster Server filesets

[Table 1-16](#) describes the VCS filesets that are included in this rolling patch:

Table 1-16 VCS filesets

Filesets
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSvcsag.rte.bff
VRTSvcs.rte.bff
VRTSvxfen.rte.bff

File System filesets

[Table 1-17](#) describes the File System filesets that are included in this rolling patch:

Table 1-17 File System filesets

Filesets
VRTSdcli.bff
VRTSvmpro.bff
VRTSvxfs.5.0.3.200-incr.bff
VRTSvxvm.05.00.0003.0200.bff

Storage Foundation Cluster File System filesets

[Table 1-18](#) describes the SFCFS filesets that are included in this rolling patch:

Table 1-18 SFCFS filesets

Filesets
VRTSaa.bff
VRTSccg.bff
VRTSdcli.bff
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSob.bff
VRTSobc33.bff
VRTSvcsag.rte.bff
VRTSvcs.rte.bff
VRTSvmpro.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.200-incr.bff
VRTSvxvm.05.00.0003.0200.bff

Storage Foundation for Oracle RAC filesets

[Table 1-19](#) describes the SF for Oracle RAC filesets that are included in this rolling patch:

Table 1-19 SF for Oracle RAC filesets

Filesets
VRTSaa.bff
VRTScavf.bff
VRTSccg.bff
VRTSdbac.rte.bff
VRTSdbed.bff
VRTSdbms3-3.0.83.4.bff
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSob.bff
VRTSobgui.bff
VRTSorgui.bff
VRTSvcsvr.bff
VRTSvcsvr.rte.bff
VRTSvcsvr.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.200-incr.bff
VRTSvxvm.05.00.0003.0200.bff

Storage Foundation filesets

[Table 1-20](#) describes the Storage Foundation filesets that are included in this rolling patch:

Table 1-20 Storage Foundation filesets

Filesets
VRTSaa.bff
VRTSccg.bff
VRTSdcli.bff
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSob.bff
VRTSobc33.bff
VRTSvcsag.rte.bff
VRTSvcs.rte.bff
VRTSvmpro.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.200-incr.bff
VRTSvxvm.05.00.0003.0200.bff

Storage Foundation for DB2 filesets

[Table 1-21](#) describes the Storage Foundation for DB2 filesets that are included in this rolling patch:

Table 1-21 Storage Foundation for DB2 filesets

Filesets
VRTSaa.bff
VRTSccg.bff
VRTSdbcom.bff
VRTSdbms3-3.0.83.3.bff
VRTSdcli.bff
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSob.bff
VRTSobc33.bff
VRTSvcsag.rte.bff
VRTSvcs.rte.bff
VRTSvmpro.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.200-incr.bff
VRTSvxvm.05.00.0003.0200.bff

Storage Foundation for Oracle filesets

[Table 1-22](#) describes the SF for Oracle filesets that are included in this rolling patch:

Table 1-22 Storage Foundation for Oracle filesets

Filesets
VRTSdbcom.bff
VRTSdbed.bff
VRTSdbms3-3.0.83.3.bff
VRTSdcli.bff
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSorgui.bff
VRTSvcsag.rte.bff
VRTSvcs.rte.bff
VRTSvmpro.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.200-incr.bff
VRTSvxvm.05.00.0003.0200.bff

Volume Manager filesets

[Table 1-23](#) describes the Volume Manager filesets that are included in this rolling patch:

Table 1-23 Volume Manager filesets

Filesets
VRTSdbms3-3.0.83.3.bff
VRTSdcli.bff
VRTSvcsvr.bff
VRTSvmman.05.00.0003.0200.bff
VRTSvmpro.bff
VRTSvxvm.05.00.0003.0200.bff

Note: The 5.0 MP3 release did not include an updated `VRTSvcsvr` fileset, which would have had the 5.0.3.0 version number. Because of that, the `VRTSvcsvr` fileset included in this release has the 5.0.3.0 version number.

Veritas Volume Replicator filesets

[Table 1-24](#) describes the Veritas Volume Replicator filesets that are included in this rolling patch:

Table 1-24 Veritas Volume Replicator filesets

Filesets
VRTSaa.bff
VRTSccg.bff
VRTSdcli.bff
VRTSob.bff
VRTSobc33.bff
VRTSvmpro.bff
VRTSvxvm.05.00.0003.0200.bff

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 RP2.

- [Installing Storage Foundation or Storage Foundation Cluster File System and 5.0 MP3 RP2](#)
- [Installing Storage Foundation for Oracle RAC and 5.0 MP3 RP2](#)

Installing Storage Foundation or Storage Foundation Cluster File System and 5.0 MP3 RP2

This section describes how to install Storage Foundation and Storage Foundation Cluster File System for the first time on a host and install 5.0 MP3 RP2.

Review the *Veritas Storage Foundation 5.0 MP3 Installation Guide* for pre-installation instructions at the following URL:

ftp://ftp.entsupport.symantec.com/pub/support/documentation/sf_install_aix.pdf

Also review the *Veritas Storage Foundation 5.0 MP3 Release Notes* for important release information at the following URL:

ftp://ftp.entsupport.symantec.com/pub/support/documentation/sf_notes_aix.pdf

To install the Storage Foundation or Storage Foundation Cluster File System and 5.0 MP3 RP2

- 1 Install the Veritas Storage Foundation or Storage Foundation Cluster File System 5.0 MP3 software from the release CD.

For example, if you are installing the Storage Foundation software, enter the following command from the top-level directory of the mounted CD:

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

The `-installonly` option is required to perform the installation without configuring the software. For other products, substitute the appropriate script for `installsf`, such as `installsfcfs` for the Storage Foundation Cluster File System software.

- 2 Review the installation prerequisites for upgrading to 5.0 MP3 RP2. See [“Prerequisites for upgrading to 5.0 MP3 RP2”](#) on page 54.

- 3 Copy the contents of the *product/patches* directory to local or shared storage.

```
# cp -r path-to-media/product/patches working_directory
```

where *path-to-media* is the path to the media.

where *product* is the product directory name such as *storage_foundation* or *sfcfs*.

where *working_directory* is your working directory.

- 4 Uncompress each of the bundled filesets.

```
# for bff in working_directory/*.gz; do gunzip $bff; done
```

- 5 Preview the pre-installation verification testing by performing the following steps:

- On one node, type:

```
# installp -pagXv -d working_directory all
```

Review the verification data and confirm that:

- No unexpected failures occur, and that
- All selected filesets are successfully verified.

- Repeat the `installp` command on all other nodes:

```
# installp -pagXv -d working_directory all
```

Symantec recommends that the local directory containing the unpacked patches is accessible from all nodes.

- Copy the list that appears under the `Selected Filesets` to a text file `filesets.to.install`. Save this file for later use.

```
*****
installp PREVIEW: installation will not actually occur.
*****
+-----+
                Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES
-----
Filesets listed in this section passed pre-installation verification
and will be installed.

Selected Filesets
-----
.
.
.
End of the Success section.
```


Installing Storage Foundation for Oracle RAC and 5.0 MP3 RP2

This section describes how to install Storage Foundation for Oracle RAC for the first time on a host and install 5.0 MP3 RP2.

