

Veritas Storage Foundation™ and High Availability Solutions Read This First

AIX

5.0 Maintenance Pack 3 Rolling Patch 1



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 Pack 1

Document version: 5.0MP3RP1.0

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

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

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

The Licensed Software and Documentation are deemed to be “commercial computer software” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

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

Third-party legal notices

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

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

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

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

<http://entsupport.symantec.com/docs/283978>

AIX is a registered trademark of IBM Corporation.

Licensing and registration

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

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

Technical support

For technical assistance, visit

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

Contents

Chapter 1	Veritas Storage Foundation and High Availability Solutions Read This First	
	System requirements	8
	Supported operating systems	8
	DB2 support	8
	Storage Foundation High Availability fixed issues	9
	Veritas Volume Manager fixed issues	9
	Veritas File System fixed issues	11
	Storage Foundation Cluster File System fixed issues	12
	Storage Foundation for Oracle fixed issues	12
	Storage Foundation for DB2 fixed issues	13
	Veritas Cluster Server fixed issues	13
	Storage Foundation and High Availability known issues	15
	Storage Foundation and High Availability known issues	15
	Veritas Volume Manager known issues	15
	Veritas File System known issues	16
	Storage Foundation Cluster File System known issues	16
	Storage Foundation for Oracle known issues	16
	Storage Foundation for DB2 known issues	17
	Veritas Cluster Server known issues	17
	Software limitations	19
	Storage Foundation for Oracle software limitations	19
	Storage Foundation for DB2 software limitations	19
	Downloading the rolling patch archive	20
	Filesets included in this rolling patch	20
	Veritas Cluster Server filesets	20
	File System filesets	21
	Storage Foundation Cluster File System filesets	21
	Storage Foundation for Oracle RAC filesets	22
	Storage Foundation filesets	23
	Storage Foundation for DB2 filesets	24
	Storage Foundation for Oracle filesets	25
	Volume Manager filesets	25
	Veritas Volume Replicator filesets	26

Installing the Veritas software for the first time	27
Installing Storage Foundation or Storage Foundation Cluster File System and 5.0 MP3 RP1	27
Installing Storage Foundation for Oracle RAC and 5.0 MP3 RP1	30
Prerequisites for upgrading to 5.0 MP3 RP1	33
Prerequisites on Veritas Cluster Server	33
Prerequisites on Storage Foundation	33
Prerequisites on Storage Foundation Cluster File System	34
Prerequisites on Storage Foundation for Oracle RAC	34
Prerequisites on Storage Foundation for Oracle	34
Prerequisites on Storage Foundation for Oracle HA	35
Prerequisites on Storage Foundation for DB2	35
Prerequisites on Storage Foundation for DB2 HA	35
Upgrading 5.0 MP3 to 5.0 MP3 RP1	36
Upgrading to 5.0 MP3 RP1 on a cluster	36
Upgrading to 5.0 MP3 RP1 on a standalone system	61
Verifying software versions	65
Removing 5.0 MP3 RP1	66
Removing 5.0 MP3 RP1 from Veritas Cluster Server	66
Removing 5.0 MP3 RP1 on Storage Foundation or Storage Foundation Cluster File System	71
Removing 5.0 MP3 RP1 on Storage Foundation for Oracle RAC	73
Documentation addendum	74
MemCPUAllocator agent	74

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 1 (RP1) AIX release.

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

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

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

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

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

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

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

System requirements

This section describes the system requirements for this release.

Supported operating systems

The 5.0 MP3 RP1 release operates on the architectures and operating systems shown below:

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

DB2 support

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

Storage Foundation High Availability fixed issues

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

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

Veritas Volume Manager fixed issues

[Table 1-1](#) describes fixed issues in the Veritas Volume Manager 5.0 MP3 RP1 release.

