

Veritas Storage Foundation™ and High Availability Solutions Read This First

Solaris

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:

For Solaris SPARC,

<http://entsupport.symantec.com/docs/283886>

For Solaris x64,

<http://entsupport.symantec.com/docs/289317>

The *Veritas Cluster Server 5.0 Release Notes* can be viewed at the following URL:

For Solaris SPARC,

<http://entsupport.symantec.com/docs/283867>

For Solaris x64,

<http://entsupport.symantec.com/docs/289324>

Solaris is a trademark of Sun Microsystems, Inc.

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
	Oracle support	8
	Storage Foundation High Availability fixed issues	9
	Veritas Volume Manager fixed issues	9
	Veritas File System fixed issues	16
	Storage Foundation Cluster File System fixed issues	18
	Storage Foundation for Oracle fixed issues	19
	Storage Foundation for DB2 fixed issues	20
	Storage Foundation for Sybase fixed issues	22
	Veritas Cluster Server fixed issues	22
	Veritas Cluster Server agents for Veritas Volume Replicator fixed issues	31
	Storage Foundation and High Availability known issues	32
	Storage Foundation and High Availability known issues	32
	Veritas Volume Manager known issues	33
	Veritas File System known issues	34
	Storage Foundation Cluster File System known issues	34
	Storage Foundation for Oracle known issues	34
	Storage Foundation for DB2 known issues	35
	Storage Foundation for Oracle RAC known issues	36
	Veritas Cluster Server known issues	36
	Software limitations	37
	Storage Foundation for Oracle software limitations	37
	Storage Foundation for DB2 software limitations	37
	Veritas Cluster Server software limitations	38
	Changes in behavior for Storage Foundation High Availability	38
	Changes in Veritas Cluster Server behavior	38
	Downloading the rolling patch archive	41
	Patches included in this rolling patch	41
	Veritas Cluster Server patches	42
	Storage Foundation patches	43
	File System patches	44

Volume Manager and Volume Replicator patches	46
Storage Foundation Cluster File System patches	48
Storage Foundation for Oracle RAC patches	51
Storage Foundation for DB2 patches	54
Storage Foundation for Oracle patches	56
Storage Foundation for Sybase patches	59
Installing the Veritas software for the first time	62
Installing Storage Foundation or Storage Foundation Cluster File System and 5.0 MP3 RP2	62
Installing Storage Foundation for Oracle RAC and 5.0 MP3 RP2	65
Prerequisites for upgrading to 5.0 MP3 RP2	67
Upgrading 5.0 MP3 to 5.0 MP3 RP2	67
Performing a phased upgrade to 5.0 MP3 RP2 on a cluster	68
Performing a full upgrade to 5.0 MP3 RP2 on a cluster	82
Upgrading to 5.0 MP3 RP2 on a standalone system	98
Verifying software versions	102
Removing 5.0 MP3 RP2	103
Removing 5.0 MP3 RP2 from Veritas Cluster Server	103
Removing 5.0 MP3 RP2 on Storage Foundation or Storage Foundation Cluster File System	106
Removing 5.0 MP3 RP2 on Storage Foundation for Oracle RAC	109
Documentation errata	112
Manual pages errata	112
Veritas Cluster Server database installation and configuration guides errata	113

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) Solaris release.

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

For Solaris SPARC:

<http://entsupport.symantec.com/docs/281987>

For Solaris x64:

<http://entsupport.symantec.com/docs/286955>

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:

- Solaris 8 (SPARC Platform 32-bit and 64-bit)
- Solaris 9 (SPARC Platform 32-bit and 64-bit)
- Solaris 10 (SPARC and x64 Platform 64-bit)

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.

Oracle support

This release of Storage Foundation for Oracle offers support for Oracle 11.1.0.7 only for Solaris SPARC.

Storage Foundation High Availability fixed issues

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

- [Veritas Volume Manager fixed issues](#)
- [Veritas File System fixed issues](#)
- [Storage Foundation Cluster File System fixed issues](#)
- [Storage Foundation for Oracle fixed issues](#)
- [Storage Foundation for DB2 fixed issues](#)
- [Storage Foundation for Oracle fixed issues](#)
- [Storage Foundation for Sybase 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.
1189199	Fixed the cause of a system panic that occurred when you unloaded the <code>vxvmp</code> driver.
1195591	Fixed the cause of a panic when a cluster had an empty RVG.
1224659	Fixed an issue in which the <code>vxconfigbackup -p</code> script sometimes created a zero-length <code>.binconfig</code> file.
1259467	Fixed an issue in which the <code>accept()</code> call entered an infinite loop.
1286298	Fixed an issue in which proper locks were not taken in all necessary places while modifying <code>last_sent_seqno</code> .

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

Incident	Description
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.
1302064	Fixed an issue in which EFI disks could not be initialized or set up after formatting the disks.
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.
1402443	Fixed the cause of a system panic in the <code>kmsg_udp_payload()</code> call.
1408367	Fixed the cause of a system panic when <code>mutex_panic()</code> was called from <code>vol_rwsleep_wrlock()</code> .
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.
1418659	Fixed an issue in which a Jumpstart installation of the 4.1 MP2 and 4.1 MP2 RP3 patches created duplicate entries in the <code>/var/svc/profile/upgrade</code> file.
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 rlinks failed to be connected, followed by <code>vxconfigd</code> hanging on a secondary node.
1437281	Fixed the cause of an error with the <code>vxdmppadm -v getdmpnode enclosure=<name></code> command when a LUN was removed incorrectly.

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

Incident	Description
1446208	Changed message V-5-1-2140 from an error message to an informational message.
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.
1452957	Fixed a panic in the <code>bcopy()</code> call from <code>dmp_recv_scslpkt()</code> .
1457132	Fixed the cause of data corruption when running the <code>vxddmpadm disable path</code> and <code>vxddmpadm disable ctrlr</code> commands.
1457758	Fixed an issue in which the <code>vxdiskadm</code> command failed to replace a disk that was removed.
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> .
1461314	DMP no longer uses the SCSI bypass on single path disks for path-suppressing TPD.
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.
1471658	Fixed the cause of a <code>vxconfigd</code> daemon core dump that occurred in the <code>priv_get_all_udid_entry()</code> call.
1471763	Fixed the cause of the following error: build_devlink_list: readlink failed for /dev/vx/rdisk/ludg: Invalid argument
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> .

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

Incident	Description
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.
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 addmir</code> command, which was used to link a snapshot volume to the source volume.
1488084	Fixed an issue in which the <code>vxddmpadm 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 <code>ON</code> triggered unexpected offlining of paths on a DMX4 array.
1508462	Fixed the cause of a <code>vxconfigd</code> hang that occurred due to a split brain condition on a cluster.

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

Incident	Description
1512352	Fixed an issue in which the <code>vxconfigstore</code> command failed with the following error: VxVM vxconfigstore ERROR V-5-2-3706 Diskgroup configuration
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.
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>vx dg split</code> command failed with the following error: Internal configuration daemon error
1544051	Fixed an issue in which the incorrect bit was being checked for an EMC Symmetrix thin device.
1586879	Improved performance of the <code>vx disk 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 CLARIIION array, which led to an I/O hang.
1589172	Fixed an issue in which the <code>vx disksetup</code> and <code>vx diskunsetup</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>vx dmpadm getsubpaths dmpnodename</code> command now validates the <code>dmpnodename</code> value before getting the subpath information.
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.

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

Incident	Description
1598706	Fixed the cause of a system crash that occurred while mirroring the rootdisk.

[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
424397	Fixed an issue with VVR RU thread not starting nio after it is created from than waiting for all replicas to have NIO's created.
853207	Fixed an issue with 4.1 vxclust reconfig step 2 timed out on joining; node, reconfiguration looping.
1058665	Fixed the <code>vxdiskunsetup</code> command failing when disk access name does not match the physical path name.
1114699	Fixed the <code>vxtask</code> command to display the resync progress subtask for shared volumes with DRL
1135462	Fixed issue that was unable not to import disk group.
1192166	Fixed the <code>vxdbg -n [newdg] deport [origdg]</code> command causing a memory leak.
1224659	Fixed an issue with the <code>vxconfigbackupd</code> script leading to 0 byte <code>binconfig</code> file being created.
1230351	Fixed a system panic in <code>vol_klog_start()</code> due to accessing freed <code>mv read_sio</code> .
1269468	Fixed an issue with <code>vxconfigd</code> core dumps.
1281274	Fixed an issue with <code>vxplex</code> core dumps during <code>vxassist addlog</code> due to DRL log length being less than 33 blocks.
1288468	Fixed an issue with <code>vxconfigd</code> sleeping and no <code>vx</code> commands were responding.
1314301	Fixed an issue with <code>vxlustart</code> .
1368737	Fixed an issue when there are no mirrors to read, <code>VOL_READ_MIRRORS ioctl</code> returns -1 instead of 1.
1373432	Fixed a system panic in <code>bcopy()</code> due to null passed in from <code>volioctl_copyin()</code> .
1374927	Fixed an issue with <code>vxvm-startup2</code> does not set <code>VISSWAP</code> flag if swap device is encapsulated and mirrored.

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

Incident	Description
1375354	Fixed an issue with <code>vxcached</code> never deletes old snaps when cache hits HWM.
1382977	Fixed a system panic due to memory allocation.
1385126	Fixed an issue with VVR I/O hanging due to the wrong generation number assignment after recovery.
1385996	Fixed a rootdisk with B0 subdisk rendering unbootable after its removed and replaced with itself.
1386980	Fixed a system panic in <code>vol_putdisk()</code> code.
1389511	Fixed issue that was unable not to force import diskgroup version 80 in VxVM 5.0.
1393030	Fixed an issue with the <code>vxdiskunsetup</code> manual page failing when the <code>dmpnode</code> is not the primary path.
1393570	Fixed a FC-Switch port failure resulting in the loss one of four paths.
1397540	Fixed an issue with the <code>vxsnap restore</code> manual page is unable to properly freeze or thaw filesystems in a CVM environment.
1397877	Enhanced the <code>vxresize</code> manual page to run from non-CVM master.
1401188	Fixed a system panic after running the <code>vxctl enable</code> or <code>vxconfigd -k</code> commands.
1402144	Fixed a system panic due to invalid pointer being passed to <code>bcopy()</code> by <code>volkio_to_kio_copy</code> .
1409986	Fixed a segmentation fault on x64 system when running the <code>vxdmadm list dmpnode all</code> command.
1409991	Fixed an issue with <code>vxclust</code> configuration caused the cluster to panic.
1410216	Fixed a secondary log error causing <code>rlink</code> disconnect after IBC unfreeze.
1412784	Fixed an issue with the system hanging while creating volumes in the guest Ldom.
1413700	Fixed an issue with the wrong label on a device lead VxVM to calculate the wrong public region size.
1414451	The <code>vxsnap</code> manual page includes <code>mirror=enclosure</code> parameter to avoid being mirrored on the same enclosure.
1416930	Fixed an issue with the <code>vxvm</code> daemon that comes online when the system is rebooted.

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

Incident	Description
1421088	Fixed a secondary panic due to a corrupted <code>volsioq_start</code> .
1424479	Fixed an issue with <code>vxdmpadm</code> dumped core when executing <code>vxdmpadm list dmpnode</code> command.
1425338	Fixed an issue with CVR fails to connect rlinks followed by <code>vxconfigd</code> hangs on secondary.
1433120	Fixed an issue with after a reboot site read policy is not honored.
1435470	Fixed an issue with cluster nodes panicking after installing 5.0 MP3.
1435681	Fixed an issue with <code>vxesd</code> looping using 100% of one CPU.
1441003	Fixed a secondary panic due to double free of message with TCP protocol and 16 connection.
1443679	Fixed an issue in FMR3, I/Os initiating DCO updates for clearing DRL async clear region may not wait for its completion.
1443748	Fixed an issue in a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load.
1444425	The <code>vxsnap prepare</code> manual page includes support for the <code>mirror=</code> attribute.

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
1370823	Fixed an issue in which running a full <code>fsck</code> did not fix a file system.
1401516	Fixed the cause of a hang that occurred after locking a file system, disconnecting the storage cable, then using <code>fsadm</code> to unlock the file system.
1412465	Fixed an issue in which the <code>vxresize</code> command failed to resize the volume, even though the file system was successfully resized.
1426951	Fixed some badly formed <code>printf()</code> statements in <code>vxm_getinfo()</code> that caused a system panic.
1441487	Changed GMS to use the standard <code>gab_api_init()</code> call to avoid a possible GAB panic.
1445511	The <code>vx_cds_control()</code> call now releases active level 1 on an error path.

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

Incident	Description
1468377	You can now shrink a file system regardless of where the structural files reside on that file system.
1484888	Fixed an issue in which the cache hit percentage shown by <code>qiostat -l</code> command was inaccurate.
1517415	Fixed the cause of a core dump when running the <code>ncheck</code> command.
1526581	<code>vx_tflush_map()</code> no longer disables the file system if a map is marked as bad, but there is no I/O error.
1588199	Fixed an issue in which <code>dm_get_allocinfo()</code> failed with the EIO error for ext4 inodes with indirect pointers.
1601187	Reverted default <code>max_seqio_extent_size</code> to 2048, from 104857.
1634788	Fixed an issue in which the <code>fsadm</code> command dumped core intermittently when trying to defragment a file system.