Review the *Veritas Storage Foundation 5.0 MP3 Installation Guide* for pre-installation instructions at the following URL:

ftp://ftp.entsupport.symantec.com/pub/support/documentation/sfrac_install_aix.pdf

Also review the *Veritas Storage Foundation 5.0 MP3 Release Notes* for important release information at the following URL:

ftp://ftp.entsupport.symantec.com/pub/support/documentation/sfrac_notes_aix.pdf

To install Storage Foundation for Oracle RAC and 5.0 MP3 RP2

- 1 Install the SF Oracle RAC 5.0 MP3 software from the release CD.
For example, if you are installing the SF Oracle RAC software, enter the following command from the top-level directory of the mounted CD:

```
# ./installsffrac -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 RP2.
See “Prerequisites for upgrading to 5.0 MP3 RP2” on page 54.
- 3 Copy the contents of the `product/patches` directory to local or shared storage.

```
# cp -r path-to-media/product/patches working_directory
```

where `path-to-media` is the path to the media.
where `product` is the product directory name such as `sfrac`.
where `working_directory` is your working directory.
- 4 Uncompress each of the bundled filesets.

```
# for bff in working_directory/*.gz; do gunzip $bff; done
```
- 5 Preview the pre-installation verification testing by performing the following steps:
 - On one node, type:

```
# installp -pagXv -d working_directory all
```

Review the verification data and confirm that:
 - No unexpected failures occur, and that
 - All selected filesets are successfully verified.
 - Repeat the `installp` command on all other nodes:

```
# installp -pagXv -d working_directory all
```

Symantec recommends that the local directory containing the unpacked patches is accessible from all nodes.

- Copy the list that appears under the Selected Filesets to a text file `filesets.to.install`. Save this file for later use.

```
*****
installp PREVIEW:  installation will not actually occur.
*****
+-----+
                Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES
-----
Filesets listed in this section passed pre-installation verification
and will be installed.

Selected Filesets
-----
.
.
.
End of the Success section.

+-----+
                BUILDDATE Verification ...
+-----+
Verifying build dates...done
FILESET STATISTICS
-----
  n Selected to be installed, of which:
  n Passed pre-installation verification
----
  n Total to be installed

RESOURCES
-----
Estimated system resource requirements for filesets being installed:
(All sizes are in 512-byte blocks)

Filesystem           Needed Space           Free Space
/usr                  760                    20062440
TOTAL:                760                    20062440
.
.
.
End of installp PREVIEW.  No apply operation has actually occurred.
*****
```

6 Apply the patches to all or selected cluster nodes.

- To apply all patches, on each node, type:

```
# installp -agXv -d working_directory all
```

- To apply only patches listed in the file `filesets.to.install`, on each node, type:

```
# installp -agXv -d working_directory -f filesets.to.install
```

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

7 Restart all the cluster nodes.

```
# shutdown -r now
```

8 Run the same installation script that you used in [step 1](#), this time specifying the `-configure` option to configure the software. For example,

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

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

See the *Veritas Storage Foundation for Oracle RAC 5.0 MP3 Installation Guide* for more information on configuring Storage Foundation and High Availability products.

Prerequisites for upgrading to 5.0 MP3 RP2

The following list describes prerequisites for upgrading to the 5.0 MP3 RP2 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 RP2 release.
- Each system must have sufficient free space to accommodate patches.

Upgrading 5.0 MP3 to 5.0 MP3 RP2

This section describes how to upgrade from 5.0 MP3 to 5.0 MP3 RP2 on a cluster or a standalone system.

- [Performing a phased upgrade to 5.0 MP3 RP2 on a cluster](#)
Use the procedures to perform a phased upgrade to 5.0 MP3 RP2 on a cluster that has SFHA, SF for Oracle HA or SF for DB2 HA, SFCFS, or Storage Foundation for Oracle RAC installed and configured.
- [Performing a full upgrade to 5.0 MP3 RP2 on a cluster](#)
Use the procedures to perform a full upgrade to 5.0 MP3 RP2 on a cluster that has VCS, SFHA, SF for Oracle HA or SF for DB2 HA, SFCFS, or Storage Foundation for Oracle RAC installed and configured.
- [Upgrading to 5.0 MP3 RP2 on a standalone system](#)
Use the procedure to upgrade to 5.0 MP3 RP2 on a system that has Storage Foundation, SF for Oracle, or SF for DB2 installed.

Performing a phased upgrade to 5.0 MP3 RP2 on a cluster

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

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

- 1 Freeze service group operations and stop VCS on the cluster.
- 2 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).
- 3 Take offline the nodes in group A and install the software patches on those nodes.
- 4 Take offline the nodes in group B and bring online the nodes in group A to restart cluster failover services.
- 5 Upgrade the nodes in group B, then bring those nodes online to join. The cluster is fully restored.

Depending on your cluster's configuration, select one of the following procedures to upgrade to 5.0 MP3 RP2:

- [Performing a phased upgrade to 5.0 MP3 RP2 on a Storage Foundation HA cluster](#)
- [Performing a phased upgrade to 5.0 MP3 RP2 on a Storage Foundation Cluster File System cluster](#)
- [Performing a phased upgrade to 5.0 MP3 RP2 on a Storage Foundation for Oracle RAC cluster](#)

Note: Symantec does not support a phased upgrade of a VCS cluster. Symantec only supports a full upgrade of a VCS cluster.

See "[Performing a full upgrade to 5.0 MP3 RP2 on a VCS cluster](#)" on page 69.

Performing a phased upgrade to 5.0 MP3 RP2 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 RP2 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 service group to another node that is running.

```
# hagr -switch service_group -to nodename
```
- 4 Make the VCS configuration writable on a node that is being upgraded:

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

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

```
# haconf -dump makero
```
- 7 Close any instance of VCS GUI that is running on the node.
- 8 Select the group of nodes that are to be upgraded first, and follow [step 9](#) through [step 22](#) for these nodes.
- 9 Stop VCS on each node in the selected group:

```
# hastop -local
```
- 10 Stop the VCS command server:

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

where `pid_of_CmdServer` is the process ID of `CmdServer`.
- 11 Stop cluster fencing, GAB, and LLT.

```
# /etc/rc.d/rc2.d/S97vxfen stop  
# /etc/rc.d/rc2.d/S92gab stop  
# /etc/rc.d/rc2.d/S701lt stop
```
- 12 If required, you can upgrade the operating system on the nodes in the selected group at this stage and patch them to a supported kernel version. See [“System requirements”](#) on page 8.
- 13 Repeat [step 9](#) through [step 11](#) if the system reboots after upgrading the operating system. You need to perform this to stop the components that started by the init scripts, if any.
- 14 Copy the contents of the `product/patches` directory to local or shared storage.

```
# cp -r path-to-media/product/patches working_directory
```

where `path-to-media` is the path to the media.
where `product` is the product directory name such as `storage_foundation_for_db2` or `storage_foundation_for_oracle`.
where `working_directory` is your working directory.