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

Incident	Description
1512521	StorageKeys: Adding compilation flag for VxVM to make function pointer references to legacy modules work in AIX 6.1 key aware code.
1455062	Fixed an issue with the master node crashing if a node leaves before responding to the MV- serialization protocol.
1449266	Fixed a panic in dmp_get_path_deferq while creating a volume.
1443752	Fixed an issue in a clustered environment the recovery of volumes having DCO v20 taking lots of time with no I/O load.
1443706	Fixed an issue in FMR3, I/Os initiating DCO updates for clearing DRL async clear region may not wait for its completion.
1441131	Fixed an issue with VxFS Corruption Detected when DCM log plex are attached with mirrored volume and VVR is not configured.
1441072	Fixed an issue with siteread policy is not honoured.
1442369	Fixed a bug in vxconfigbackupd script leading to 0 byte binconfig file being created.
1441020	Fixed a secondary panic due to double free of message with TCP protocol and 16 connection.
1440837	Fixed a panic due to memory allocation failure in interrupt context.

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

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

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

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

Veritas File System fixed issues

[Table 1-2](#) describes fixed issues in the Veritas File System 5.0 MP3 RP1 release.

Table 1-2 Veritas File System 5.0 MP3 RP1 fixed issues

Incident	Description
1517337	Fixed a panic issue with lockcount of an FS thread upon thread_terminate[]].
1487928	Fixed the build script to compile packages to enable the storage keys feature.
1423867	Fixed an issue in which the vx_convndata_files() call could take more than 10 minutes to complete.
1415188	Fixed a full fsck core dump that was due to running out of swap space and a malloc failure.

Table 1-2 Veritas File System 5.0 MP3 RP1 fixed issues

Incident	Description
1414178	Fixed an issue with VxFS using too much CPU while looking for odd-sized extents (vxi_alloc_fail).
1412465	Fixed a vxresize command failure to resize a volume, but the command could resize the file system.
1412160	Fixed a core dump caused by a VxFS function call while setting DST attributes.
1400046	Fixed an issue with the fsapadm enforceckpt mount_point command that resulted in a core dump.
1398904	Fixed an issue with VxFS filesystems temporarily hang in vx_delay().
1366772	Fixed a performance issue.

Storage Foundation Cluster File System fixed issues

Table 1-3 describes fixed issues in the Storage Foundation Cluster File System 5.0 MP3 RP1 release.

Table 1-3 Storage Foundation Cluster File System 5.0 MP3 RP1 fixed issues

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

Storage Foundation for Oracle fixed issues

Table 1-4 describes fixed issues in the Storage Foundation for Oracle 5.0 MP3 RP1 release.

Table 1-4 Storage Foundation for Oracle 5.0 MP3 RP1 fixed issues

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

Table 1-4 Storage Foundation for Oracle 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.
1425256	Support flashsnap CVM slave.

Storage Foundation for DB2 fixed issues

[Table 1-5](#) describes fixed issues in the Storage Foundation for DB2 5.0 MP3 RP1 release.

Table 1-5 Storage Foundation for DB2 5.0 MP3 RP1 fixed issues

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

Veritas Cluster Server fixed issues

[Table 1-6](#) describes fixed issues in the Veritas Cluster Server 5.0 MP3 RP1 release.

Table 1-6 Veritas Cluster Server 5.0 MP3 RP1 fixed issues

Incidents	Description
1510002	Fixed an issue where the HostMonitor tool misreported CPU as busy when it is idle in an AIX 5.0 MP3 Micropartition environment.