[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
1413494	Fixed a failure of the <code>umount -f</code> command to unmount a VxFS file system.
1414175	Improved VxFS performance.
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 caused by running out of swap space, which resulted in a malloc failure.
1417973	Eliminated a benign error that occurred on globally- mounted VxFS file systems in a SunCluster environment when using the <code>scswitch</code> command or <code>mount</code> command.
1423867	Optimized <code>vx_convnode_data_files()</code> .
1428661	Improved the performance of <code>fsadm resize</code> on SFCFS.
1433066	Fixed a case of looping in <code>vx_do_putpage()</code> due to having a page beyond <code>i_wsize</code> .
1434438	Fixed a panic in <code>vx_unlockmap()</code> due to a null <code>ml_tranp</code> pointer.

Table 1-4 Veritas File System 5.0 MP3 RP1 fixed issues

Incident	Description
1437490	The <code>fsclustadm's lltdb.c</code> is now mult-threaded safe for CFSMountAgent.

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
1518713	The <code>vxfsckd -n</code> command now initializes the <code>nthrs</code> variable.
1531031	Fixed an issue in which quota hard limits could be exceeded on a clustered file system.
1539892	Fixed an issue in which a clustered file system that was mounted on one node required <code>fsck</code> to be run.
1556159	Fixed an issue in which adding a file system to a diskgroup caused the monitor to label the <code>cvmvoldg</code> resource as offline, which in turn caused other CFS file systems to become offline.
1591783	Optimized <code>getattr()</code> to operate faster when binaries are mmaped from many nodes.
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
1447197	Fixed an issue after a 5.0 MP3 upgrade, CFSMountAgent restarts and is not sending alive messages.

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> .
1530125	Fixed an issue in which the owner of the following directories was changed when installing <code>VRTSdbms</code> 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 <code>vxdbms</code> server, in which an attacker could see the name and port of the server.

[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.
1435906	Fixed JumpStart problem of VxDBMS package perl scripts are not executable.

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
1481426	Fixed an issue in which the owner of the following directories was changed when installing patches or packages for the Storage Foundation for DB2 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 vxsnapadm_50.log file, which is used for trace vxsnapadm issues.

Table 1-9 Storage Foundation for DB2 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 DB2 5.0 or 5.0 MP3 releases: <ul style="list-style-type: none"> ■ /etc ■ /etc/default ■ /etc/init.d ■ /etc/rc2.d ■ /opt
1651363	Fixed a security issue with the vxdbms 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 /tmp.
1434688	Storage Foundation for DB2 is no longer creating world writable files under /tmp.
1435527	Improved boot time for DBEDAgent startup script.
1435906	Fixed JumpStart problem with VxDBMS package perl scripts are not executable.

Storage Foundation for Sybase fixed issues

[Table 1-11](#) describes fixed issues in the Storage Foundation for Sybase 5.0 MP3 RP2 release.

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

Incidents	Description
1481426	<p>Fixed an issue in which the owner of the following directories was changed when installing patches or packages for the Storage Foundation for Sybase 5.0 or 5.0 MP3 releases:</p> <ul style="list-style-type: none"> ■ /etc ■ /etc/default ■ /etc/init.d ■ /etc/rc2.d ■ /opt

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	<p>[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.</p>
1362407	<p>[LLT] Fixed an issue in which the <code>lltdump</code> command failed to display all the LLT packets and produces the following error:</p> <pre>bash-3.00# /opt/VRTSllt/lltdump -f /dev/bge2 CR C 60425 S 2559 D 00 P 000 rdy 0000 seq 000001dc len 0000 lltdump: cannot read messages on /dev/bge2: Error 0</pre> <p>The <code>lltdump</code> command gets control and data information from <code>dlpi</code> streams read head queue. The initial buffer size passed to get control information was 36. The latest <code>dlpi</code> drivers like <code>bge</code> and <code>nge</code> have control information that is larger than 36. Insufficient buffer size for control information produces the error message "Cannot read messages ". The buffer size was increased from 36 to 64.</p>

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

Incidents	Description
1368385	[Agents] Fixed an issue in which DiskGroupSnap does not work if layered volumes are used. VxVM creates layered volumes by default, in case of larger volumes spanning multiple disks. The agent expects each volume to have a plex at each site but VxVM does not assign a site tag to plex and there is only one top level plex. Thus, the agent reports that the configuration is invalid. This was a limitation in the original agent when no layered volumes were supported.
1377324	[Agents] Fixed a parsing error which caused an error message to appear in the <code>/var/VRTSvcs/log/tmp/Oracle-0</code> file.
1451717	[VCS] Fixed an issue in which the correct error message was not displayed if the value of non-existing attribute was queried for a node from the remote cluster. The command <code>hasys -value <sys_from_remote_cluster> JunkAttribute</code> produces a core dump.
1465956	[VCS] Fixed an issue in which you cannot delete a system even if it has no service group configured on it. Whenever a system is added, it is added to the SystemList of the VCShmg group (if HostMonitorLogLvl is enabled). While deleting the system from the cluster, VCS should silently delete this from the SystemList of VCShmg. However, it produces an error. VCS now lets you delete the system without displaying any error.
1482806	[GAB] Fixed an issue in which uninstalling GAB produced the following error "Error in removing the gab entry in the <code>/etc/devlinks.tab</code> " when the GAB module was not loaded in the kernel.
1487725	[Agents] Fixed an issue in which the zone agent monitor script failed with an unexpected error. In the month of December, the Zone agent monitor would fail with the message: "Month '12' out of range 0..11 at <code>/opt/VRTSvcs/bin/Zone/monitor</code> line 164". The Zone agent monitor code was not setting the <code>timelocal()</code> function properly. Correct monitor code. Note that the issue is related only to a specific month of the year.
1469788/ 1469787	[LLT] Fixed an issue in which LLT cannot be unloaded and returns the error message "Module LLT is in use" even when the system was shutdown.
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	[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.
1522568	[Agents] Fixed an issue in which the agent framework crashed while setting the resource name for the dependent attribute.
1528584	[Agents] Fixed an issue where the system performance dropped when a large number of application resources are configured and the Application agent searches the process table continuously.
1537111	[VCS] VCS issues warning messages with <code>ha</code> commands on a ZFS root file system due to the <code>prcntl()</code> function being called with a NULL <code>sched_class</code> .
1537141	[Agents] Fixed an issue in which the Mount agent leaks memory despite the installation of the 5.0MP3HF1 patch.
1538208	[VCS] Fixed an issue in which the value of attribute <code>HostUtilization</code> is not 0 even after <code>HostMonitor</code> resource is deleted.
1539087	[Agents] Fixed an issue in which the agent framework seems to be leaking memory during message logging.
1540807	[GAB] Fixed an issue in which the error number returned by the <code>gab_receive()</code> function in the GAB library is wrong. The <code>gab_receive()</code> function returns -1, but the error number was set to 0.
1542326	[Agents] Fixed an issue in which the IPMultiNICB agent crashes and produces core dump when monitoring an IP address that is brought up outside of VCS control. An IP address brought up outside of VCS control, e.g., as a part of a non-global zone configuration, can be monitored by an IPMultiNICB resource. Such a configuration exercises a code path in the agent which causes a core dump. Source code agent to fix the problem.
1542382	[Agents] Fixed an issue in which starting the Mount agent created a defunct process.

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

Incidents	Description
1542391	<p>[Agents] Fixed an issue in which VCS indicated that the zone was online when it was not active by modifying the zone agent for better monitoring. The Zone agent uses the RUNNING state to determine if a non-global zone resource is online. A non-global zone can go into the running state even before all the services inside the non-global zone are started. Added the BootState attribute to determine at what level the non-global zone is considered to be online: single-user, multi-user, or multi-user-server.</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.</p> <p>See “New attribute EntryPointTimeout” on page 40.</p>
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 39.</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 VCShmg group with a HostMonitor type resource VCShm. 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>
1600452	<p>[Fencing] Fixed an issue in which the script to shutdown fencing (vxfen) produces an unexpected error message.</p>

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

Incidents	Description
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>
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>vxfsenwap</code> command. If the keys on coordinator disks are accidentally cleared, they can be refreshed using the <code>vxfsenwap</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>vxfsenwap</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>vxfsenwap</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 38.</p>
1632806/ 1677496	<p>[GAB] Fixed an issue in which panic results when clients access the <code>gab_api</code> pointer through <code>GAB_API_INIT</code>.</p>

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

Incidents	Description
1633781	[VCS] Fixed an issue in which the NFS resource goes to faulted state even after it is restarted if rpcbind/portmap daemon is restarted. During the online monitoring of the NFS resource, if the rpcbind/portmap daemon is restarted, the NFS resource monitor entry point detects the resource as offline unexpectedly. This triggers the clean entry point for the resource. The clean entry point gets executed successfully and thereafter, the NFS resource tries to restart itself. The monitor entry point after the restart again detects the NFS resource as offline and the resource goes to FAULTED state. The clean entry point is used to check whether the server daemons are running or not. If the server daemons are running, it does nothing and exits successfully. However, the running daemons do not indicate that they are registered with rpcbind/portmap. The rpcbind/portmap restart terminates the registrations of all RPC daemons. So the RPC service daemons must be restarted whenever the rpcbind/portmap restarts itself. Thus, the monitor was returning offline even when the daemons were running. The clean entry point now always restarts the server daemons. If the server daemons are running, it kills the running daemons.
1633973	[VCS] Fixed an issue in which the node does not test the Authority attribute before bringing the faulted service group online, leading to concurrency violations and the service group being taken offline on the disaster recovery site.
1634924	[VCS] Fixed an issue in which the engine logs indicated CPU usage even after the HostMonitor resource is deleted.
1635792	[VCS] Fixed an issue in which the Zpool monitor returned unknown when ZFS filesystem snapshot was created. The Zpool agent monitor checks if all the ZFS file systems are mounted. If the Zpool agent monitor does not find a file system mounted, it sets the UNKNOWN state flag. Thus, ZFS snapshots are not mounted and this results in the UNKNOWN flag being set for the ZPool resource. If the ZFS file system is a snapshot, the check for mounted status is not done and hence, the UNKNOWN state flag is not set.
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.

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

Incidents	Description
1638725	<p>[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.</p>
1668609	<p>[Agents] Fixed an issue in which the Proxy agent is updated to allow the target resource to be probed before scheduling the first probe of the Proxy resource.</p>
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 commad: <code>hagrp -switch Hard_ParentSG -to sysB</code> switches only the hard parent group along with the child group. When the <code>hargp -switch</code> command is executed with any of the following options:</p> <ul style="list-style-type: none"> i) <code>hagrp -switch SG_parent -any</code> ii) <code>hagrp -switch SG_parent -to <sys></code> <p>The parent group switches (while the child group is online) only in the case of a hard dependency. The switch does not happen in the case of soft or firm dependency. The switch operation succeeds for an OLH parent, if only the parent group is online. The child group has no other parents online. The OLH parent and child group can have other parents. However, the OLH child group is always a leaf node.</p>
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>

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

Incidents	Description
1677412	[Agents] Fixed an issue so that when the SystemList of the service group is modified, you do not start all agents but only the required agents. The agent that was stopped by a user on a system gets restarted even if the group has no resource of that agent type, when the SystemList is modified to add that system. On SystemList modification to add new systems in SystemList, the engine starts all the agents without ensuring if the group has a resource of that type. Code changes so that only agents for which the group has resources are started whenever the SystemList is modified to add a new system.
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.
1713201	[Agents] Fixed an issue in which the Oracle agent starts Oracle with a non-default Oracle userid but the monitor function does not detect it as online. When you have a dummy user that belongs to the same group as the Oracle binaries and is a part of the Owner attribute, the Oracle agent starts Oracle but the monitor function does not detect it as online. This happens because the IDof the Owner attribute and the id of the /proc/PID/object/a.out file are checked. The a.out file is the same as the \$ORACLE_HOME/bin/oracle binary. Since these two do not match, the agent detects it as online. The user ID of \$ORACLE_HOME/bin/oracle binary was matched to that of the /proc/PID/object/a.out file. If these two user ids matched, you cache the cookie and proceed with the next process.

[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.

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

Incidents	Description
1392826	<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># haconf -makerw # hagrps -modify <i>servicegroup_name</i> PreOnline 1 # haconf -dump -makero</pre>
1394624	LLT: fixed an issue where the lldlv thread spun indefinitely.
1395905	Changes implemented to close device file for device vxddmpconfig.
1397692	Removed a condition where VCS engine clients hung in connect when the target system was down.
1397738	Support provided for Solaris 8 and Solaris 9 branded zones.
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>hagrps -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.
1427100	Fixed an issue where LDom CfgFile did not work with LDom 1.0.3.
1457429	Removed the VCS NOTICE V-16-1-53021 message after the <code>hastart</code> command is run.