```
      n Total to be installed
RESOURCES
-----
Estimated system resource requirements for filesets being installed:
      (All sizes are in 512-byte blocks)
Filesystem           Needed Space           Free Space
-----
/usr                  760                    20062440
TOTAL:                760                    20062440
.
.
.
End of installp PREVIEW.  No apply operation has actually occurred.
*****
```

- 17 Apply the patches to all or selected cluster nodes.
 - To apply all patches, on each node, type:
installp -agXv -d *working_directory* all
 - To apply only patches listed in the file *filesets.to.install*, on each node, type:
installp -agXv -d *working_directory* -f *filesets.to.install*Review the summaries at the end of each run and confirm that all of the intended patches were successfully applied.
- 18 After all of the nodes in the selected group are upgraded, shut down and reboot each of the nodes. After the nodes come up, application failover capability is available for that group of nodes.
- 19 Make the VCS configuration writable again from any node in the selected group:
haconf -makerw
- 20 Unfreeze the service group operations on each node for which you upgraded the operating system:
hasys -unfreeze -persistent *nodename*
- 21 Make the VCS configuration read-only:
haconf -dump -makero
- 22 Switch the service group to the original node:
hagrp -switch *service_group* -to *nodename*
- 23 Repeat [step 9](#) through [step 22](#) for the second group of nodes.

- 24 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 RP2 installation.
- For more information see the “[Software limitations](#)” on page 36 about older backups failing to be restored using the DBED scripts.
- See the *Veritas Storage Foundation for Oracle Administrator's Guide* for the `sfa_rept_adm` command.
- For more information see “[Storage Foundation for Oracle fixed issues](#)” on page 20 or “[Storage Foundation for DB2 fixed issues](#)” on page 22 for incident 1425261.

Performing a phased upgrade to 5.0 MP3 RP2 on a Storage Foundation Cluster File System cluster

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

To perform a phased upgrade to 5.0 MP3 RP2 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, switch the service group to another node that is running.

```
# hagrps -switch service_group -to nodename
```
- 4 From any node in the cluster, make the VCS configuration writable:

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

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

```
# haconf -dump -makero
```
- 7 Select a group of nodes that are to be upgraded first, and follow [step 8](#) through [step 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
```

- a 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:

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

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

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

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

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

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

- 11 Stop 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 On each node in the selected group, 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
```

- 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 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/S701lt stop
```

16 If required, you can upgrade the operating system of nodes in the selected group at this stage, and patch the nodes to a supported kernel version. See “[System requirements](#)” on page 8.

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 Copy the contents of the *product/patches* directory to local or shared storage.

`# cp -r path-to-media/product/patches working_directory`

where *path-to-media* is the path to the media.

where *product* is the product directory name such as *sfcfs*.

where *working_directory* is your working directory.

19 Uncompress each of the bundled filesets.

`# for bff in working_directory/*.gz; do gunzip $bff; done`

20 Preview the pre-installation verification testing by performing the following steps:

- On one node, type:

`# installp -pagXv -d working_directory all`

Review the verification data and confirm that:

- No unexpected failures occur.
- All selected filesets are successfully verified.
- Repeat the `installp` command on all other nodes:
`# installp -pagXv -d working_directory all`
 Symantec recommends that the local directory containing the unpacked patches is accessible from all nodes.
- Copy the list that appears under the Selected Filesets to a text file `filesets.to.install`. Save this file for later use.

```
*****
installp PREVIEW: installation will not actually occur.
*****
+-----+
                Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES
-----
Filesets listed in this section passed pre-installation verification
and will be installed.
```

```

Selected Filesets
-----
.
.
.
End of the Success section.
+-----+
                BUILDDATE Verification ...
+-----+
Verifying build dates...done
FILESET STATISTICS
-----
    n Selected to be installed, of which:
    n Passed pre-installation verification
----
    n Total to be installed

RESOURCES
-----
Estimated system resource requirements for filesets being installed:
      (All sizes are in 512-byte blocks)
      Filesystem           Needed Space           Free Space
      /usr                  760                   20062440
      TOTAL:                760                   20062440
      .
      .
      .
End of installp PREVIEW.  No apply operation has actually occurred.
*****

```

21 Apply the patches to the nodes in the selected group.

- To apply all patches, on each node, type:
installp -agXv -d *working_directory* all
- To apply only patches listed in the file `filesets.to.install`, on each node, type:
installp -agXv -d *working_directory* -f *filesets.to.install*

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

22 On the second group of nodes, stop the failover service group. Downtime starts for failover service groups.

```
# hagrp -offline failover_service_group
```

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

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

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

```
# haconf -makerw
```
- 26 Enter the following command on each node in the selected group to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent nodename
```
- 27 Make the configuration read-only:

```
# haconf -dump -makero
```
- 28 Switch the service group to the original node:

```
# hagrps -switch service_group -to nodename
```
- 29 Bring the CVM service group online on each node in the selected group:

```
# hagrps -online cvm -sys nodename
```
- 30 Restart 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:

```
# vxrvgs -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 8](#) through [step 33](#) for the second group of nodes.

Performing a phased upgrade to 5.0 MP3 RP2 on a Storage Foundation for Oracle RAC cluster

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

To upgrading to 5.0 MP3 RP2 on a 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 Switch the service group to another node that is running.

```
# hagrps -switch service_group -to nodename
```
- 4 Select the group of nodes that are to be upgraded first, and follow [step 5](#) through [step 23](#) for these nodes.

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

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

On stopping CRS if any gsd relevant process remains active, you must stop that process manually.
- 6 All Oracle users must stop all applications using the CFS mounts not under VCS control.
 - a To verify that no processes use the CFS mount point, enter the following command:

```
# fuser -c mount_point
```
 - b Stop any processes using a CFS mount point with the mechanism provided by the application.
- 7 All Oracle users must unmount any CFS file systems not under VCS control on all nodes. To unmount CFS file systems not under VCS control:
 - a Log in as root.
 - b Remove the existing unused modules in the kernel memory and in the library memory.

```
# /usr/sbin/slibclean
```
 - c Determine the file systems to unmount by checking the output of the mount file. For example:

```
# mount | grep vxfs | grep cluster
```
 - d By specifying the mount point for the file system, unmount each file system that is listed in the output:

```
# umount mount_point
```
- 8 Stop VCS on each node in the group being upgraded:

```
# hastop -local
```
- 9 Stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of *CmdServer*.
- 10 Stop VCSMM and LMX if they are running:

```
# /etc/init.d/vcsmm.rc stop  
# /etc/init.d/lmx.rc stop
```
- 11 Unregister CFS from GAB.

```
# fsclustadm cfsdeinit
```

12 Stop cluster fencing, ODM, and GAB:

```
# /etc/rc.d/rc2.d/S97vxfen stop
# /etc/rc.d/rc2.d/S99odm stop
# /etc/rc.d/rc2.d/S92gab stop
# /etc/rc.d/rc2.d/S701lt stop
```

13 If required, you can upgrade the operating system of the nodes of the sub-cluster at this stage, and patch them to a supported kernel version.

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

See “[System requirements](#)” on page 8.

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

```
# mount | grep vxfs
```

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

```
# umount /mount-point
```

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

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

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

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

```
# fsck -V vxfs filesystem-device
# mount -V vxfs filesystem-device mountpoint
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.

There may be a pending large fileset clone removal extended operation if the `umount` command fails with the following error:

```
file system device busy
```

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

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

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

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

- 15 Stop activity to all VxVM volumes.

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

- 16 On each node of the sub-cluster, stop all VxVM volumes by entering the following command for each disk group:

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

Verify that no volumes remain open:

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

- 17 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

- 18 Copy the contents of the *product/patches* directory to local or shared storage.