Table 1-6 Veritas Cluster Server 5.0 MP3 RP1 fixed issues

Incidents	Description
1469381	<p>Fixed an issue where the Share agent was 10x slower on 5.0 MP1 with 300+ Share resources in a service group.</p> <p>Note: This fix changes basic VCS functionality, it is critically important for you to implement these changes for all service groups that contain NFSRestart resources.</p> <p>You must set the value of the PreOnline attribute to 1 for all service groups that contain NFSRestart resources. Failure to set the service group's PreOnline attribute to a value of 1 results in broken NFSRestart resource configurations.</p> <p>The <code>ha</code> commands to change this attribute are:</p> <pre># haconf -makerw # hagrps -modify servicegroup_name PreOnline 1 # haconf -dump -makero</pre>
1457429	Removed the VCS NOTICE V-16-1-53021 message after the <code>hastart</code> command is run.
1424927	Optimized GAB connect messages.
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.
1404384	Global groups can switch over to a node where WAC is not running, when PreSwitch is set to 1 and HAD runs properly.
1403471	Reduced time for global cluster fault detection.
1398750	<p>Added the MemCPUAllocator agent.</p> <p>See “MemCPUAllocator agent” on page 74.</p>
1397692	Removed a condition where VCS engine clients hung in connect when the target system was down.
1396639	Return code for SCSI commands are now logged before re-using a variable.
1395905	Changes implemented to close device file for device <code>vxddmconfig</code> .
1394624	LLT: fixed an issue where the <code>lltdlv</code> thread spun indefinitely.
1379299	LLT: fixed <code>llt_recordmac()</code> messages.

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](#)
- [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 RP1 hosts on Storage Foundation Manager 1.1.1 can be viewed at the following URL:

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

Veritas Volume Manager known issues

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

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

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

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

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

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

Workaround

There is no known workaround.

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

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

Workaround

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

Veritas File System known issues

There are no known issues in 5.0 MP3 RP1 release of Veritas File System.

Storage Foundation Cluster File System known issues

There are no known issues in 5.0 MP3 RP1 release of Storage Foundation Cluster File System.

Storage Foundation for Oracle known issues

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

The database fails over during Flashsnap operations (1469310)

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

Workaround

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

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

Storage Foundation for DB2 known issues

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

The database fails over during Flashsnap operations (1475719)

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

Workaround

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

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

Veritas Cluster Server known issues

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

Unloading of LLT is broken (1111685)

All operations that call for unloading LLT require the nodes to be rebooted.

Workaround

You cannot reload LLT and dependent drivers.

GAB membership to VCS (Port h) may not form if VCS is brought online before the Port a membership forms (1509742)

In certain situations where "Port h" may not get seeded, `ha` commands cannot work. Some examples of these situations follow:

- The cluster is starting up from a completely powered off state, and one of the nodes in the cluster does not come up.
- In a Minimal Downtime Upgrade when you unconfigure and reconfigure GAB on the upgraded nodes, then "Port h" for VCS engine does not get seeded.

As a result of this issue the VCS engine runs, but when you issue any `ha` commands, you receive errors.

For example, if you `grep` for the `had` process from `ps` output, you see output similar to the following:

```
# ps -ef | grep -i had
root      6612      1  0 12:20 ?          00:00:00
/opt/VRTSvcs/bin/had
root      6637      1  0 12:20 ?          00:00:00
/opt/VRTSvcs/bin/hashadow
root      7228  7192  0 12:23 pts/1    00:00:00 grep -i had
```

When you attempt an `ha` command—for example `hastatus`—you get error messages similar to the following:

```
# hastatus -summ
VCS ERROR V-16-1-10600 Cannot connect to VCS engine
VCS WARNING V-16-1-11046 Local system not available
```

Workaround

Perform the following steps to resolve this issue:

- 1 On the node where you see the `ha` command failure, stop VCS.
Note that you may have to kill the `had` process using the `kill` command.
- 2 Ensure that "Port a" membership is formed on that node. Run the `gabconfig` command with the `-a` option to verify "Port a" membership.

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen      220f03 membership 01
```
- 3 If "Port a" membership is not yet formed, run the following command on that node.