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
1255362	The RVG Snapshot agent now picks up volumes that are not in the RVG.
1295115	Enabled the fdsetup wizard to set up a firedrill SG in a secured VVR-GCO environment.
1433149	Fixed issues related to the OnlineTimeout attribute with RVGPrimary and RVGSharedPri agents.
1671357	Enabled the RVGPrimary agent to migrate a VVR primary to secondary in the case of having multiple secondaries.

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.

Veritas Volume Manager does not update boot-archive on SPARC Solaris 10 Update 6 and up (1471606)

The SPARC version of Solaris 10 Update 6 adds use of boot-archive. VxVM does not update the boot-archive after installation. Usually this limitation causes no issues; however, there could be an issue if an unclean shutdown occurs prior to the next clean shutdown.

Workaround

After installation has completed, run the following command:

```
/sbin/bootadm update-archive
```

A reboot is required in some situations if the root disk is encapsulated on Solaris 10 (1190522)

Veritas Volume Manager is not registering and unregistering SMF scripts correctly. This issue may occur if the root disk is encapsulated on Solaris 10. To achieve proper behavior, a reboot is required.

Workaround

A reboot is required in the following situations:

- after removing the VxVM 5.0 MP3 RP2 using the `patchrm` command
- after installing 5.0 MP3 and before installing the VxVM 5.0 MP3 RP2

Autotagging can cause the reattach of a site to fail in a Campus Cluster (1470548)

Using the autotagging feature for a Campus Cluster can cause the site attach to fail. VxVM displays an error message such as the following:

```
VxVM vxdg ERROR V-5-1-10128 tagid already assigned to disk
```

Workaround

Tag disks manually; do not use autotagging.

STK6x50 array in A/PF mode can get spurious path failures (1471740)

For an STK6x50 array in A/PF mode, if any open is done on any paths during failover, the open may fail. The paths may be marked as failed. The DMP node may then go into failed state, potentially causing any plexes associated with the node to become detached.

Workaround

If a plex becomes detached, manually clear the FAILING flag on the disk and reattach the plex.

DS4x00 array A/PF failover on one node can cause I/O failure on other nodes in cluster (1471794)

SCSI-3 PR keys are not getting registered correctly on secondary paths of the cluster. Any I/O on the device will fail and can cause the plex to be detached.

Workaround

The DS4x00 array is only supported in A/PC mode.

Veritas File System known issues

There are no known issues for Veritas File System in this release.

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.

patchrm fails when removing VRTSdbcom, VRTSdbed, and VRTSdb2ed patches (1726470)

The `patchrm` command fails because of missing backout data when removing the 5.0 MP3 RP2 versions of the `VRTSdbcom`, `VRTSdbed`, and `VRTSdb2ed` patches. See Sun CR 6581364 for more information.

Workaround

To uninstall the patches:

- 1 Uninstall the `VRTSdbcom`, `VRTSdbed`, or `VRTSdb2ed` package, as appropriate to your setup.
- 2 Install 5.0 GA version of the `VRTSdbcom`, `VRTSdbed`, or `VRTSdb2ed` package.
- 3 Run the `dbed_patch_50ga` script:

```
# cd /net/release/re/release_train/sol/5.0MP3/SxRT-5.0MP3-GA/\
dvd1-sol_sparc/storage_foundation_for_oracle/scripts
# dbed_patch_50ga
```
- 4 Install the 5.0 MP3 patch for the `VRTSdbcom`, `VRTSdbed`, or `VRTSdb2ed` package.

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.

Joining a new node to the cluster may fail (1390591)

If you have a RAC cluster that has fencing enabled and a Sun StorageTek 2540 machine configured in A/PF mode, joining a new node to the cluster may fail if the cluster has a failover in progress.

Workaround

There is no known workaround.

Veritas Cluster Server known issues

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

The redundant nfs_restart trigger executes unnecessary hares -list commands (1542334)

The `hatrigger` script calls the `nfs_restart` trigger after being brought online. This issues unnecessary `hares -list` commands. The `nfs_restart` trigger has been replaced by the `NFSRestart` agent and is redundant.

Workaround:

Remove the obsolete `nfs_restart` trigger from the `hatrigger` script by moving it from the `/opt/VRTSvcs/bin/triggers` directory to some other directory. VCS then does not execute any of the `nfs_*` triggers and instead issues a log message "Failed to send trigger for triggername; script does not exist".

Mount agent reports unexpected offline for the Mount resource and dumps core. (1537141)

The Mount agent reports unexpected offline for Mount resource of FSType=nfs. It also dumps core. The getmntany() system call used by the Mount agent leaks memory. The memory leak increases with local zones and SecondLevelMonitor=1. This is due to the existence of a documented Sun bug 6813502 - "mntfs will leak mappings when called from a forking MT program"

Workaround:

Contact Sun Microsystems for further guidance.

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 19 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 20 for incident 1425261.

Veritas Cluster Server software limitations

Following is a known limitations in the 5.0 MP3 RP2 release of Veritas Cluster Server:

When using Live Upgrade to perform a product upgrade VRTSvcs is not removed successfully (1485016)

If you attempt to use Live Upgrade for a product upgrade of VCS, the VRTSvcs package does not get removed successfully from the Alternate Boot Environment.

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.
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 example, on Solaris, if the EPClass attribute is TS*, and the value of the EPPriority attribute is set to 0, then the base priority for entry points is set to 59 by the operating system. Similarly on Solaris, if scheduling class is RT*, then base priority is 100.</p> <p>If the value of this attribute varies from -60 to 60 (except 0 and -1), it increases or decreases the base priority by the configured value. For example, on Solaris, if EPClass is set to TS* and EPPriority is set to -20, then the scheduling priority of the entry point would be 39 (59 base value and - 20 configured value).</p> <p>*TS (for Solaris) = TimeShare scheduling class *RT (for Solaris) = RealTime scheduling class</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.

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:

For Solaris SPARC,

<http://entsupport.symantec.com/docs/281987>

For Solaris x64,

<http://entsupport.symantec.com/docs/286955>

Patches included in this rolling patch

This section describes the Solaris SPARC and x64 patches included in this rolling patch.

- [Veritas Cluster Server patches](#)
- [Storage Foundation patches](#)
- [File System patches](#)
- [Volume Manager and Volume Replicator patches](#)
- [Storage Foundation Cluster File System patches](#)
- [Storage Foundation for Oracle RAC patches](#)
- [Storage Foundation for DB2 patches](#)
- [Storage Foundation for Oracle patches](#)
- [Storage Foundation for Sybase patches](#)

Veritas Cluster Server patches

This sections describes the VCS Solaris SPARC and x64 patches.

Solaris SPARC

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

Table 1-16 VCS 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
123722-02	VRTSat 5.0MP3: Maintenance Patch for Authentication Server	39000	X	X	X
139356-02	Contains fixes for: VRTSllt, VRTSgab, VRTSvxfen, VRTSvcs, and VRTSvcsag	97000	X		
139357-02	Contains fixes for: VRTSllt, VRTSgab, VRTSvxfen, VRTSvcs, and VRTSvcsag	97000		X	
139358-02	Contains fixes for: VRTSvcs and VRTSvcsag	88000			X
139359-02	Contains fixes for: VRTSllt, VRTSgab, and VRTSvxfen	8000			X

Solaris x64

[Table 1-17](#) describes the Solaris x64 VCS patches that are included in this rolling patch:

Table 1-17 VCS 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
139360-02	Contains fixes for: VRTSllt, VRTSgab, and VRTSvxfen	7000	X
139361-02	Contains fixes for: VRTSvcs and VRTSvcsag	86000	X

Storage Foundation patches

This sections describes the Storage Foundation Solaris SPARC and x64 patches.

Solaris SPARC

[Table 1-18](#) describes the Solaris SPARC Storage Foundation patches that are included in this rolling patch:

Table 1-18 SF 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
140657-01	VRTSdsa 5.0MP3RP2: Maintenance Patch for VRTSdsa 5.0	20	X	X	X
140661-01	VRTSobgui 5.0MP3RP2: Maintenance Patch for VEA GUI	64000	X	X	X
141279-01	VRTSmapro 5.0MP3RP2: Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	26	X	X	X

See [Table 1-20](#) for VxFS patches, [Table 1-22](#) for VxVM patches, and [Table 1-24](#) for SFCFS patches that are also included in the Storage Foundation package.

Solaris x64

[Table 1-19](#) describes the Solaris x64 Storage Foundation patches that are included in this rolling patch:

Table 1-19 SF 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
127362-03	VRTSddlpr 5.0MP3RP2_x86: Rolling Patch 02 for VRTSddlpr 5.0 MP3	150000	X
141280-01	VRTSmapro 5.0MP3RP2_x86: Rolling Patch for 5.0MP3 for Solaris 10	13	X

See [Table 1-21](#) for VxFS patches, [Table 1-23](#) for VxVM patches, and [Table 1-25](#) for SFCFS patches that are also included in the Storage Foundation package.

File System patches

This sections describes the File System Solaris SPARC and x64 patches.

Solaris SPARC

[Table 1-20](#) describes the Solaris SPARC File System patches that are included in this rolling patch:

Table 1-20 VxFS 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
123200-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.8	13516	X		
123201-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.9	13516		X	
123202-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.10	15938			X
123722-02	VRTSat 5.0MP3: Maintenance Patch for Authentication Server	39000	X	X	X
139737-01	VRTSdcli 5.0MP3RP1: Rolling Patch 01 for VRTSdcli 5.0MP3	26300	X	X	X
139741-02	VRTSob 5.0MP3RP2: Maintenance Patch for VEA Server	6000	X	X	X
139742-02	VRTSobc33 5.0MP3RP2: Maintenance Patch for VEA Server	23000	X	X	X
139743-01	VRTSaa 5.0MP3RP1: Maintenance Patch for VRTSaa	1000	X	X	X
139744-01	VRTSccg 5.0MP3RP1: Maintenance Patch for VRTSccg	200	X	X	X

Solaris x64

Table 1-21 describes the Solaris x64 File System patches that are included in this rolling patch:

Table 1-21 VxFS 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
127337-03	VRTSvxfs 5.0MP3RP_x86: Rolling Patch for File System 5.0MP3	13516	X
139738-01	VRTSdcli 5.0MP3RP2_x86: Rolling Patch 01 for VRTSdcli 5.0MP3	31000	X
139745-02	VRTSob_x86 5.0MP3RP2: Maintenance Patch for VEA Server	14000	X
139746-02	VRTSobc33_x86 5.0MP3RP2: Maintenance Patch for VEA Server	69000	X
139747-01	VRTSaa_x86 5.0MP3RP1: Maintenance Patch for VRTSaa	2000	X
139748-01	VRTSccg_x86 5.0MP3RP2: Maintenance Patch for VRTSccg	340	X
140658-01	VRTSdsa 5.0MP3RP2: Maintenance Patch for VRTSdsa 5.0	25	
140662-01	VRTSobgui 5.0MP3RP2_x86: Maintenance Patch for VEA GUI	114000	

Volume Manager and Volume Replicator patches

This sections describes the Volume Manager and Volume Replicator Solaris SPARC and x64 patches.

Solaris SPARC

[Table 1-22](#) describes the Solaris SPARC Volume Manager and Volume Replicator patches that are included in this rolling patch:

Table 1-22 VxVM and VVR 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
122058-12	VRTSvxvm 5.0MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3	184000	X	X	X
123722-02	VRTSat 5.0MP3: Maintenance Patch for Authentication Server	39000	X	X	X
123823-05	VRTSddlpr 5.0MP3RP2: Rolling patch 02 for VRTSddlpr 5.0 MP3	13954	X	X	X
139354-01	VRTSvmmman 5.0MP3RP1: Rolling Patch 01 for Volume Manager 5.0MP3	530	X	X	X
139737-01	VRTSdcli 5.0MP3RP1: Rolling Patch 01 for VRTSdcli 5.0MP3	26300	X	X	X
139739-01	VRTSvmpro 5.0MP3RP1: Rolling Patch 01 for VRTSvmpro 5.0	19000	X	X	X
139741-02	VRTSob 5.0MP3RP2: Maintenance Patch for VEA Server	6000	X	X	X
139742-02	VRTSobc33 5.0MP3RP2: Maintenance Patch for VEA Server	23000	X	X	X
139743-01	VRTSaa 5.0MP3RP2: Maintenance Patch for VRTSaa	1000	X	X	X
139744-01	VRTSccg 5.0MP3RP2: Maintenance Patch for VRTSccg	200	X	X	X
140657-01	VRTSdsa 5.0MP3RP2: Maintenance Patch for VRTSdsa 5.0	20	X	X	X
140661-01	VRTSobgui 5.0MP3RP2: Maintenance Patch for VEA GUI	64000	X	X	X

Table 1-22 VxVM and VVR 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
141745-01	VCS Agents for VVR 5.0: MP3RP2 for VVR 5.0	110	X	X	X

Solaris x64

[Table 1-23](#) describes the Solaris x64 Volume Manager and Volume Replicator patches that are included in this rolling patch:

Table 1-23 VxVM and VVR 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
127336-03	VM 5.0_x64_MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3_x86	110000	X
127362-03	VRTSddlpr 5.0MP3RP2_x86: Rolling Patch 02 for VRTSddlpr 5.0 MP3	150000	X
128091-02	VRTSvcsvr 5.0 MP3 RP2: Rolling patch 02 for VCS agents for Veritas Volume Replicator 5.0	200	X
139355-01	VRTSvmman 5.0MP3RP1: Rolling Patch 01 for Volume Manager 5.0MP3_x86	920	X
139738-01	VRTSdcli 5.0MP3RP2_x86: Rolling Patch 01 for VRTSdcli 5.0MP3	31000	X
139740-01	VRTSvmpro 5.0MP3RP2_x86: Rolling Patch 01 for VRTSvmpro 5.0	29000	X
139745-02	VRTSob_x86 5.0MP3RP2: Maintenance Patch for VEA Server	14000	X
139746-02	VRTSobc33_x86 5.0MP3RP2: Maintenance Patch for VEA Server	69000	X
139747-01	VRTSaa_x86 5.0MP3RP2: Maintenance Patch for VRTSaa	2000	X
139748-01	VRTSccg_x86 5.0MP3RP2: Maintenance Patch for VRTSccg	340	X

Storage Foundation Cluster File System patches

This sections describes the Storage Foundation Cluster File System Solaris SPARC and x64 patches.