```
# cp -r path-to-media/product/patches working_directory
```

where *path-to-media* is the path to the media.

where *product* is the product directory name such as *sfrac*.

where *working_directory* is your working directory.

- 19 Uncompress each of the bundled filesets.

```
# for bff in working_directory/*.gz; do gunzip $bff; done
```

- 20 Apply all patches on each node:

```
# installp -agXv -d working_directory all
```

- 21 After all of the nodes in the sub-cluster are upgraded, repeat through [step 20](#) to upgrade the remaining part of the cluster.

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

```
# shutdown -r now
```

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

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

- ```
hasys -unfreeze -persistent nodename
```
- 24 Switch the service group to the original node:
- ```
# hagr -switch service_group -to nodename
```
- 25 On the second group of nodes, stop the failover service group. Downtime starts for failover service groups.
- ```
hagr -offline failover_service_group
```
- 26 Make the VCS configuration writable again from any node second group:
- ```
# haconf -makerw
```
- 27 Enter the following command on each node in the second group to freeze HA service group operations for failover service group:
- ```
hasys -freeze -persistent nodename
```
- 28 Make the configuration read-only:
- ```
# haconf -dump -makero
```
- 29 On the second group of nodes, shutdown VCS. Execute following command on one of the nodes.
- ```
hastop -all -force
```
- 30 On the first group of nodes, start VCS on each of the nodes:
- ```
# hastart
```
- Downtime ends for the failover service groups.
- 31 If CRS is not controlled by VCS, use the following command on each node to start CRS.
- ```
/etc/init.crs start
```
- 32 Check if the VEA service restarted:
- ```
# /opt/VRTS/bin/vxsvcctl status
```
- If the VEA service is not running, restart it:
- ```
/opt/VRTS/bin/vxsvcctl start
```
- 33 After all of the services are up on the first group of nodes, stop VCS. Execute following command on one of the nodes.
- ```
# hastop -all -force
```
- 34 On the second group of nodes, start VCS on each of the nodes
- ```
hastart
```
- 35 Repeat [step 5](#) through [step 21](#) for the second group of nodes.
- 36 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node in second group.

- 37 If CRS is not controlled by VCS, use the following command on each node in second group to start CRS:  

```
/etc/init.crs start
```
- 38 Check if the VEA service was restarted:  

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

If the VEA service is not running, restart it:

```
/opt/VRTS/bin/vxsvcctl start
```
- 39 Make the VCS configuration writable again from any node second group:  

```
haconf -makerw
```
- 40 Enter the following command on each node in the second group to unfreeze HA service group operations for failover service group:  

```
hasys -unfreeze -persistent nodename
```
- 41 Make the configuration read-only:  

```
haconf -dump -makero
```
- 42 On the first group of nodes, start VCS on each of the nodes  

```
hastart
```
- 43 Relink Oracle's CRS and database libraries for Storage Foundation for Oracle RAC:
  - a Run the following command:  

```
/opt/VRTS/install/installsfrac -configure
```
  - b Choose the correct relinking option for your version of Oracle:
    - Relink Storage Foundation for Oracle RAC for Oracle 9i
    - Relink Storage Foundation for Oracle RAC for Oracle 10g Release 1
    - Relink Storage Foundation for Oracle RAC for Oracle 10g Release 2
    - Relink Storage Foundation for Oracle RAC for Oracle 11g
- 44 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 RP2 installation.  

For more information see the “[Software limitations](#)” on page 36 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 20 or “[Storage Foundation for DB2 fixed issues](#)” on page 22 for incident 1425261.

## Performing a full upgrade to 5.0 MP3 RP2 on a cluster

Performing a full 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 RP2:

- [Performing a full upgrade to 5.0 MP3 RP2 on a VCS cluster](#)
- [Performing a full upgrade to 5.0 MP3 RP2 on a Storage Foundation HA cluster](#)
- [Performing a full upgrade to 5.0 MP3 RP2 on a Storage Foundation Cluster File System cluster](#)
- [Performing a phased upgrade to 5.0 MP3 RP2 on a Storage Foundation for Oracle RAC cluster](#)

---

**Note:** Symantec does not support a phased upgrade of a VCS cluster.

---

### Performing a full upgrade to 5.0 MP3 RP2 on a VCS cluster

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

#### To perform pre-upgrade tasks

- 1 Copy the contents of the product directory `cluster_server/patches` to local or shared storage.  

```
cp -r path-to-media/cluster_server/patches working_directory
```
- 2 Uncompress each of the bundled filesets.  

```
for bff in working_directory/*.gz; do gunzip $bff; done
```
- 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 service groups except ClusterService. On any node, type:

```
hagr -freeze $grp -persistent
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 a shutdown of VCS cluster and drivers

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

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

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

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

- 3 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.

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

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

- 5 Halt and unload vxfen. On each node, perform the following steps:

- Unload the vxfen driver:

```
/etc/methods/vxfenext -stop
```

- Confirm that the vxfen driver is unloaded:

```
sleep 3; /etc/methods/vxfenext -status
vxfen: unloaded
```

- 6 Unconfigure GAB. On each node, type:  
`# /sbin/gabconfig -U`
- 7 Halt and unload the GAB driver. On each node, perform the following steps:
  - Unload the GAB kernel module:  
`# /etc/methods/gabkext -stop`
  - Confirm that the GAB driver is unloaded:  
`# sleep 3; /etc/methods/gabkext -status`  
`gab: unloaded`
- 8 Unconfigure LLT. On each node, type:  
`# /sbin/lltconfig -U`
- 9 Unload the LLT portable streams driver. On each node, type:  
`# /usr/sbin/strload -u -d /usr/lib/drivers/pse/llt`

---

**Note:** The unloading of LLT is broken for this release (Etrack 1111685). As a result, all operations that require the unloading of LLT require that you reboot the nodes. Reloading LLT and dependent drivers is not viable.

---

### To preview and apply the upgrade

- 1 Preview the pre-installation verification testing by performing the following steps:
  - On one node, type:  
`# installp -pagXv -d working_directory all`  
 Review the verification data and confirm that:
    - No unexpected failures occur, and that
    - All selected filesets are successfully verified.
  - Repeat the installp command on all other nodes:  
`# installp -pagXv -d working_directory all`  
 Symantec recommends that the local directory containing the unpacked patches is accessible from all nodes.
  - Copy the list of successful filesets (in bold below) to a text file filesets.to.install. Save this file for later use.

```

installp PREVIEW: installation will not actually occur.

+-----+
 Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES
```

```

Filesets listed in this section passed pre-installation verification
and will be installed.

Selected Filesets

VRTSgab.rte 5.0.3.200 # Veritas Group Membership and...
VRTSllt.rte 5.0.3.200 # Veritas Low Latency Transpor...
VRTSvc.s.rte 5.0.3.200 # Veritas Cluster Server 5.0MP...
VRTSvc.sag.rte 5.0.3.200 # Veritas Cluster Server 5.0MP...
VRTSvxfen.rte 5.0.3.200 # Veritas I/O Fencing 5.0MP3RP...

End of the Success section.
+-----+
 BUILDDATE Verification ...
+-----+
Verifying build dates...done
FILESET STATISTICS

 5 Selected to be installed, of which:
 5 Passed pre-installation verification

 5 Total to be installed

RESOURCES

Estimated system resource requirements for filesets being installed:
(All sizes are in 512-byte blocks)

Filesystem Needed Space Free Space
/usr 760 20062440
TOTAL: 760 20062440

(...omitted...)

End of installp PREVIEW. No apply operation has actually occurred.