```
# gabconfig -x
```
- 4 Once you have ensured "Port a" membership, restart VCS.

Contact Symantec Technical Support for information on the hotfix for this issue.

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 RP1. Otherwise, prior backups cannot be restored using the 5.0 MP3 RP1 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 12 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 RP1. Otherwise, prior backups cannot be restored using the 5.0 MP3 RP1 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 13 for incident 1425261.

Downloading the rolling patch archive

The patches comprising the 5.0 MP3 RP1 release are available for download from the Veritas website. After downloading the 5.0 MP3 RP1 file, use the `tar -z` command to uncompress and extract the archive.

For the 5.0 MP3 RP1 download archive and instructions, see the following TechNote on the Symantec Technical Support website:

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

Filesets included in this rolling patch

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

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

Veritas Cluster Server filesets

[Table 1-7](#) describes the VCS filesets that are included in this rolling patch:

Table 1-7 VCS filesets

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

File System filesets

Table 1-8 describes the File System filesets that are included in this rolling patch:

Table 1-8 File System filesets

Filesets
VRTSdcli.bff
VRTSvmpro.bff
VRTSvxfs.5.0.3.100-incr.bff
VRTSvxvm.05.00.0003.0100.bff

Storage Foundation Cluster File System filesets

Table 1-9 describes the SFCFS filesets that are included in this rolling patch:

Table 1-9 SFCFS filesets

Filesets
VRTSaa.bff
VRTSccg.bff
VRTSdcli.bff
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSob.bff
VRTSobc33.bff
VRTSvcsag.rte.bff
VRTSvcs.rte.bff
VRTSvmpro.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.100-incr.bff
VRTSvxvm.05.00.0003.0100.bff

Storage Foundation for Oracle RAC filesets

Table 1-10 describes the SF for Oracle RAC filesets that are included in this rolling patch:

Table 1-10 SF for Oracle RAC filesets

Filesets
VRTSaa.bff
VRTSccg.bff
VRTSdbcom.bff
VRTSdbed.bff
VRTSdbms3-3.0.83.3.bff
VRTSdcli.bff
VRTSgab.rte.bff
VRTSlit.rte.bff
VRTSob.bff
VRTSobc33.bff
VRTSorgui.bff
VRTSvcsg.rte.bff
VRTSvcsg.rte.bff
VRTSvmpro.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.100-incr.bff
VRTSvxvm.05.00.0003.0100.bff

Storage Foundation filesets

Table 1-11 describes the Storage Foundation filesets that are included in this rolling patch:

Table 1-11 Storage Foundation filesets

Filesets
VRTSaa.bff
VRTSccg.bff
VRTSdbms3-3.0.83.3.bff
VRTSdcli.bff
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSob.bff
VRTSobc33.bff
VRTSvcsg.rte.bff
VRTSvcsg.rte.bff
VRTSvmpro.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.100-incr.bff
VRTSvxvm.05.00.0003.0100.bff

Storage Foundation for DB2 filesets

Table 1-12 describes the Storage Foundation for DB2 filesets that are included in this rolling patch:

Table 1-12 Storage Foundation for DB2 filesets

Filesets
VRTSaa.bff
VRTSccg.bff
VRTSd2gui.bff
VRTSdb2ed.bff
VRTSdbcom.bff
VRTSdbms3-3.0.83.3.bff
VRTSdcli.bff
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSob.bff
VRTSobc33.bff
VRTSvcsag.rte.bff
VRTSvcs.rte.bff
VRTSvmpro.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.100-incr.bff
VRTSvxvm.05.00.0003.0100.bff

Storage Foundation for Oracle filesets

Table 1-13 describes the SF for Oracle filesets that are included in this rolling patch:

Table 1-13 Storage Foundation for Oracle filesets

Filesets
VRTSdbcom.bff
VRTSdbed.bff
VRTSdbms3-3.0.83.3.bff
VRTSdcli.bff
VRTSgab.rte.bff
VRTSllt.rte.bff
VRTSorgui.bff
VRTSvcsag.rte.bff
VRTSvcs.rte.bff
VRTSvmpro.bff
VRTSvxfen.rte.bff
VRTSvxfs.5.0.3.100-incr.bff
VRTSvxvm.05.00.0003.0100.bff

Volume Manager filesets

Table 1-14 describes the Volume Manager filesets that are included in this rolling patch:

Table 1-14 Volume Manager filesets

Filesets
VRTSdbms3-3.0.83.3.bff
VRTSdcli.bff
VRTSvmpro.bff
VRTSvxvm.05.00.0003.0100.bff

Veritas Volume Replicator filesets

Table 1-15 describes the Veritas Volume Replicator filesets that are included in this rolling patch:

Table 1-15 Veritas Volume Replicator filesets

Filesets
VRTSaa.bff
VRTSccg.bff
VRTSdcli.bff
VRTSob.bff
VRTSobc33.bff
VRTSvmpro.bff
VRTSvxvm.05.00.0003.0100.bff

Installing the Veritas software for the first time

This section describes how to install a Storage Foundation and High Availability Solutions product for the first time on a host and install 5.0 MP3 RP1.

- [Installing Storage Foundation or Storage Foundation Cluster File System and 5.0 MP3 RP1](#)
- [Installing Storage Foundation for Oracle RAC and 5.0 MP3 RP1](#)

Installing Storage Foundation or Storage Foundation Cluster File System and 5.0 MP3 RP1

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 RP1.

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

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

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

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

To install the Storage Foundation or Storage Foundation Cluster File System and 5.0 MP3 RP1

- 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 `installsfdfs` for the Storage Foundation Cluster File System software.

- 2 Review the installation prerequisites for upgrading to 5.0 MP3 RP1. See [“Prerequisites for upgrading to 5.0 MP3 RP1”](#) on page 33.

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

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

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

where *product* is the product directory name such as

storage_foundation or *sfcfs*.

where *working_directory* is your working directory.

- 4 Uncompress each of the bundled filesets.

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

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

- On one node, type:

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

Review the verification data and confirm that:

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

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

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

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

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