Solaris SPARC

[Table 1-24](#) describes the Solaris SPARC Storage Foundation Cluster File System patches that are included in this rolling patch:

Table 1-24 SFCFS 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
122058-12	VRTSvxvm 5.0MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3	184000	X	X	X
123088-03	VRTSgms 5.0MP3RP2: Rolling Patch for GMS 5.0 MP3 - Sun 5.8	97	X		
123089-03	VRTSgms 5.0MP3RP2: Rolling Patch for GMS 5.0 MP3 - Sun 5.9	97		X	
123090-03	VRTSgms 5.0MP3RP2: Rolling Patch for GMS 5.0 MP3 - Sun 5.10	97			X
123200-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.8	13516	X		
123201-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.9	13516		X	
123202-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.10	15938			X
123722-02	VRTSat 5.0MP3: Maintenance Patch for Authentication Server	39000	X	X	X
123823-05	VRTSddlpr 5.0MP3RP2: Rolling patch 02 for VRTSddlpr 5.0 MP3	13954	X	X	X
139354-01	VRTSvmmman 5.0MP3RP2: Rolling Patch 01 for Volume Manager 5.0MP3	530	X	X	X
139356-02	SunOS 5.8: fixes for gab, llt, vxfen, vcs, vcsag	97000	X		
139357-02	SunOS 5.9: fixes for gab, llt, vxfen, vcs, vcsag	97000		X	

Table 1-24 SFCFS 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
139358-02	SunOS 5.10: fixes for vcs, vcsag	88000			X
139359-02	SunOS 5.10: fixes for gab, llt, vxfen	8000			X
139737-01	VRTSdcli 5.0MP3RP2: Rolling Patch 01 for VRTSdcli 5.0MP3	26300	X	X	X
139739-01	VRTSvmpro 5.0MP3RP2: Rolling Patch 01 for VRTSvmpro 5.0	19000	X	X	X
139741-02	VRTSob 5.0MP3RP2: Maintenance Patch for VEA Server	6000	X	X	X
139742-02	VRTSobc33 5.0MP3RP2: Maintenance Patch for VEA Server	23000	X	X	X
139743-01	VRTSaa 5.0MP3RP2: Maintenance Patch for VRTSaa	1000	X	X	X
139744-01	VRTSccg 5.0MP3RP2: Maintenance Patch for VRTSccg	200	X	X	X
139753-02	VRTScavf 5.0MP3RP2: Maintenance Patch for Cluster Server Agents 5.0	405	X		
139754-02	VRTScavf 5.0MP3RP2: Maintenance Patch for Cluster Server Agents 5.0	405		X	
139755-02	VRTScavf 5.0MP3RP2: Maintenance Patch for Cluster Server Agents 5.0	405			X
141745-01	VCS Agents for VVR 5.0: MP3RP2 for VVR 5.0	110	X	X	X

Solaris x64

[Table 1-25](#) describes the Solaris x64 Storage Foundation Cluster File System patches that are included in this rolling patch:

Table 1-25 SFCFS 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
127336-03	VM 5.0_x64_MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3_x86	110000	X

Table 1-25 SFCFS 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
127337-03	VRTSvxfs 5.0MP3RP_x86: Rolling Patch for File System 5.0MP3	13516	X
127341-03	VRTSgms 5.0MP3RP2_x86: Rolling Patch for GMS 5.0 MP3	90	X
127362-03	VRTSddlpr 5.0MP3RP2_x86: Rolling Patch 02 for VRTSddlpr 5.0 MP3	150000	X
128091-02	VRTSvcsvr 5.0 MP3 RP2: Rolling patch 02 for VCS agents for Veritas Volume Replicator 5.0	200	X
139355-01	VRTSvmman 5.0MP3RP1: Rolling Patch 01 for Volume Manager 5.0MP3_x86	920	X
139360-02	SunOS 5.10: fixes for gab, llt, vxfen	7000	X
139361-02	SunOS 5.10: fixes for vcs, vcsag	86000	X
139738-01	VRTSdcli 5.0MP3RP1_x86: Rolling Patch 01 for VRTSdcli 5.0MP3	31000	X
139740-01	VRTSvmpro 5.0MP3RP1_x86: Rolling Patch 01 for VRTSvmpro 5.0	29000	X
139745-02	VRTSob_x86 5.0MP3RP2: Maintenance Patch for VEA Server	14000	X
139746-02	VRTSobc33_x86 5.0MP3RP2: Maintenance Patch for VEA Server	69000	X
139747-01	VRTSaa_x86 5.0MP3RP2: Maintenance Patch for VRTSaa	2000	X
139748-01	VRTSccg_x86 5.0MP3RP2: Maintenance Patch for VRTSccg	340	X
139756-02	VRTScavf 5.0MP3RP2_x86: Maintenance Patch for Cluster Server Agents 5.0	9451	X
140658-01	VRTSdsa 5.0MP3RP2: Maintenance Patch for VRTSdsa 5.0	25	X
140662-01	VRTSobgui 5.0MP3RP2_x86: Maintenance Patch for VEA GUI	114000	X

Storage Foundation for Oracle RAC patches

This sections describes the Storage Foundation for Oracle RAC Solaris SPARC and x64 patches.

Solaris SPARC

[Table 1-26](#) describes the Solaris SPARC Storage Foundation for Oracle RAC patches that are included in this rolling patch:

Table 1-26 SF for Oracle RAC 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
122058-12	VRTSvxvm 5.0MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3	184000	X	X	X
123200-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.8	13516	X		
123201-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.9	13516		X	
123202-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.10	15938			X
123722-02	VRTSat 5.0MP3: Maintenance Patch for Authentication Server	39000	X	X	X
139354-01	VRTSvmman 5.0MP3RP1: Rolling Patch 01 for Volume Manager 5.0MP3	530	X	X	X
139356-02	SunOS 5.8: fixes for gab, llt, vxfen, vcs, vcsag	97000	X		
139357-02	SunOS 5.9: fixes for gab, llt, vxfen, vcs, vcsag	97000		X	
139358-02	SunOS 5.10: fixes for vcs, vcsag	88000			X
139359-02	SunOS 5.10: fixes for gab, llt, vxfen	8000			X
139362-01	VRTSdbms3 5.0MP3RP2: Rolling Patch for Solaris 8, 9 and 10	64	X	X	X
139366-03	VRTSdbcom 5.0MP3RP2 Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	2971	X	X	X

Table 1-26 SF for Oracle RAC 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
139367-01	VRTSdbed 5.0MP3RP2 Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	3150	X	X	X
139737-01	VRTSdcli 5.0MP3RP1: Rolling Patch 01 for VRTSdcli 5.0MP3	26300	X	X	X
139739-01	VRTSvmpro 5.0MP3RP1: Rolling Patch 01 for VRTSvmpro 5.0	19000	X	X	X
139741-02	VRTSob 5.0MP3RP2: Maintenance Patch for VEA Server	6000	X	X	X
139742-02	VRTSobc33 5.0MP3RP2: Maintenance Patch for VEA Server	23000	X	X	X
139743-01	VRTSaa 5.0MP3RP1: Maintenance Patch for VRTSaa	1000	X	X	X
139744-01	VRTSccg 5.0MP3RP1: Maintenance Patch for VRTSccg	200	X	X	X
139753-02	VRTScavf 5.0MP3RP2: Maintenance Patch for Cluster Server Agents 5.0	405	X		
139754-02	VRTScavf 5.0MP3RP2: Maintenance Patch for Cluster Server Agents 5.0	405		X	
139755-02	VRTScavf 5.0MP3RP2: Maintenance Patch for Cluster Server Agents 5.0	405			X
141745-01	VCS Agents for VVR 5.0: MP3RP2 for VVR 5.0	110	X	X	X

Solaris x64

[Table 1-27](#) describes the Solaris x64 Storage Foundation for Oracle RAC patches that are included in this rolling patch:

Table 1-27 SF for Oracle RAC 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
127336-03	VM 5.0_x64_MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3_x86	110000	X

Table 1-27 SF for Oracle RAC 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
127337-03	VRTSvxfs 5.0MP3RP_x86: Rolling Patch for File System 5.0MP3	13516	X
128091-02	VRTSvcsvr 5.0 MP3 RP2: Rolling patch 02 for VCS agents for Veritas Volume Replicator 5.0	200	X
139355-01	VRTSvmman 5.0MP3RP1: Rolling Patch 01 for Volume Manager 5.0MP3_x86	920	X
139360-02	SunOS 5.10: fixes for gab, llt, vxfen	7000	X
139361-02	SunOS 5.10: fixes for vcs, vcsag	86000	X
139363-02	VRTSdbms3 5.0MP3RP2_x86: Rolling Patch for Solaris 10	55	X
139371-02	VRTSdbcom 5.0MP3RP2_x86: Rolling Patch for 5.0 MP3	9846	X
139372-02	VRTSdbed 5.0MP3RP2_x86: Rolling Patch for 5.0MP3	3150	X
139738-01	VRTSdcli 5.0MP3RP2_x86: Rolling Patch 01 for VRTSdcli 5.0MP3	31000	X
139740-01	VRTSvmpro 5.0MP3RP2_x86: Rolling Patch 01 for VRTSvmpro 5.0	29000	X
139745-02	VRTSob_x86 5.0MP3RP2: Maintenance Patch for VEA Server	14000	X
139746-02	VRTSobc33_x86 5.0MP3RP2: Maintenance Patch for VEA Server	69000	X
139747-01	VRTSaa_x86 5.0MP3RP2: Maintenance Patch for VRTSaa	2000	X
139748-01	VRTSccg_x86 5.0MP3RP2: Maintenance Patch for VRTSccg	340	X
139756-01	VRTScavf 5.0MP3RP2_x86: Maintenance Patch for Cluster Server agents 5.0	9451	X

Storage Foundation for DB2 patches

This sections describes the Storage Foundation for DB2 Solaris SPARC patches.

Solaris SPARC

[Table 1-28](#) describes the Solaris SPARC Storage Foundation for DB2 patches that are included in this rolling patch:

Table 1-28 SF for DB2 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
122058-12	VRTSvxvm 5.0MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3	184000	X	X	X
123200-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.8	13516	X		
123201-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.9	13516		X	
123202-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.10	15938			X
123722-02	VRTSat 5.0MP3: Maintenance Patch for Authentication Server	39000	X	X	X
139354-01	VRTSvmman 5.0MP3RP1: Rolling Patch 01 for Volume Manager 5.0MP3	530	X	X	X
139356-02	SunOS 5.8: fixes for gab, llt, vxfen, vcs, vcsag	97000	X		
139357-02	SunOS 5.9: fixes for gab, llt, vxfen, vcs, vcsag	97000		X	
139358-02	SunOS 5.10: fixes for vcs, vcsag	88000			X
139359-02	SunOS 5.10: fixes for gab, llt, vxfen	8000			X
139362-02	VRTSdbms3 5.0MP3RP2: Rolling Patch for Solaris 8, 9 and 10	64	X	X	X
139366-03	VRTSdbcom 5.0MP3RP2 Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	2971	X	X	X

Table 1-28 SF for DB2 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
139369-02	VRTSdb2ed 5.0MP3RP2 Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	2971	X	X	X
139370-02	VRTSd2gui 5.0MP3RP2: Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	21669	X	X	X
139737-01	VRTSdcli 5.0MP3RP1: Rolling Patch 01 for VRTSdcli 5.0MP3	26300	X	X	X
139739-01	VRTSvmpro 5.0MP3RP1: Rolling Patch 01 for VRTSvmpro 5.0	19000	X	X	X
139741-02	VRTSob 5.0MP3RP2: Maintenance Patch for VEA Server	6000	X	X	X
139742-02	VRTSobc33 5.0MP3RP2: Maintenance Patch for VEA Server	23000	X	X	X
139743-01	VRTSaa 5.0MP3RP1: Maintenance Patch for VRTSaa	1000	X	X	X
139744-01	VRTSccg 5.0MP3RP1: Maintenance Patch for VRTSccg	200	X	X	X
141745-01	VCS Agents for VVR 5.0: MP3RP2 for VVR 5.0	110	X	X	X

Storage Foundation for Oracle patches

This sections describes the Storage Foundation for Oracle Solaris SPARC and x64 patches.