```

- 2 Apply the patches to all or selected cluster nodes.
  - To apply all patches, on each node, type:  
# **installp -agXv -d *working\_directory* all**
  - To apply only patches listed in the file `filesets.to.install`, on each node, type:  
# **installp -agXv -d *working\_directory* -f *filesets.to.install***Review the summaries at the end of each run and confirm that all of the intended patches were successfully applied.

### 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
hagr -list | sort -u +0b -1 | \
 while read grp sys ; do
 hagr -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:  

```
hagr -online ClusterService -sys system
```

where `system` is the node name.

## Performing a full upgrade to 5.0 MP3 RP2 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 RP2 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 Close any instance of VCS GUI that is running on the node.
- 7 Stop VCS:  

```
hastop -local
```

- 8 Stop the VCS command server:

```
ps -ef | grep CmdServer
```

```
kill -9 pid_of_CmdServer
```

where *pid\_of\_CmdServer* is the process ID of CmdServer.
- 9 Stop cluster fencing, GAB, and LLT.

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

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

```
/etc/rc.d/rc2.d/S701lt stop
```
- 10 If required, you can upgrade the operating system on the nodes at this stage and patch them to a supported kernel version.  
See “[System requirements](#)” on page 8.
- 11 Repeat [step 7](#) through [step 9](#) if the system reboots after upgrading the operating system. You need to perform this to stop the components that started by the init scripts, if any.
- 12 Copy the contents of the *product/patches* directory to local or shared storage.

```
cp -r path-to-media/product/patches working_directory
```

where *path-to-media* is the path to the media.  
where *product* is the product directory name such as *storage\_foundation\_for\_db2* or *storage\_foundation\_for\_oracle*.  
where *working\_directory* is your working directory.
- 13 Uncompress each of the bundled filesets.

```
for bff in working_directory/*.gz; do gunzip $bff; done
```
- 14 Preview the pre-installation verification testing by performing the following steps:
  - On one node, type:

```
installp -pagXv -d working_directory all
```

Review the verification data and confirm that:
    - No unexpected failures occur, and that
    - All selected filesets are successfully verified.
  - Repeat the `installp` command on all other nodes:

```
installp -pagXv -d working_directory all
```

Symantec recommends that the local directory containing the unpacked patches is accessible from all nodes.

- Copy the list that appears under the Selected Filesets to a text file filesets.to.install. Save this file for later use.

```

installp PREVIEW: installation will not actually occur.

+-----+
 Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES

 Filesets listed in this section passed pre-installation verification
 and will be installed.

 Selected Filesets

 .
 .
 .
End of the Success section.

+-----+
 BUILDDATE Verification ...
+-----+
Verifying build dates...done
FILESET STATISTICS

 n Selected to be installed, of which:
 n Passed pre-installation verification

 n Total to be installed
RESOURCES

 Estimated system resource requirements for filesets being installed:
 (All sizes are in 512-byte blocks)
 Filesystem Needed Space Free Space
 /usr 760 20062440
 TOTAL: 760 20062440
 .
 .
 .
End of installp PREVIEW. No apply operation has actually occurred.

```

- 15 Apply the patches to all or selected cluster nodes.
  - To apply all patches, on each node, type:  

```
installp -agXv -d working_directory all
```
  - To apply only patches listed in the file `filesets.to.install`, on each node, type:  

```
installp -agXv -d working_directory -f filesets.to.install
```Review the summaries at the end of each run and confirm that all of the intended patches were successfully applied.
- 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 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 RP2 installation.  
For more information see the [“Software limitations”](#) on page 36 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 20 or [“Storage Foundation for DB2 fixed issues”](#) on page 22 for incident 1425261.

## Performing a full upgrade to 5.0 MP3 RP2 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 RP2 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 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:  

```
mount | grep 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
```

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

```
vxrvrg -g diskgroup stop rvg_name
```
  - c 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:

```
hstop -local
```

- 12 Stop the VCS command server:

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

where *pid\_of\_CmdServer* is the process ID of CmdServer.

- 13 Stop cluster 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
```

- 14 If required, you can upgrade the operating system of the nodes at this stage, and patch them to a supported kernel version.

See “[System requirements](#)” on page 8.

- 15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

- 16 Copy the contents of the *product/patches* directory to local or shared storage.

```
cp -r path-to-media/product/patches working_directory
```

where *path-to-media* is the path to the media.

where *product* is the product directory name such as *sfcfs*.

where *working\_directory* is your working directory.

- 17 Uncompress each of the bundled filesets.

```
for bff in working_directory/*.gz; do gunzip $bff; done
```

- 18 Preview the pre-installation verification testing by performing the following steps:

- On one node, type:

```
installp -pagXv -d working_directory all
```

Review the verification data and confirm that:

- No unexpected failures occur.
- All selected filesets are successfully verified.
- Repeat the `installp` command on all other nodes:  

```
installp -pagXv -d working_directory all
```

Symantec recommends that the local directory containing the unpacked patches is accessible from all nodes.
- Copy the list that appears under the Selected Filesets to a text file `filesets.to.install`. Save this file for later use.

```

installp PREVIEW: installation will not actually occur.

+-----+
 Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES

 Filesets listed in this section passed pre-installation verification
 and will be installed.

 Selected Filesets

 .
 .
 .
End of the Success section.
+-----+
 BUILDDATE Verification ...
+-----+
Verifying build dates...done
FILESET STATISTICS

 n Selected to be installed, of which:
 n Passed pre-installation verification

 n Total to be installed

RESOURCES

 Estimated system resource requirements for filesets being installed:
 (All sizes are in 512-byte blocks)

 Filesystem Needed Space Free Space
 /usr 760 20062440
 TOTAL: 760 20062440

 .
 .
 .
End of installp PREVIEW. No apply operation has actually occurred.

```

- 19 Apply the patches to all or selected cluster nodes.
  - To apply all patches, on each node, type:  

```
installp -agXv -d working_directory all
```
  - To apply only patches listed in the file `filesets.to.install`, on each node, type:  

```
installp -agXv -d working_directory -f filesets.to.install
```Review the summaries at the end of each run and confirm that all of the intended patches were successfully applied.
- 20 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.
- 21 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node.
- 22 Make the VCS configuration writable again from any node:  

```
haconf -makerw
```
- 23 Enter the following command on each node to unfreeze HA service group operations:  

```
hasys -unfreeze -persistent nodename
```
- 24 Make the configuration read-only:  

```
haconf -dump -makero
```
- 25 Bring the CVM service group online on each node:  

```
hagrps -online cvm -sys nodename
```
- 26 Restart all the volumes by entering the following command for each disk group:  

```
vxvol -g diskgroup startall
```
- 27 If you stopped any RVGs in [step 10](#), restart each RVG:  

```
vxrvrg -g diskgroup start rvg_name
```
- 28 Remount all VxFS file systems on all nodes:  

```
mount /filesystem
```
- 29 Remount all Storage Checkpoints on all nodes:  

```
mount /checkpoint_name
```

## Performing a full upgrade to 5.0 MP3 RP2 on a Storage Foundation for Oracle RAC cluster

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

### To upgrading to 5.0 MP3 RP2 on a 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 If CRS is not controlled by VCS, use the following command on each node to stop CRS.  