```
*****
installp PREVIEW:  installation will not actually occur.
*****
```

```
+-----+
                        Pre-installation Verification...
+-----+
```

```
Verifying selections...done
```

```
Verifying requisites...done
```

```
Results...
```

```
SUCSESSES
```

```
-----
```

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

```
Selected Filesets
```

```
-----
```

```
.
.
.
```

End of the Success section.

```

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

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

```

6 Apply the patches to all or selected cluster nodes.

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

7 Restart the host.

8 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 `installsfscfs` 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 RP1

This section describes how to install Storage Foundation for Oracle RAC for the first time on a host and install 5.0 MP3 RP1.

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

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

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

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

To install Storage Foundation for Oracle RAC and 5.0 MP3 RP1

- 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 RP1.
See “Prerequisites for upgrading to 5.0 MP3 RP1” on page 33.
- 3 Copy the contents of the *product/patches* directory to local or shared storage.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

6 Apply the patches to all or selected cluster nodes.

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

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

7 Restart all the cluster nodes.

```
# shutdown -r now
```

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

```
# cd /opt/VRTS/install
# ./installsfrac -configure [-rsh] node1 node2 ... nodeN
```

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

Prerequisites for upgrading to 5.0 MP3 RP1

This section describes the prerequisites for upgrading 5.0 MP3 RP1.

- [Prerequisites on Veritas Cluster Server](#)
- [Prerequisites on Storage Foundation](#)
- [Prerequisites on Storage Foundation Cluster File System](#)
- [Prerequisites on Storage Foundation for Oracle RAC](#)

Prerequisites on Veritas Cluster Server

You must have the following operating system and VCS release installed on the system before you upgrade to 5.0 MP3 RP1:

Table 1-16 Operating system and VCS release

Operating system	Product release
AIX 5.3 TL7 with SP2 or later is required	VCS 5.0 MP3
AIX 6.1 TL0 with SP4 or later is required	VCS 5.0 MP3

Prerequisites on Storage Foundation

You must have the following operating system and Storage Foundation release installed on the system before you upgrade to 5.0 MP3 RP1:

Table 1-17 Operating system and Storage Foundation release

Operating system	Product release
AIX 5.3 TL7 with SP2 or later	Storage Foundation 5.0 MP3
AIX 6.1 TL0 with SP4 or later	Storage Foundation 5.0 MP3

Prerequisites on Storage Foundation Cluster File System

You must have the following operating system and SFCFS release installed on the system before you upgrade to 5.0 MP3 RP1:

Table 1-18 Operating system and SFCFS release

Operating system	Product release
AIX 5.3 TL7 with SP2 or later	SFCFS 5.0 MP3
AIX 6.1 TL0 with SP4 or later	SFCFS 5.0 MP3

Prerequisites on Storage Foundation for Oracle RAC

You must have the following operating system and SF for Oracle RAC release installed on the system before you upgrade to 5.0 MP3 RP1:

Table 1-19 Operating system and SFCFS for Oracle RAC release