Solaris SPARC

[Table 1-29](#) describes the Solaris SPARC Storage Foundation for Oracle patches that are included in this rolling patch:

Table 1-29 SF for Oracle 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
122058-12	VRTSvxvm 5.0MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3	184000	X	X	X
123200-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.8	13516	X		
123201-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.9	13516		X	
123202-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.10	15938			X
123722-02	VRTSat 5.0MP3: Maintenance Patch for Authentication Server	39000	X	X	X
139354-01	VRTSvmman 5.0MP3RP1: Rolling Patch 01 for Volume Manager 5.0MP3	530	X	X	X
139356-02	SunOS 5.8: fixes for gab, llt, vxfen, vcs, vcsag	97000	X		
139357-02	SunOS 5.9: fixes for gab, llt, vxfen, vcs, vcsag	97000		X	
139358-02	SunOS 5.10: fixes for vcs, vcsag	88000			X
139359-02	SunOS 5.10: fixes for gab, llt, vxfen	8000			X
139362-01	VRTSdbms3 5.0MP3RP2: Rolling Patch for Solaris 8, 9 and 10	64	X	X	X
139366-03	VRTSdbcom 5.0MP3RP2 Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	2971	X	X	X

Table 1-29 SF for Oracle 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
139367-02	VRTSdbed 5.0MP3RP2 Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	3150	X	X	X
139368-01	VRTSorgui 5.0MP3RP2 Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	18974	X	X	X
139737-01	VRTSdcli 5.0MP3RP1: Rolling Patch 01 for VRTSdcli 5.0MP3	26300	X	X	X
139739-01	VRTSvmpro 5.0MP3RP1: Rolling Patch 01 for VRTSvmpro 5.0	19000	X	X	X
139741-02	VRTSob 5.0MP3RP2: Maintenance Patch for VEA Server	6000	X	X	X
139742-02	VRTSobc33 5.0MP3RP2: Maintenance Patch for VEA Server	23000	X	X	X
139743-01	VRTSaa 5.0MP3RP1: Maintenance Patch for VRTSaa	1000	X	X	X
139744-01	VRTSccg 5.0MP3RP1: Maintenance Patch for VRTSccg	200	X	X	X
141745-01	VCS Agents for VVR 5.0: MP3RP2 for VVR 5.0	110	X	X	X

Solaris x64

[Table 1-30](#) describes the Solaris x64 Storage Foundation for Oracle patches that are included in this rolling patch:

Table 1-30 SF for Oracle 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
127336-03	VM 5.0_x64_MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3_x86	110000	X
127337-03	VRTSvxfs 5.0MP3RP_x86: Rolling Patch for File System 5.0MP3	13516	X
128091-02	VRTSvcsvr 5.0 MP3 RP2: Rolling patch 02 for VCS agents for Veritas Volume Replicator 5.0	200	X

Table 1-30 SF for Oracle 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
139355-01	VRTSvmmman 5.0MP3RP1: Rolling Patch 01 for Volume Manager 5.0MP3_x86	920	X
139360-02	SunOS 5.10: fixes for gab, llt, vxfen	7000	X
139361-02	SunOS 5.10: fixes for vcs, vcsag	86000	X
139363-01	VRTSdbms3 5.0MP3RP2_x86: Rolling Patch for Solaris 10	55	X
139371-01	VRTSdbcom 5.0MP3RP2_x86: Rolling Patch for 5.0 MP3	9846	X
139372-02	VRTSdbed 5.0MP3RP2_x86: Rolling Patch for 5.0MP3	3150	X
139373-01	VRTSorgui 5.0MP3RP2_x86: Rolling Patch for Solaris 10	9747	X
139738-01	VRTSdcli 5.0MP3RP1_x86: Rolling Patch 01 for VRTSdcli 5.0MP3	31000	X
139740-01	VRTSvmpro 5.0MP3RP1_x86: Rolling Patch 01 for VRTSvmpro 5.0	29000	X
139745-02	VRTSob_x86 5.0MP3RP2: Maintenance Patch for VEA Server	14000	X
139746-02	VRTSobc33_x86 5.0MP3RP2: Maintenance Patch for VEA Server	69000	X
139747-01	VRTSaa_x86 5.0MP3RP2: Maintenance Patch for VRTSaa	2000	X
139748-01	VRTSccg_x86 5.0MP3RP2: Maintenance Patch for VRTSccg	340	X

Storage Foundation for Sybase patches

This sections describes the Storage Foundation for Sybase Solaris SPARC and x64 patches.

Solaris SPARC

[Table 1-31](#) describes the Solaris SPARC Storage Foundation for Sybase patches that are included in this rolling patch:

Table 1-31 SF for Sybase 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
122058-12	VRTSvxvm 5.0MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3	184000	X	X	X
123200-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.8	13516	X		
123201-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.9	13516		X	
123202-05	VRTSvxfs 5.0MP3RP2: Rolling Patch for File System 5.0MP3 - Sun 5.9	15938			X
123722-02	VRTSat 5.0MP3: Maintenance Patch for Authentication Server	39000	X	X	X
139354-01	VRTSvmman 5.0MP3RP2: Rolling Patch 01 for Volume Manager 5.0MP3	530	X	X	X
139356-02	SunOS 5.8: fixes for gab, llt, vxfen, vcs, vcsag	97000	X		
139357-02	SunOS 5.9: fixes for gab, llt, vxfen, vcs, vcsag	97000		X	
139358-02	SunOS 5.10: fixes for vcs, vcsag	88000			X
139359-02	SunOS 5.10: fixes for gab, llt, vxfen	8000			X
139737-01	VRTSdcli 5.0MP3RP1: Rolling Patch 01 for VRTSdcli 5.0MP3	26300	X	X	X
139739-01	VRTSvmpro 5.0MP3RP1: Rolling Patch 01 for VRTSvmpro 5.0	19000	X	X	X

Table 1-31 SF for Sybase 5.0 MP3 RP2 Solaris SPARC patches

Patches	Description	Size in kb	Solaris 8	Solaris 9	Solaris 10
139741-02	VRTSob 5.0MP3RP2: Maintenance Patch for VEA Server	6000	X	X	X
139742-02	VRTSobc33 5.0MP3RP2: Maintenance Patch for VEA Server	23000	X	X	X
139743-01	VRTSaa 5.0MP3RP1: Maintenance Patch for VRTSaa	1000	X	X	X
139744-01	VRTSccg 5.0MP3RP1: Maintenance Patch for VRTSccg	200	X	X	X
141272-01	VRTSybed 5.0MP3RP2: Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	22	X	X	X
141279-01	VRTSmapro 5.0MP3RP2: Rolling Patch for 5.0MP3 for Solaris 8, 9 and 10	26	X	X	X
141745-01	VCS Agents for VVR 5.0: MP3RP2 for VVR 5.0	110	X	X	X

Solaris x64

[Table 1-32](#) describes the Solaris x64 Storage Foundation for Sybase patches that are included in this rolling patch:

Table 1-32 SF for Sybase 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
127336-03	VM 5.0_x64_MP3RP2: Rolling Patch 02 for Volume Manager 5.0MP3_x86	110000	X
127337-03	VRTSvxfs 5.0MP3RP_x86: Rolling Patch for File System 5.0MP3	13516	X
128091-02	VRTSvcsvr 5.0 MP3 RP2: Rolling patch 02 for VCS agents for Veritas Volume Replicator 5.0	200	X
139355-01	VRTSvmman 5.0MP3RP1: Rolling Patch 01 for Volume Manager 5.0MP3_x86	920	X
139360-02	SunOS 5.10: fixes for gab, llt, vxfen	7000	X

Table 1-32 SF for Sybase 5.0 MP3 RP2 Solaris x64 patches

Patches	Description	Size in kb	Solaris 10
139361-02	SunOS 5.10: fixes for vcs, vcsag	86000	X
139738-01	VRTSdcli 5.0MP3RP1_x86: Rolling Patch 01 for VRTSdcli 5.0MP3	31000	X
139740-01	VRTSvmpro 5.0MP3RP1_x86: Rolling Patch 01 for VRTSvmpro 5.0	29000	X
139745-02	VRTSob_x86 5.0MP3RP2: Maintenance Patch for VEA Server	14000	X
139746-02	VRTSobc33_x86 5.0MP3RP2: Maintenance Patch for VEA Server	69000	X
139747-01	VRTSaa_x86 5.0MP3RP2: Maintenance Patch for VRTSaa	2000	X
139748-01	VRTSccg_x86 5.0MP3RP2: Maintenance Patch for VRTSccg	340	X
141280-01	VRTSmapro 5.0MP3RP2_x86: Rolling Patch for 5.0MP3 for Solaris 10	13	X
141281-01	VRTSsybed 5.0MP3RP2_x86: Rolling Patch for 5.0MP3 for Solaris 10	9	X

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://exftpp.symantec.com/pub/support/products/Foundation_Suite/306937.pdf](http://exftpp.symantec.com/pub/support/products/Foundation_Suite/306937.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_sol.pdf](http://ftp.entsupport.symantec.com/pub/support/documentation/sf_notes_sol.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 67.

3 Run the one of the following commands to upgrade to 5.0 MP3 RP2 manually.

- For Storage Foundation on Solaris 8 SPARC:

```
# patchadd -M patch_dir 122058-12 123088-03 123200-05 \  
123722-02 123823-05 139354-01 139356-02 139737-01 \  
139739-01 139741-02 139742-02 139743-01 139744-01 \  
139753-02 140657-01 140661-01 141279-01 141745-01
```

- For Storage Foundation on Solaris 9 SPARC:

```
# patchadd -M patch_dir 122058-12 123089-03 123201-05 \  
123722-02 123823-05 139354-01 139357-02 139737-01 \  
139739-01 139741-02 139742-02 139743-01 139744-01 \  
139754-02 140657-01 140661-01 141279-01 141745-01
```

- For Storage Foundation on Solaris 10 SPARC:

```
# patchadd -M patch_dir 122058-12 123090-03 123202-05 \  
123722-02 123823-05 139354-01 139358-02 139359-02 \  
139737-01 139739-01 139741-02 139742-02 139743-01 \  
139744-01 139754-02 140657-01 140661-01 141279-01 141745-01
```

- For Storage Foundation on Solaris 10 x64:

```
# patchadd -M patch_dir 127336-03 127337-03 127362-03 \  
128091-02 139355-01 139360-02 139361-02 139738-01 \  
139740-01 139745-02 139746-02 139747-01 139748-01 \  
140658-01 140662-01 141280-01
```

- For Storage Foundation Cluster File System on Solaris 8 SPARC:

```
# patchadd -M patch_dir 122058-12 123088-03 123200-05 \  
123722-02 123823-05 139354-01 139356-02 139737-01 \  
139739-01 139741-02 139742-02 139743-01 139744-01 \  
139753-02 141745-01
```

- For Storage Foundation Cluster File System on Solaris 9 SPARC:

```
# patchadd -M patch_dir 122058-12 123089-03 123201-05 \  
123722-02 123823-05 139354-01 139357-02 139737-01 \  
139739-01 139741-02 139742-02 139743-01 139744-01 \  
139754-02 141745-01
```

- For Storage Foundation Cluster File System on Solaris 10 SPARC:

```
# patchadd -M patch_dir 122058-12 123090-03 123202-05 \  
123722-02 123823-05 139354-01 139358-02 139359-02 \  
139737-01 139739-01 139741-02 139742-02 139743-01 \  
139744-01 139755-02 141745-01
```

- For Storage Foundation Cluster File System on Solaris 10 x64:

```
# patchadd -M patch_dir 127336-03 127337-03 127341-03 \  
127362-03 128091-02 139355-01 139360-02 139361-02 \  
139738-01 139740-01 139745-02 139746-02 139747-01 \  
139748-01 139756-02 140658-01 140662-01
```

where *patch_dir* is the name of the patch directory where the patch resides.

See [“Patches included in this rolling patch”](#) on page 41.

- 4 Reboot the nodes:

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

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

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

For other products, substitute the appropriate script for `installsf` such as `installsfcfs` for the Storage Foundation Cluster File System software. See the *Veritas Storage Foundation 5.0 MP3 Installation Guide* for more information on configuring Storage Foundation and High Availability products.

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_sol.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_sol.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 67.
- 3 On each node, run one of the following commands to upgrade to 5.0 MP3 RP2 manually.