```
/etc/init.crs stop
```

On stopping CRS if any gsd relevant process remains active, you must stop that process manually.
- 4 All Oracle users must stop all applications using the CFS mounts not under VCS control.
  - a Verify that no processes are using the CFS mount point:  

```
fuser -c mount_point
```
  - b Stop any processes using a CFS mount point with the mechanism provided by the application.
- 5 All Oracle users must unmount any CFS file systems not under VCS control on all nodes. To unmount CFS file systems not under VCS control:
  - a Log in as root.
  - b Remove the existing unused modules in the kernel memory and in the library memory.  

```
/usr/sbin/slibclean
```
  - c Determine the file systems to unmount by checking the output of the mount file. For example:  

```
mount | grep vxfs | grep cluster
```
  - d By specifying the mount point for the file system, unmount each file system that is listed in the output:  

```
umount mount_point
```
- 6 Stop VCS:  

```
hastop -local
```
- 7 Stop the VCS command server:  

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

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

- 8 Stop VCSMM and LMX if they are running:

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

- 9 Unregister CFS from GAB.

```
fsclustadm cfsdeinit
```

- 10 Stop cluster fencing, ODM, and GAB:

```
/etc/rc.d/rc2.d/S97vxfen stop
/etc/rc.d/rc2.d/S99odm stop
/etc/rc.d/rc2.d/S92gab stop
/etc/rc.d/rc2.d/S701lt stop
```

- 11 If required, you can upgrade the operating system of the nodes of the sub-cluster at this stage, and patch them to a supported kernel version.

---

**Note:** If you are upgrading a Storage Foundation for Oracle RAC cluster, you must upgrade the nodes of the sub-cluster at this stage to one of the operating system versions that this RP release supports.

---

See “[System requirements](#)” on page 8.

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

```
mount | grep vxfs
```

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

```
umount /mount-point
```

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

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

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

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

```
fsck -V vxfs filesystem-device
mount -V vxfs filesystem-device mountpoint
umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.

There may be a pending large fileset clone removal extended operation if the `umount` command fails with the following error:

```
file system device busy
```

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

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

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

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

- 13 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.

- 14 On each node of the sub-cluster, stop all VxVM volumes by entering the following command for each disk group:

```
vxvol -g diskgroup stopall
```

Verify that no volumes remain open:

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

- 15 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

- 16 Copy the contents of the *product/patches* directory to local or shared storage.

```
cp -r path-to-media/product/patches working_directory
```

where *path-to-media* is the path to the media.

where *product* is the product directory name such as *sfrac*.

where *working\_directory* is your working directory.

- 17 Uncompress each of the bundled filesets.

```
for bff in working_directory/*.gz; do gunzip $bff; done
```

- 18 Apply all patches on each node:

```
installp -agXv -d working_directory all
```

- 19 After all of the nodes in the sub-cluster are upgraded, repeat through [step 20](#) to upgrade the remaining part of the cluster.

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

```
shutdown -r now
```

- 20 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node.
- 21 Enter the following command on each node to unfreeze HA service group operations:  

```
hasys -unfreeze -persistent nodename
```
- 22 Start VCS on each of the nodes:  

```
hastart
```

Downtime ends for the failover service groups.
- 23 If CRS is not controlled by VCS, use the following command on each node to start CRS.  

```
/etc/init.crs start
```
- 24 Check if the VEA service restarted:  

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

If the VEA service is not running, restart it:  

```
/opt/VRTS/bin/vxsvcctl start
```
- 25 After all of the services are up, stop VCS. Execute following command on one of the nodes.  

```
hstop -all -force
```
- 26 Start VCS on each of the nodes:  

```
hastart
```
- 27 Relink Oracle's CRS and database libraries for Storage Foundation for Oracle RAC:
  - a Run the following command:  

```
/opt/VRTS/install/installsfrac -configure
```
  - b Choose the correct relinking option for your version of Oracle:
    - Relink Storage Foundation for Oracle RAC for Oracle 9i
    - Relink Storage Foundation for Oracle RAC for Oracle 10g Release 1
    - Relink Storage Foundation for Oracle RAC for Oracle 10g Release 2
    - Relink Storage Foundation for Oracle RAC for Oracle 11g

- 28** 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 RP2 installation.
- For more information see the [“Software limitations”](#) on page 36 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 20 or [“Storage Foundation for DB2 fixed issues”](#) on page 22 for incident 1425261.

## Upgrading to 5.0 MP3 RP2 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 upgrading to 5.0 MP3 RP2 on a standalone system

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
- 3 If required, you can upgrade the system at this stage, and patch it to a supported kernel version.
- 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:
  - a Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
  - b Use the `vrxvg stop` command to stop each RVG individually:

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

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

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

---

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

---

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

```
vxvol -g diskgroup stopall
```

Verify that no volumes remain open:

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

## 9 Check if the VEA service is running:

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

If the VEA service is running, stop it:

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

10 Copy the contents of the *product/patches* directory to local or shared storage.

```
cp -r path-to-media/product/patches working_directory
```

where *path-to-media* is the path to the media.

where *product* is the product directory name such as *storage\_foundation*.

where *working\_directory* is your working directory.

## 11 Uncompress each of the bundled filesets.

```
for bff in working_directory/*.gz; do gunzip $bff; done
```

## 12 Preview the pre-installation verification testing by performing the following steps:

## ■ On one node, type:

```
installp -pagXv -d working_directory all
```

Review the verification data and confirm that:

## ■ No unexpected failures occur, and that

## ■ All selected filesets are successfully verified.

■ Repeat the `installp` command on all other nodes:

```
installp -pagXv -d working_directory all
```

Symantec recommends that the local directory containing the unpacked patches is accessible from all nodes.

■ Copy the list that appears under the Selected Filesets to a text file `filesets.to.install`. Save this file for later use.

```

installp PREVIEW: installation will not actually occur.

+-----+
+-----+
Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES

Filesets listed in this section passed pre-installation verification
and will be installed.
```

```

Selected Filesets

.
.
.
End of the Success section.
+-----+
 BUILDDATE Verification ...
+-----+
Verifying build dates...done
FILESET STATISTICS

 n Selected to be installed, of which:
 n Passed pre-installation verification

 n Total to be installed

RESOURCES

Estimated system resource requirements for filesets being installed:
 (All sizes are in 512-byte blocks)
Filesystem Needed Space Free Space
/usr 760 20062440
TOTAL: 760 20062440
.
.
.
End of installp PREVIEW. No apply operation has actually occurred.

```

**13** Apply the patches to all or selected cluster nodes.

- To apply all patches, on each node, type:  
# **installp -agXv -d *working\_directory* all**
  - To apply only patches listed in the file `filesets.to.install`, on each node, type:  
# **installp -agXv -d *working\_directory* -f *filesets.to.install***
- Review the summaries at the end of each run and confirm that all of the intended patches were successfully applied.