Note: You must add patch 139742-02 prior to adding patch 139741-02 or the installation will fail. If you are installing patch 122058-12 on Solaris 10 Update 5, you must also install patch 128306-05 or higher.

- For Storage Foundation for Oracle RAC on Solaris 8 SPARC:

```
# patchadd -M patch_dir 122058-12 123200-05 123722-02 \  
139354-01 139356-02 139362-01 139366-03 139367-01 \  
139737-01 139739-01 139742-02 139741-02 139743-01 \  
139744-01 139753-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 9 SPARC:

```
# patchadd -M patch_dir 122058-12 123201-05 123722-02 \  
139354-01 139357-02 139362-01 139366-03 139367-01 \  
139737-01 139739-01 139742-02 139741-02 139743-01 \  
139744-01 139754-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 10 SPARC:

```
# patchadd -M patch_dir 122058-12 123202-05 123722-02 \  
128306-05 139354-01 139358-02 139359-02 139362-01 \  
139366-03 139367-01 139737-01 139739-01 139742-02 \  
139741-02 139743-01 139744-01 139755-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 10 x64:

```
# patchadd -M patch_dir 127336-03 127337-03 127341-03 \  
127362-03 128091-02 139355-01 139360-02 139361-02 \  
139363-02 139371-02 139372-02 139738-01 139740-01 \  
139745-02 139746-02 139747-01 139748-01 139756-02 \  
140658-01 140662-01 141288-02
```

where *patch_dir* is the name of the patch directory where the patch resides.

See [“Patches included in this rolling patch”](#) on page 41.

- 4 Restart the hosts.
- 5 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 82.

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.

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

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

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

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

```
# 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/init.d/vxfen stop  
# /etc/init.d/gab stop  
# /etc/init.d/S7011t 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 Run the one of the following commands to upgrade to 5.0 MP3 RP2.

Note: For Storage Foundation for Oracle RAC, you must add patch 139742-02 prior to adding patch 139741-02 or the installation will fail. If you are installing patch 122058-12 on Solaris 10 Update 5, you must also install patch 128306-05 or higher.

- For Storage Foundation on Solaris 8 SPARC:

```
# patchadd -M patch_dir 122058-12 123088-03 123200-05 \  
123722-02 123823-05 139354-01 139356-02 139737-01 \  
139739-01 139741-02 139742-02 139743-01 139744-01 \  
139753-02 140657-01 140661-01 141279-01 141745-01
```
- For Storage Foundation on Solaris 9 SPARC:

```
# patchadd -M patch_dir 122058-12 123089-03 123201-05 \  
123722-02 123823-05 139354-01 139357-02 139737-01 \  
139739-01 139741-02 139742-02 139743-01 139744-01 \  
139754-02 140657-01 140661-01 141279-01 141745-01
```
- For Storage Foundation on Solaris 10 SPARC:

```
# patchadd -M patch_dir 122058-12 123090-03 123202-05 \  
123722-02 123823-05 139354-01 139358-02 139359-02 \  
139737-01 139739-01 139741-02 139742-02 139743-01 \  
139744-01 139754-02 140657-01 140661-01 141279-01 141745-01
```
- For Storage Foundation on Solaris 10 x64:

```
# patchadd -M patch_dir 127336-03 127337-03 127362-03 \  
128091-02 139355-01 139360-02 139361-02 139738-01 \  
139740-01 139745-02 139746-02 139747-01 139748-01 \  
140658-01 140662-01 141280-01
```
- For Storage Foundation for Oracle RAC on Solaris 8 SPARC:

```
# patchadd -M patch_dir 122058-12 123200-05 123722-02 \  
139354-01 139356-02 139362-01 139366-03 139367-01 \  
139737-01 139739-01 139742-02 139741-02 139743-01 \  
139744-01 139753-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 9 SPARC:

```
# patchadd -M patch_dir 122058-12 123201-05 123722-02 \  
139354-01 139357-02 139362-01 139366-03 139367-01 \  
139737-01 139739-01 139742-02 139741-02 139743-01 \  
139744-01 139754-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 10 SPARC:

```
# patchadd -M patch_dir 122058-12 123202-05 123722-02 \  
128306-05 139354-01 139358-02 139359-02 139362-01 \  
139366-03 139367-01 139737-01 139739-01 139742-02 \  
139741-02 139743-01 139744-01 139755-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 10 x64:

```
# patchadd -M patch_dir 127336-03 127337-03 127341-03 \  
127362-03 128091-02 139355-01 139360-02 139361-02 \  
139363-02 139371-02 139372-02 139738-01 139740-01 \  
139745-02 139746-02 139747-01 139748-01 139756-02 \  
140658-01 140662-01 141288-02
```
- For Storage Foundation for DB2 on Solaris 8 SPARC:

```
# patchadd -M patch_dir 122058-12 123200-05 123722-02 \  
139354-01 139356-02 139362-02 139366-03 139369-02 \  
139370-02 139737-01 139739-01 139741-02 139742-02 \  
139743-01 139744-01 141745-01
```

- For Storage Foundation for DB2 on Solaris 9 SPARC:

```
# patchadd -M patch_dir 122058-12 123201-05 123722-02 \  
139354-01 139357-02 139362-02 139366-03 139369-02 \  
139370-02 139737-01 139739-01 139741-02 139742-02 \  
139743-01 139744-01 141745-01
```
- For Storage Foundation for DB2 on Solaris 10 SPARC:

```
# patchadd -M patch_dir 122058-12 123202-05 123722-02 \  
139354-01 139358-02 139359-02 139362-02 139366-03 \  
139369-02 139370-02 139737-01 139739-01 139741-02 \  
0139742-02 139743-01 139744-01 141745-01
```

where *patch_dir* is the name of the patch directory where the patch resides.

See [“Patches included in this rolling patch”](#) on page 41.

- 15 After all of the nodes in the selected group are upgraded, shut down and reboot each of the nodes. After the nodes come up, application failover capability is available for that group of nodes.
- 16 Make the VCS configuration writable again from any node in the selected group:

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

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

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

```
# hagrps -switch service_group -to nodename
```
- 20 Repeat [step 9](#) through [step 19](#) for the second group of nodes.
- 21 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 37 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 19 or [“Storage Foundation for DB2 fixed issues”](#) on page 20 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.

```
# hagrp -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 30](#) for these nodes.
- 8 On each node in the selected group, enter the following command to check if any Storage Checkpoints are mounted:

```
# df | 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:

```
# df | 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 `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.

- 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 ODM, cluster fencing, GAB, and LLT in the following order:
- ```
/etc/init.d/odm stop
/etc/init.d/vxfen stop
/etc/init.d/gab stop
/etc/init.d/llt 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 Run one of the following commands to upgrade to 5.0 MP3 RP2.

- For Storage Foundation Cluster File System on Solaris 8 SPARC:  

```
patchadd -M patch_dir 122058-12 123088-03 123200-05 \
123722-02 123823-05 139354-01 139356-02 139737-01 \
139739-01 139741-02 139742-02 139743-01 139744-01 \
139753-02 141745-01
```
- For Storage Foundation Cluster File System on Solaris 9 SPARC:  

```
patchadd -M patch_dir 122058-12 123089-03 123201-05 \
123722-02 123823-05 139354-01 139357-02 139737-01 \
139739-01 139741-02 139742-02 139743-01 139744-01 \
139754-02 141745-01
```
- For Storage Foundation Cluster File System on Solaris 10 SPARC:  

```
patchadd -M patch_dir 122058-12 123090-03 123202-05 \
123722-02 123823-05 139354-01 139358-02 139359-02 \
139737-01 139739-01 139741-02 139742-02 139743-01 \
139744-01 139755-02 141745-01
```
- For Storage Foundation Cluster File System on Solaris 10 x64:  

```
patchadd -M patch_dir 127336-03 127337-03 127341-03 \
127362-03 128091-02 139355-01 139360-02 139361-02 \
139738-01 139740-01 139745-02 139746-02 139747-01 \
139748-01 139756-02 140658-01 140662-01
```

where *patch\_dir* is the name of the patch directory where the patch resides.

See [“Patches included in this rolling patch”](#) on page 41.

- 19 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.
- 20 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.
- 21 If necessary, reinstate any missing mount points in the `/etc/vfstab` file on each node.
- 22 Make the VCS configuration writable again from any node in the selected group:  

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

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

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

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

- 26 Bring the CVM service group online on each node in the selected group:  
`# hagrps -online cvm -sys nodename`
- 27 Restart all the volumes by entering the following command for each disk group:  
`# vxvol -g diskgroup startall`
- 28 If you stopped any RVGs in [step 10](#), restart each RVG:  
`# vxrvrg -g diskgroup start rvg_name`
- 29 Remount all VxFS file systems on all nodes in the selected group:  
`# mount /filesystem`
- 30 Remount all Storage Checkpoints on all nodes in the selected group:  
`# mount /checkpoint_name`
- 31 Repeat [step 8](#) through [step 30](#) 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 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 the group of nodes that are to be upgraded first, and follow [step 8](#) through [step 29](#) for these nodes.

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

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

```
/etc/init.d/init.crs stop
```
- 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 stop
/etc/init.d/lmx stop
```
- 11 Stop cluster fencing, ODM, and GAB:
 

```
/etc/init.d/vxfen stop
/etc/init.d/odm stop
/etc/init.d/gab stop
```
- 12 On each node, unload the vxfen, LMX, GAB, VCSMM, GMS, and GLM kernel modules if they are still loaded:
  - a Verify if the vxfen kernel module is loaded. For example:
 

```
modinfo|grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.0MP3)
```

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

```
modunload -i 210
```
  - b Verify if the LMX kernel module is loaded. For example:
 

```
modinfo | grep lmx
239 ffffffff1253000 13a30 236 1 lmx (LLT Mux '5.0MP3')
```

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

```
modunload -i 239
```
  - c Verify if the VCSMM kernel module is loaded. For example:
 

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

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

```
modunload -i 312
```
  - d Verify if the GMS kernel module is loaded. For example:
 

```
modinfo | grep gms
311 78289c91 4867 292 1 vxgms (VxGMS 5.1.0.0.TOT_nightly Solar)
```

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

```
modunload -i 311
```

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

```
modinfo | grep glm
310 78b68000 24268 291 1 vxglm (VxGLM 5.1,REV=TOT
SunOS 5.9)
```

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

```
modunload -i 310
```

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

```
modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.0MP3)
```

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

```
modunload -i 149
```

### 13 Stop LLT:

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

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

```
modinfo|grep llt
147 50ca4000 d6bc 110 1 llt (LLT 5.0MP3)
```

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

```
modunload -i 147
```

- 14 Check if the root disk of each node of the sub-cluster is under VxVM control by running this command.

```
df -v /
```

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

- a Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes

`mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

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

---

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

---

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

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

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

- 15** 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.

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

```
df | grep vxfs
```

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

```
umount /filesystem
```

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

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

There may be a pending large fileset clone removal extended operation if the `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
```

- 17 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.

- 18 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
```

- 19 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
```

- 20 On each node of the sub-cluster, run one of the following commands to upgrade to 5.0 MP3 RP2.

---

**Note:** You must add patch 139742-02 prior to adding patch 139741-02 or the installation will fail. If you are installing patch 122058-12 on Solaris 10 Update 5, you must also install patch 128306-05 or higher.

---

- For Storage Foundation for Oracle RAC on Solaris 8 SPARC:

```
patchadd -M patch_dir 122058-12 123200-05 123722-02 \
139354-01 139356-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139753-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 9 SPARC:

```
patchadd -M patch_dir 122058-12 123201-05 123722-02 \
139354-01 139357-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139754-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 10 SPARC:  

```
patchadd -M patch_dir 122058-12 123202-05 123722-02 \
128306-05 139354-01 139358-02 139359-02 139362-01 \
139366-03 139367-01 139737-01 139739-01 139742-02 \
139741-02 139743-01 139744-01 139755-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 10 x64:  

```
patchadd -M patch_dir 127336-03 127337-03 127341-03 \
127362-03 128091-02 139355-01 139360-02 139361-02 \
139363-02 139371-02 139372-02 139738-01 139740-01 \
139745-02 139746-02 139747-01 139748-01 139756-02 \
140658-01 140662-01 141288-02
```

where *patch\_dir* is the name of the patch directory where the patch resides.

See “[Storage Foundation for Oracle RAC patches](#)” on page 51.

- 21 After all of the nodes in the sub-cluster are upgraded, repeat [step 8](#) 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.

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

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

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

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

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

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

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

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

```
vxvol -g diskgroup startall
```



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

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

- 28 Remount all VxFS file systems on all nodes:

```
mount /filesystem
```

- 29 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

- 30 Switch the service group to the original node:

```
hagr -switch service_group -to nodename
```

- 31 Repeat [step 8](#) through [step 29](#) for the second group of nodes.