**14** Shut down and restart the system.

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

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

```
vxvol -g diskgroup startall
```

**17** If you stopped any RVGs in [step 6](#), restart each RVG:

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

18 Remount all VxFS file systems and Storage Checkpoints:

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

19 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

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 RP2 installation.

For more information see the [“Software limitations”](#) on page 36 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 20 or [“Storage Foundation for DB2 fixed issues”](#) on page 22 for incident 1425261.

## Verifying software versions

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

```
ls1pp -L VRTS*
```

## Removing 5.0 MP3 RP2

Roll back of the 5.0 MP3 RP2 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 RP2 to the release 5.0 MP3 version for Veritas Cluster Server.

- [Removing 5.0 MP3 RP2 from Veritas Cluster Server](#)
- [Removing 5.0 MP3 RP2 on Storage Foundation or Storage Foundation Cluster File System](#)
- [Removing 5.0 MP3 RP2 on Storage Foundation for Oracle RAC](#)

## Removing 5.0 MP3 RP2 from Veritas Cluster Server

Use the following procedure to remove VCS 5.0 MP3 RP2 from your cluster manually.

### To remove 5.0 MP3 RP2 from VCS manually

- 1 Verify that all of the VCS 5.0MP3RP2 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 5.0.3.200
VRTS11t.rte 5.0.3.200
VRTSvcs.rte 5.0.3.200
VRTSvcsag.rte 5.0.3.200
VRTSvxfen.rte 5.0.3.200
```

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

```
nohdr='^Z$'
while read pkg ver; do
 lsipp -l $pkg | egrep -v "$nohdr"
 nohdr='^ Fileset +Level State '
done < fileset.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<br>VRTSgab.rte    | 5.0.3.200 | APPLIED | Veritas Group Membership and Atomic Broadcast 5.0MP3RP2 by Symantec 12/15/2008-11:49:23 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505 |
| Path: /etc/objrepos<br>VRTSgab.rte        | 5.0.3.200 | APPLIED | Veritas Group Membership and Atomic Broadcast 5.0MP3RP2 by Symantec 12/15/2008-11:49:23 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505 |
| -----                                     |           |         |                                                                                                                                          |
| Path: /usr/lib/objrepos<br>VRTSllt.rte    | 5.0.3.200 | APPLIED | Veritas Low Latency Transport 5.0MP3RP2 by Symantec 12/15/2008-11:49:11 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505                 |
| Path: /etc/objrepos<br>VRTSllt.rte        | 5.0.3.200 | APPLIED | Veritas Low Latency Transport 5.0MP3RP2 by Symantec 12/15/2008-11:49:11 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505                 |
| -----                                     |           |         |                                                                                                                                          |
| Path: /usr/lib/objrepos<br>VRTSvc.s.rte   | 5.0.3.200 | APPLIED | Veritas Cluster Server 5.0MP3RP2 by Symantec 12/15/2008-11:48:10 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505                        |
| Path: /etc/objrepos<br>VRTSvc.s.rte       | 5.0.3.200 | APPLIED | Veritas Cluster Server 5.0MP3RP2 by Symantec 12/15/2008-11:48:10 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505                        |
| -----                                     |           |         |                                                                                                                                          |
| Path: /usr/lib/objrepos<br>VRTSvc.sag.rte | 5.0.3.200 | APPLIED | Veritas Cluster Server 5.0MP3RP2 Bundled Agents by Symantec                                                                              |

```

Path: /usr/lib/objrepos
VRTSvxfen.rte 5.0.3.200 APPLIED Veritas I/O Fencing 5.0MP3RP2
 by Symantec
 12/15/2008-11:49:40 Build
 Env:AIX ogma 3 5 0001D4CAD300
 IBM,9115-505
```

```
Path: /etc/objrepos
VRTSvxfen.rte 5.0.3.200 APPLIED Veritas I/O Fencing 5.0MP3RP2
 by Symantec
 12/15/2008-11:49:40 Build
 Env:AIX ogma 3 5 0001D4CAD300
 IBM,9115-505
```

- 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 -freeze $grp -persistent
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 Unconfigure vxfen if the VCS cluster uses the fencing option. On each node, type:  

```
/sbin/vxfenconfig -U
```
- 15 Halt and unload vxfen. On each node, perform the following steps:
  - Unload the vxfen driver:  

```
/etc/methods/vxfenext -stop
```
  - Confirm that the vxfen driver is unloaded:  

```
sleep 3; /etc/methods/vxfenext -status
```

vxfen: unloaded
- 16 Unconfigure GAB. On each node, type:  

```
/sbin/gabconfig -U
```
- 17 Halt and unload the GAB driver. On each node, perform the following steps:
  - Unload the GAB kernel module:  

```
/etc/methods/gabkext -stop
```
  - Confirm that the GAB driver is unloaded:  

```
sleep 3; /etc/methods/gabkext -status
```

gab: unloaded
- 18 Unconfigure LLT. On each node, type:  

```
/sbin/lltconfig -U
```
- 19 Halt and unload the LLT portable streams driver. On each node, perform the following steps:
  - Unload the LLT streams driver:  

```
/usr/sbin/strload -u -d /usr/lib/drivers/pse/llt
```
  - Confirm that the LLT streams driver is unloaded:  

```
/usr/sbin/strload -q -d /usr/lib/drivers/pse/llt
```

- 20 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](#).
- 21 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.
- 22 Reboot all nodes in the cluster.
- 23 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 safely ignore the warning about the failure to unfreeze the ClusterService group.
- 24 Bring the ClusterService service group online, if necessary. On any node, type:  

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

where `system` is the node name.

## Removing 5.0 MP3 RP2 on Storage Foundation or Storage Foundation Cluster File System

You can use the following procedure to uninstall 5.0 MP3 RP2 on Storage Foundation or Storage Foundation Cluster File System (SFCFS).

### To uninstall 5.0 MP3 RP2 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:
  - a Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
  - b Use the `vxrvg stop` command to stop each RVG individually:  

```
vxrvg -g diskgroup stop rvg_name
```
  - c On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:  

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

---

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

---

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

```
vxvol -g diskgroup stopall
```

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

```
vxprint -Aht -e v_open
```
- 8 Stop VCS along with all the resources. Then, stop the remaining resources manually:  

```
/etc/rc.d/rc2.d/vcs stop
```

9 Uninstall VCS:

```
cd /opt/VRTS/install
./uninstallvcs [-usersh]
```

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

```
rm /etc/vxfenmode
```

11 Unmount /dev/odm:

```
umount /dev/odm
```

12 Unload the ODM module:

```
modinfo | grep odm
modunload -i 154
```

13 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
```

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

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

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

---

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

---

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 RP2 on Storage Foundation for Oracle RAC

You can use the following procedure to uninstall the 5.0 MP3 RP2 on Storage Foundation for Oracle RAC systems.

### To uninstall the 5.0 MP3 RP2 on SF Oracle RAC

- 1 If CRS is not controlled by VCS, use the following command on each node to stop CRS:

```
/etc/init.crs stop
```

- 2 Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped. In the `gabconfig -a` command output, the VCS engine or high availability daemon (HAD) `port h` is not displayed. This indicates that VCS has been stopped.

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

Sample output:

```
GAB Port Memberships
=====
Port a gen 5c3d0b membership 01
Port b gen 5c3d10 membership 01
Port d gen 5c3d0c membership 01
Port o gen 5c3d0f membership 01
```

- 3 Uninstall Storage Foundation for Oracle RAC.

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

See the *Veritas Storage Foundation for Oracle RAC 5.0 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 for AIX*.

### MemCPUAllocator agent

Use the MemCPUAllocator agent to allocate CPU and memory to an IBM AIX dedicated partition. 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. Configure this resource as a leaf node in the service group dependency tree.

For prerequisites and other important information about this agent, refer to:

“[MemCPUAllocator agent notes](#)” on page 101

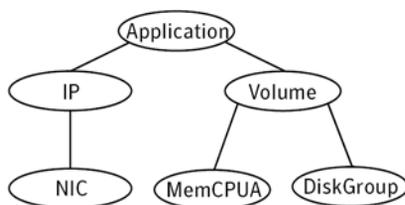
### Platform

AIX

### 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 | <p>Checks that the online agent function succeeded. If it succeeded, then the monitor agent function reports the resource state as <code>ONLINE</code>. If it did not succeed, then the monitor agent function reports the resource state as <code>OFFLINE</code>.</p> <p>If the agent is not able to allocate the required resources during the online agent function, the subsequent monitor reports <code>OFFLINE</code> and the resource faults. Because the resource is a leaf node, VCS engine stops bringing other resources online and marks the group as <code>FAULTED</code>. 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.</p> |

## Attributes

**Table 1-25** 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. The agent tries to connect to any HMC on this list in the order that they are specified.</p> <p>Type-dimension: string-vector</p> <p>Example: HMC = { myhmc1, myhmc2 }</p> |

**Table 1-26** Optional attributes

| <b>Optional attribute</b> | <b>Description</b>                                                                                                                                                                                                                                                        |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. 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. 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 101.
- See [“Dynamic resource allocation scenarios”](#) on page 102.
- See [“Configuring MemCPUAllocator”](#) on page 106.

## 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 1-27** 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.

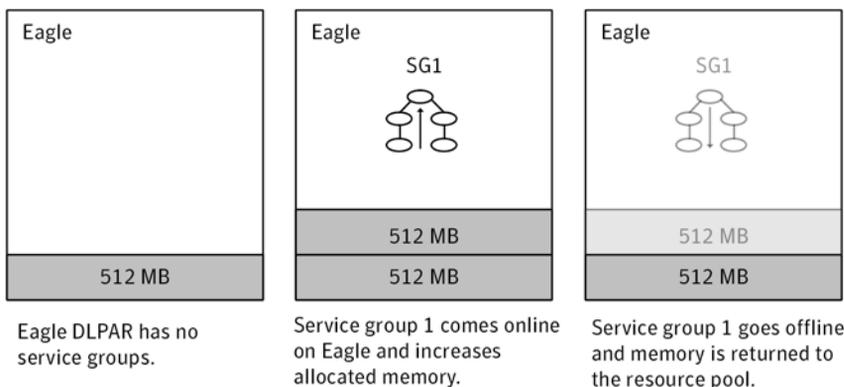
**Table 1-28** The memory that is required for service group SG1 and SG2

| 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.

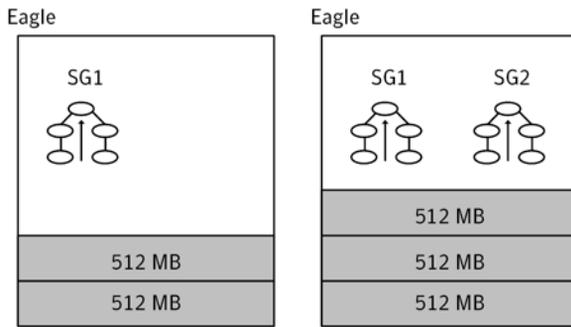
**Figure 1-2** Bringing a service group online and taking it offline on a DLPAR



### 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.

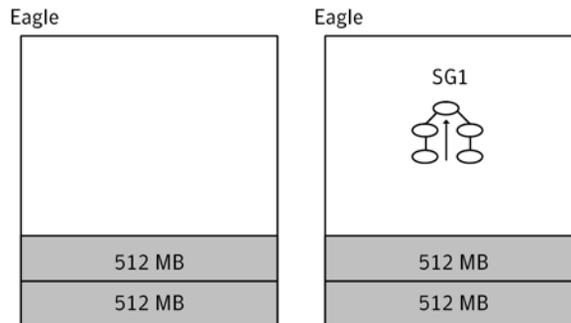
Figure 1-3 Bringing another service group online on a DLPAR

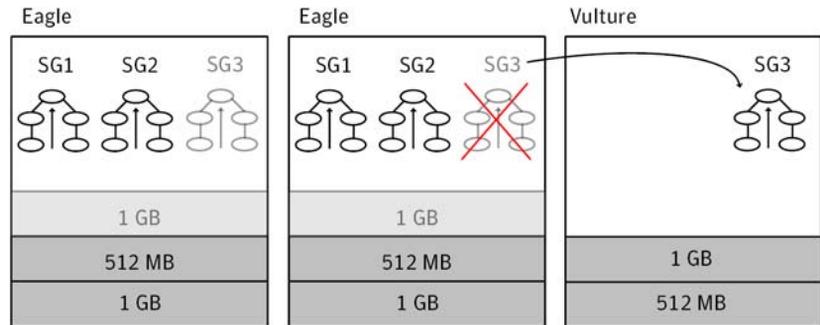


### Scenario 3: DLPAR has required resources

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 memory



**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, "Sample service group for a MemCPUAllocator resource, where the MemCPUA resource represents the MemCPUAllocator resource,"](#) on page 98.

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.

### 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.
 Use this keyword to remove a disk physically from the system, or
 to clean up a disk when you physically remove the disk from the
 system. See to the Veritas Volume Manager Administrator's Guide
 for more information.
```

This keyword does not exclude the disk from VxVM usage. To exclude the disk from VxVM usage, use the `vxddmpadm` command.

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.

The following options can be specified to restrict the ASL configuration to specific devices:

`scandisks [!]ctlr=controller_list`

Selects devices that are connected to the logical controllers specified as a comma-separated list. If you prepend a `!` to `ctlr`, all devices are selected except those that are connected to the specified controllers.

`scandisks [!]device=device_list`

Selects the devices that are specified as a comma-separated list. If you prepend a `!` to `device`, all devices except those listed are discovered.

`scandisks fabric`

Selects fabric devices only, such as devices that have the `DDI_NT_BLOCK_FABRIC` property set.

`scandisks new`

Selects new disks (that is, disks not known to VxVM).

`scandisks [!]pctlr=physical_controller_list`

Selects devices that are connected to the physical controllers specified as a list of items separated by `+` characters. If you prepend a `!` to `pctlr`, all devices are selected except those that are connected to the specified physical controllers.

## Veritas Cluster Server Installation Guide errata

You can find an updated version of this guide on the Symantec support website [http://www.symantec.com/enterprise/support/assistance\\_care.jsp](http://www.symantec.com/enterprise/support/assistance_care.jsp). 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

You can find an updated version of the following guides on the Symantec support website [http://www.symantec.com/enterprise/support/assistance\\_care.jsp](http://www.symantec.com/enterprise/support/assistance_care.jsp):

- 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