- 32 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

- 33 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 37 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 19 or [“Storage Foundation for DB2 fixed issues”](#) on page 20 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 a full upgrade to 5.0 MP3 RP2 on VCS cluster

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

```
hagr -state
```
- 3 Take the ClusterService service group offline if it is running. On any node, type:  

```
hagr -offline -force ClusterService -sys nodename
```
- 4 Make the VCS configuration writable. On any node, type:  

```
haconf -makerw
```
- 5 Freeze all service groups. On any node, type:  

```
hagr -freeze service_group -persistent
```

where `service_group` is the name of the service group. Note that the ClusterService group cannot be frozen.

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

```
haconf -dump -makero
```
- 7 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
```
- 8 Shut down VCS. On any node, type:  

```
/opt/VRTSvcs/bin/hastop -all -force
```
- 9 Shut down CmdServer. On each node, type:  

```
/opt/VRTSvcs/bin/CmdServer -stop
```
- 10 Verify that VCS has shut down. On any node, type:  

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

The output resembles:

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

The output shows no membership for port h.
- 11 For Solaris 10, on nodes that run non-global zones, check if the non-global zones are in the running state. Boot the non-global zones that are not in the running state.
  - Check the zone's state. On each node, type:  

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

```
zoneadm -z zone boot
```

where zone is the name of the non-global zone.
- 12 Unconfigure vxfen if the VCS cluster uses the fencing option. On each node, type:  

```
/sbin/vxfenconfig -U
```
- 13 Unload vxfen. On each node, perform the following steps:
  - Identify the vxfen kernel module, for example:  

```
modinfo|grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.0MP3)
```
  - Unload vxfen using the module number.  

```
modunload -i 210
```
- 14 Unconfigure GAB. On each node, type:  

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

- 15 Unload GAB. On each node, perform the following steps:
  - Identify the GAB kernel module. For example:

```
modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.0MP3)
```
  - Unload GAB using the module number:

```
modunload -i 149
```
- 16 Unconfigure LLT. On each node, perform the following steps:
  - Type:

```
/sbin/lltconfig -U
```
  - Type **y** on each node in response to the message.
- 17 Unload LLT. On each node, perform the following steps:
  - Identify the LLT kernel module. For example:

```
modinfo | grep llt
147 50ca4000 d6bc 110 1 llt (LLT 5.0MP3)
```
  - Unload LLT using the module number:

```
modunload -i 147
```
- 18 Change directory to the Veritas Cluster Server patches directory on the disc.
- 19 Add the VCS 5.0 MP3 RP2 patches. On each node, type:
  - For Solaris SPARC 8:

```
patchadd 139356-01
```
  - For Solaris SPARC 9:

```
patchadd 139357-01
```
  - For Solaris SPARC 10:

```
patchadd 139359-01
patchadd 139358-01
```
  - For Solaris SPARC 8, 9, 10, add the 5.0 MP3 RP2 Authentication Service patch.

```
patchadd 123722-02
```
  - For Solaris x64:

```
patchadd 139360-01
patchadd 139361-01
```
- 20 Verify that the patches have been installed. On each node, type:

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

  - a If the cluster has NFS or NFSREstart resources, 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
```

- 21 If you received any error messages when you unloaded the LLT, GAB, or VXFEN modules, reboot all the nodes in the cluster.
- 22 If you do not perform step 21, start the following VCS components manually. On each node, type:

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

You do not have to start vxfen unless you use the fencing option.
- 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 -unfreeze service_group -persistent
haconf -dump -makero
```

where *service\_group* is the name of the service group.
- 24 Bring online the ClusterService service group, if necessary. On any node type:

```
hagrps -online ClusterService -sys nodename
```

## 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/init.d/vxfen stop
/etc/init.d/gab stop
/etc/init.d/S70llt 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 Run the one of the following commands to upgrade to 5.0 MP3 RP2:

---

**Note:** For Storage Foundation for Oracle RAC, you must add patch 139742-02 prior to adding patch 139741-02 or the installation will fail. If you are installing patch 122058-12 on Solaris 10 Update 5, you must also install patch 128306-05 or higher.

---

- For Storage Foundation on Solaris 8 SPARC:
 

```
patchadd -M patch_dir 122058-12 123088-03 123200-05 \
123722-02 123823-05 139354-01 139356-02 139737-01 \
139739-01 139741-02 139742-02 139743-01 139744-01 \
139753-02 140657-01 140661-01 141279-01 141745-01
```
- For Storage Foundation on Solaris 9 SPARC:
 

```
patchadd -M patch_dir 122058-12 123089-03 123201-05 \
123722-02 123823-05 139354-01 139357-02 139737-01 \
139739-01 139741-02 139742-02 139743-01 139744-01 \
139754-02 140657-01 140661-01 141279-01 141745-01
```
- For Storage Foundation on Solaris 10 SPARC:
 

```
patchadd -M patch_dir 122058-12 123090-03 123202-05 \
123722-02 123823-05 139354-01 139358-02 139359-02 \
139737-01 139739-01 139741-02 139742-02 139743-01 \
139744-01 139754-02 140657-01 140661-01 141279-01 141745-01
```
- For Storage Foundation on Solaris 10 x64:
 

```
patchadd -M patch_dir 127336-03 127337-03 127362-03 \
128091-02 139355-01 139360-02 139361-02 139738-01 \
139740-01 139745-02 139746-02 139747-01 139748-01 \
140658-01 140662-01 141280-01
```

- For Storage Foundation for Oracle RAC on Solaris 8 SPARC:
 

```
patchadd -M patch_dir 122058-12 123200-05 123722-02 \
139354-01 139356-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139753-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 9 SPARC:
 

```
patchadd -M patch_dir 122058-12 123201-05 123722-02 \
139354-01 139357-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139754-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 10 SPARC:
 

```
patchadd -M patch_dir 122058-12 123202-05 123722-02 \
128306-05 139354-01 139358-02 139359-02 139362-01 \
139366-03 139367-01 139737-01 139739-01 139742-02 \
139741-02 139743-01 139744-01 139755-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 10 x64:
 

```
patchadd -M patch_dir 127336-03 127337-03 127341-03 \
127362-03 128091-02 139355-01 139360-02 139361-02 \
139363-02 139371-02 139372-02 139738-01 139740-01 \
139745-02 139746-02 139747-01 139748-01 139756-02 \
140658-01 140662-01 141288-02
```
- For Storage Foundation for DB2 on Solaris 8 SPARC:
 

```
patchadd -M patch_dir 122058-12 123200-05 123722-02 \
139354-01 139356-02 139362-02 139366-03 139369-02 \
139370-02 139737-01 139739-01 139741-02 139742-02 \
139743-01 139744-01 141745-01
```
- For Storage Foundation for DB2 on Solaris 9 SPARC:
 

```
patchadd -M patch_dir 122058-12 123201-05 123722-02 \
139354-01 139357-02 139362-02 139366-03 139369-02 \
139370-02 139737-01 139739-01 139741-02 139742-02 \
139743-01 139744-01 141745-01
```
- For Storage Foundation for DB2 on Solaris 10 SPARC:
 

```
patchadd -M patch_dir 122058-12 123202-05 123722-02 \
139354-01 139358-02 139359-02 139362-02 139366-03 \
139369-02 139370-02 139737-01 139739-01 139741-02 \
0139742-02 139743-01 139744-01 141745-01
```

where *patch\_dir* is the name of the patch directory where the patch resides.

See “[Patches included in this rolling patch](#)” on page 41.

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

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

```
haconf -makerw
```

- 15 Unfreeze the service group operations on each node:  

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

```
haconf -dump -makero
```
- 17 If you are currently using backup and restore for the DBED repository. Perform a full backup of the DBED repository database after completing the 5.0 MP3 RP2 installation.  
For more information see the [“Software limitations”](#) on page 37 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 19 or [“Storage Foundation for DB2 fixed issues”](#) on page 20 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:  

```
df | 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:  

```
df | 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:  

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

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

```
kill -9 pid_of_CmdServer
```

where `pid_of_CmdServer` is the process ID of `CmdServer`.
- 13 Stop ODM, cluster fencing, GAB, and LLT in the following order:  

```
/etc/init.d/odm stop
```

```
/etc/init.d/vxfen stop
```

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

```
/etc/init.d/llt 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/vxsvcctrl status
```

If the VEA service is running, stop it:  

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

16 Run one of the following commands to upgrade to 5.0 MP3 RP2.

- For Storage Foundation Cluster File System on Solaris 8 SPARC:  

```
patchadd -M patch_dir 122058-12 123088-03 123200-05 \
123722-02 123823-05 139354-01 139356-02 139737-01 \
139739-01 139741-02 139742-02 139743-01 139744-01 \
139753-02 141745-01
```
- For Storage Foundation Cluster File System on Solaris 9 SPARC:  

```
patchadd -M patch_dir 122058-12 123089-03 123201-05 \
123722-02 123823-05 139354-01 139357-02 139737-01 \
139739-01 139741-02 139742-02 139743-01 139744-01 \
139754-02 141745-01
```
- For Storage Foundation Cluster File System on Solaris 10 SPARC:  

```
patchadd -M patch_dir 122058-12 123090-03 123202-05 \
123722-02 123823-05 139354-01 139358-02 139359-02 \
139737-01 139739-01 139741-02 139742-02 139743-01 \
139744-01 139755-02 141745-01
```
- For Storage Foundation Cluster File System on Solaris 10 x64:  

```
patchadd -M patch_dir 127336-03 127337-03 127341-03 \
127362-03 128091-02 139355-01 139360-02 139361-02 \
139738-01 139740-01 139745-02 139746-02 139747-01 \
139748-01 139756-02 140658-01 140662-01
```

where *patch\_dir* is the name of the patch directory where the patch resides.

See “[Patches included in this rolling patch](#)” on page 41.

17 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.

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

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

20 Make the VCS configuration writable again from any node:  

```
haconf -makerw
```

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

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

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

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

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

```
vxrvg -g diskgroup start rvg_name
```
- 26 Remount all VxFS file systems on all nodes:  

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

```
mount /checkpoint_name
```

## Performing a full upgrade to 5.0 MP3 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 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 Stop Oracle and CRS on each node of the sub-cluster.
  - If CRS is controlled by VCS, log in as superuser on each system in the cluster and enter the following command:

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

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

```
/etc/init.d/lmx stop
```
- 9 Stop cluster fencing, ODM, and GAB:

```
/etc/init.d/vxfen stop
```

```
/etc/init.d/odm stop
```

```
/etc/init.d/gab stop
```
- 10 On each node, unload the vxfen, LMX, GAB, VCSMM, GMS, and GLM kernel modules if they are still loaded:
  - a Verify if the vxfen kernel module is loaded. For example:

```
modinfo|grep vxfen
```

210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.0MP3)  
If the vxfen kernel module is loaded then unload it. For example:

```
modunload -i 210
```
  - b Verify if the LMX kernel module is loaded. For example:

```
modinfo | grep lmx
```

239 ffffffff1253000 13a30 236 1 lmx (LLT Mux '5.0MP3')  
If the LMX kernel module is loaded then unload it. For example:

```
modunload -i 239
```
  - c Verify if the VCSMM kernel module is loaded. For example:

```
modinfo | grep vcsmm
```

312 78bc0000 43ae8 293 1 vcsmm (VRTSvcsmm 5.1)  
If the VCSMM kernel module is loaded then unload it. For example:

```
modunload -i 312
```

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

```
modinfo | grep gms
311 78289c91 4867 292 1 vxgms (VxGMS
5.1.0.0.TOT_nightly Solar)
```

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

```
modunload -i 311
```

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

```
modinfo | grep glm
310 78b68000 24268 291 1 vxglm (VxGLM 5.1,REV=TOT
SunOS 5.9)
```

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

```
modunload -i 310
```

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

```
modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.0MP3)
```

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

```
modunload -i 149
```

## 11 Stop LLT:

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

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

```
modinfo|grep llt
147 50ca4000 d6bc 110 1 llt (LLT 5.0MP3)
```

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

```
modunload -i 147
```

## 12 Check if the root disk of each node of the sub-cluster is under VxVM control by running this command.

```
df -v /
```

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

- a Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes

`mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

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

---

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

---

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

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

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

- 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:

```
df | grep vxfs
```

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

```
umount /filesystem
```

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

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

There may be a pending large fileset clone removal extended operation if the `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 On each node of the sub-cluster, run one of the following commands to upgrade to 5.0 MP3 RP2.

---

**Note:** You must add patch 139742-02 prior to adding patch 139741-02 or the installation will fail. If you are installing patch 122058-12 on Solaris 10 Update 5, you must also install patch 128306-05 or higher.

---

- For Storage Foundation for Oracle RAC on Solaris 8 SPARC:

```
patchadd -M patch_dir 122058-12 123200-05 123722-02 \
139354-01 139356-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139753-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 9 SPARC:

```
patchadd -M patch_dir 122058-12 123201-05 123722-02 \
139354-01 139357-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139754-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 10 SPARC:
 

```
patchadd -M patch_dir 122058-12 123202-05 123722-02 \
128306-05 139354-01 139358-02 139359-02 139362-01 \
139366-03 139367-01 139737-01 139739-01 139742-02 \
139741-02 139743-01 139744-01 139755-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 10 x64:
 

```
patchadd -M patch_dir 127336-03 127337-03 127341-03 \
127362-03 128091-02 139355-01 139360-02 139361-02 \
139363-02 139371-02 139372-02 139738-01 139740-01 \
139745-02 139746-02 139747-01 139748-01 139756-02 \
140658-01 140662-01 141288-02
```

where *patch\_dir* is the name of the patch directory where the patch resides.

See “[Storage Foundation for Oracle RAC patches](#)” on page 51.

- 19 After all of the nodes in the sub-cluster are upgraded, repeat [step 6](#) through [step 18](#) to upgrade the remaining part of the cluster.

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

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

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

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

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

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

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

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

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

```
vxvol -g diskgroup startall
```



- 25 If CRS is not controlled by VCS, enter the following command on each node to start CRS.
- ```
# /etc/init.d/init.crs start
```
- 26 Remount all VxFS file systems on all nodes:
- ```
mount /filesystem
```
- 27 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
```
- 28 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
- 29 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 37 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 19 or [“Storage Foundation for DB2 fixed issues”](#) on page 20 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 Check if the root disk is under VxVM control by running this command:

```
df -v /
```

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

- a Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

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

---

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

---

- b Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.
- 4 If required, you can upgrade the system at this stage, and patch it to a supported kernel version.
  - 5 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:
  - 6 Unmount all Storage Checkpoints and file systems:

```
df | grep vxfs
```

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

- 7 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:
  - a Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
  - b Use the `vrxvrg stop` command to stop each RVG individually:
 

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

```
vrxrlink -g diskgroup status rlink_name
```

---

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

---

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

```
vxvol -g diskgroup stopall
```

 Verify that no volumes remain open:
 

```
vxprint -Aht -e v_open
```
- 10 Check if the VEA service is running:
 

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

 If the VEA service is running, stop it:
 

```
/opt/VRTS/bin/vxsvcctl stop
```
- 11 Use one of the following commands to upgrade to 5.0 MP3 RP2.

---

**Note:** For Storage Foundation for Oracle RAC, you must add patch 139742-02 prior to adding patch 139741-02 or the installation will fail. If you are installing patch 122058-12 on Solaris 10 Update 5, you must also install patch 128306-05 or higher.

---

- For Storage Foundation on Solaris 8 SPARC:
 

```
patchadd -M patch_dir 122058-12 123088-03 123200-05 \
123722-02 123823-05 139354-01 139356-02 139737-01 \
139739-01 139741-02 139742-02 139743-01 139744-01 \
139753-02 140657-01 140661-01 141279-01 141745-01
```
- For Storage Foundation on Solaris 9 SPARC:
 

```
patchadd -M patch_dir 122058-12 123089-03 123201-05 \
123722-02 123823-05 139354-01 139357-02 139737-01 \
139739-01 139741-02 139742-02 139743-01 139744-01 \
139754-02 140657-01 140661-01 141279-01 141745-01
```

- For Storage Foundation on Solaris 10 SPARC:  

```
patchadd -M patch_dir 122058-12 123090-03 123202-05 \
123722-02 123823-05 139354-01 139358-02 139359-02 \
139737-01 139739-01 139741-02 139742-02 139743-01 \
139744-01 139754-02 140657-01 140661-01 141279-01 141745-01
```
- For Storage Foundation on Solaris 10 x64:  

```
patchadd -M patch_dir 127336-03 127337-03 127362-03 \
128091-02 139355-01 139360-02 139361-02 139738-01 \
139740-01 139745-02 139746-02 139747-01 139748-01 \
140658-01 140662-01 141280-01
```
- For Storage Foundation for Oracle RAC on Solaris 8 SPARC:  

```
patchadd -M patch_dir 122058-12 123200-05 123722-02 \
139354-01 139356-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139753-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 9 SPARC:  

```
patchadd -M patch_dir 122058-12 123201-05 123722-02 \
139354-01 139357-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139754-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 10 SPARC:  

```
patchadd -M patch_dir 122058-12 123202-05 123722-02 \
128306-05 139354-01 139358-02 139359-02 139362-01 \
139366-03 139367-01 139737-01 139739-01 139742-02 \
139741-02 139743-01 139744-01 139755-02 141745-01
```
- For Storage Foundation for Oracle RAC on Solaris 10 x64:  

```
patchadd -M patch_dir 127336-03 127337-03 127341-03 \
127362-03 128091-02 139355-01 139360-02 139361-02 \
139363-02 139371-02 139372-02 139738-01 139740-01 \
139745-02 139746-02 139747-01 139748-01 139756-02 \
140658-01 140662-01 141288-02
```
- For Storage Foundation for DB2 on Solaris 8 SPARC:  

```
patchadd -M patch_dir 122058-12 123200-05 123722-02 \
139354-01 139356-02 139362-02 139366-03 139369-02 \
139370-02 139737-01 139739-01 139741-02 139742-02 \
139743-01 139744-01 141745-01
```
- For Storage Foundation for DB2 on Solaris 9 SPARC:  

```
patchadd -M patch_dir 122058-12 123201-05 123722-02 \
139354-01 139357-02 139362-02 139366-03 139369-02 \
139370-02 139737-01 139739-01 139741-02 139742-02 \
139743-01 139744-01 141745-01
```

- For Storage Foundation for DB2 on Solaris 10 SPARC:

```
patchadd -M patch_dir 122058-12 123202-05 123722-02 \
139354-01 139358-02 139359-02 139362-02 139366-03 \
139369-02 139370-02 139737-01 139739-01 139741-02 \
0139742-02 139743-01 139744-01 141745-01
```

where *patch\_dir* is the name of the patch directory where the patch resides.

See “[Patches included in this rolling patch](#)” on page 41.

- 12 Shut down and restart the system.
- 13 If necessary, reinstate any missing mount points in the */etc/vfstab* file.
- 14 Restart all the volumes by entering the following command for each disk group:

```
vxvol -g diskgroup startall
```
- 15 If you stopped any RVGs in [step 7](#), restart each RVG:

```
vxrvg -g diskgroup start rvg_name
```
- 16 Remount all VxFS file systems and Storage Checkpoints:

```
mount /filesystem
mount /checkpoint_name
```
- 17 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

```
/opt/VRTS/bin/vxsvcctrl start
```
- 18 If you need to re-encapsulate and mirror the root disk, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.
- 19 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 37 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 19 or “[Storage Foundation for DB2 fixed issues](#)” on page 20 for incident 1425261.

## Verifying software versions

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

```
pkginfo -l VRTSvlic patch-id_name patch-id_name ...
```

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

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

```
hagrps -offline -force ClusterService -sys system
```
- 3 Make the VCS configuration writable. On any node, type:  

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

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

where *service\_group* is the name of the service group. Note that the ClusterService group cannot be frozen.
- 5 Save the configuration (main.cf) file with the groups frozen. On any node, type:  

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

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

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

```
/opt/VRTSvcs/bin/hastop -all -force
```
- 8 Shut down CmdServer. On each node, type:  

```
/opt/VRTSvcs/bin/CmdServer -stop
```
- 9 Verify that VCS has shut down. On any node, type:  

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

The output resembles:  
GAB Port Memberships  
Port a gen 23dc0001 membership 01

The output shows no membership for port h.
- 10 For Solaris 10, on nodes that run non-global zones, check if the non-global zones are in the running state. Boot the non-global zones that are not in the running state.
  - Check the zone's state. On each node, type:  

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

```
zoneadm -z zone boot
```

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

---

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

---

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

```
/sbin/vxfenconfig -U
```
- 12 Unload vxfen. On each node, perform the following steps:
  - Identify the vxfen kernel module, for example:  

```
modinfo | grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.0MP3RP2)
```
  - Unload vxfen using the module number.  

```
modunload -i 210
```
- 13 Unconfigure GAB. On each node, type:  

```
/sbin/gabconfig -U
```
- 14 Unload GAB. On each node, perform the following steps:
  - Identify the GAB kernel module. For example:  

```
modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.0MP3RP2)
```
  - Unload GAB using the module number:  

```
modunload -i 149
```



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

- Type:  

```
/sbin/lltconfig -U
```
- Type **y** on each node in response to the message.

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

- Identify the LLT kernel module. For example:  

```
modinfo | grep llt
147 50ca4000 d6bc 110 1 llt (LLT 5.0MP3RP2)
```
- Unload LLT using the module number:  

```
modunload -i 147
```

17 Remove the VCS 5.0 MP3 RP2 patches. On each node, type:

- For Solaris SPARC 8:  

```
patchrm 139356-01
```
- For Solaris SPARC 9:  

```
patchrm 139357-01
```
- For Solaris SPARC 10:  

```
patchrm 139358-01
patchrm 139359-01
```
- For Solaris x64:  

```
patchrm 139361-01
patchrm 139360-01
```

---

**Note:** For Solaris SPARC 8, 9, 10, if you must remove the 5.0 MP3 RP2 Authentication Service patch (123722-02), you must uninstall the entire VCS product stack, then reinstall VCS.

---

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

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

19 If the LLT, GAB, or VXFEN modules cannot be stopped or unloaded following the patch removal, reboot all nodes in the cluster.

20 If you do not perform [step 19](#), start the VCS components manually. On each node, type:

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

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

- 21 After VCS has started, perform the following steps:
  - Verify all resources have been probed. On any node, type:

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

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

where *service\_group* is the name of the service group.
- 22 Bring online the ClusterService service group, if necessary. On any node type:

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

where *system* is the node name.

## Removing 5.0 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 Check if the root disk is under VxVM control by running this command:

```
df -v /
```

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

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

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

---

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

---

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

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

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

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

```
df | grep vxfs
```

- 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 `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.

---

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

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

```
vxvol -g diskgroup stopall
```

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

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

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

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

- 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/vxsvcctrl status
```

If the VEA service is running, stop it:

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

```
cd /opt/VRTS/install
./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 Stop Oracle and CRS on each node of the sub-cluster.
  - If CRS is controlled by VCS, log in as superuser on each system in the sub-cluster and enter the following command:

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

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

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

```
hastop -local
```
- 2 Stop cluster fencing, VCSMM, LMX, ODM, and GAB:

```
/etc/init.d/vxfen stop
/etc/init.d/vcsmm stop
/etc/init.d/lmx stop
/etc/init.d/odm stop
/etc/init.d/gab stop
```
- 3 On each node, unload the vxfen, LMX, GAB, VCSMM, GMS, and GLM kernel modules if they are still loaded.
  - a Verify if the vxfen kernel module is loaded. For example:

```
modinfo|grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 5.0MP3)
```

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

```
modunload -i 210
```
  - b Verify if the LMX kernel module is loaded. For example:

```
modinfo | grep lmx
239 ffffffff1253000 13a30 236 1 lmx (LLT Mux
'5.0MP3')
```

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

```
modunload -i 239
```
  - c Verify if the VCSMM kernel module is loaded. For example:

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

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

```
modunload -i 312
```

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

```
modinfo | grep gms
311 78289c91 4867 292 1 vxgms (VxGMS
5.1.0.0.TOT_nightly Solar)
```

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

```
modunload -i 311
```

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

```
modinfo | grep glm
310 78b68000 24268 291 1 vxglm (VxGLM 5.1,REV=TOT
SunOS 5.9)
```

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

```
modunload -i 310
```

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

```
modinfo | grep gab
149 50cc6000 2b451 112 1 gab (GAB device 5.0MP3)
```

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

```
modunload -i 149
```

- 4 Stop LLT:

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

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

```
modinfo|grep llt
147 50ca4000 d6bc 110 1 llt (LLT 5.0MP3)
```

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

```
modunload -i 147
```

- 5 Uninstall the Storage Foundation for Oracle RAC 5.0 MP3 RP2 patches.

- For Storage Foundation for Oracle RAC on Solaris 8 SPARC:

```
patchrm 122058-12 123200-05 123722-02 \
139354-01 139356-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139753-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 9 SPARC:

```
patchrm 122058-12 123201-05 123722-02 \
139354-01 139357-02 139362-01 139366-03 139367-01 \
139737-01 139739-01 139742-02 139741-02 139743-01 \
139744-01 139754-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 10 SPARC:

```
patchrm 122058-12 123202-05 123722-02 \
128306-05 139354-01 139358-02 139359-02 139362-01 \
139366-03 139367-01 139737-01 139739-01 139742-02 \
139741-02 139743-01 139744-01 139755-02 141745-01
```

- For Storage Foundation for Oracle RAC on Solaris 10 x64:  
# **patchrm 127336-03 127337-03 127341-03 \  
127362-03 128091-02 139355-01 139360-02 139361-02 \  
139363-02 139371-02 139372-02 139738-01 139740-01 \  
139745-02 139746-02 139747-01 139748-01 139756-02 \  
140658-01 140662-01 141288-02**

See “[Storage Foundation for Oracle RAC patches](#)” on page 51.

Refer to the README file for the patch details.

- 6 After removing the patches, reboot the nodes:

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

- 7 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 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 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