Operating system	Product release
AIX 5.3 TL7 with SP2 or later	SF for Oracle RAC 5.0 MP3
AIX 6.1 TL0 with SP4 or later	SF for Oracle RAC 5.0 MP3

Prerequisites on Storage Foundation for Oracle

You must have the following operating system and SF for Oracle release installed on the system before you upgrade to 5.0 MP3 RP1:

Table 1-20 Operating system and SFCFS for Oracle release

Operating system	Product release
AIX 5.3 TL7 with SP2 or later	SF for Oracle 5.0 MP3
AIX 6.1 TL0 with SP4 or later	SF for Oracle 5.0 MP3

Prerequisites on Storage Foundation for Oracle HA

You must have the following operating system and SF for Oracle HA release installed on the system before you upgrade to 5.0 MP3 RP1:

Table 1-21 Operating system and SFCFS for Oracle HA release

Operating system	Product release
AIX 5.3 TL7 with SP2 or later	SF for Oracle HA5.0 MP3
AIX 6.1 TL0 with SP4 or later	SF for Oracle HA 5.0 MP3

Prerequisites on Storage Foundation for DB2

You must have the following operating system and SF for DB2 release installed on the system before you upgrade to 5.0 MP3 RP1:

Table 1-22 Operating system and SFCFS for DB2 release

Operating system	Product release
AIX 5.3 TL7 with SP2 or later	SF for DB2 5.0 MP3
AIX 6.1 TL0 with SP4 or later	SF for DB2 5.0 MP3

Prerequisites on Storage Foundation for DB2 HA

You must have the following operating system and SF for DB2 HA release installed on the system before you upgrade to 5.0 MP3 RP1:

Table 1-23 Operating system and SFCFS for DB2 HA release

Operating system	Product release
AIX 5.3 TL7 with SP2 or later	SF for DB2 HA 5.0 MP3
AIX 6.1 TL0 with SP4 or later	SF for DB2 HA 5.0 MP3

Upgrading 5.0 MP3 to 5.0 MP3 RP1

This section describes how to upgrade from 5.0 MP3 to 5.0 MP3 RP1 on a cluster or a standalone system.

- [Upgrading to 5.0 MP3 RP1 on a cluster](#)
Use the procedures to upgrade to 5.0 MP3 RP1 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 RP1 on a standalone system](#)
Use the procedure to upgrade to 5.0 MP3 RP1 on a system that has Storage Foundation, SF for Oracle, or SF for DB2 installed.

Upgrading to 5.0 MP3 RP1 on a cluster

Upgrading 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 upgrading 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, and leave a group of one or more nodes running.
- 3 Take the first group offline and install the software patches.
- 4 Bring the first group (with the newly installed patches) online to restart cluster failover services.
- 5 Upgrade the nodes in the second group, and bring them online. The cluster is fully restored.

Depending on your cluster's configuration, select one of the following procedures to upgrade to 5.0 MP3 RP1:

- [Upgrading to 5.0 MP3 RP1 on VCS cluster](#)
- [Upgrading to 5.0 MP3 RP1 on VCS cluster using a minimal downtime upgrade](#)
- [Upgrading to 5.0 MP3 RP1 on a Storage Foundation HA cluster](#)
- [Upgrading to 5.0 MP3 RP1 on a Storage Foundation Cluster File System cluster](#)
- [Upgrading to 5.0 MP3 RP1 on a Storage Foundation for Oracle RAC cluster](#)

Upgrading to 5.0 MP3 RP1 on VCS cluster

Use the following procedures to upgrade your cluster to VCS 5.0 MP3 RP1.

To perform pre-upgrade tasks

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

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

```
# for bff in working_directory/*.gz; do gunzip $bff; done
```
- 3 List the service groups in your cluster and their status. On any node, type:

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

```
# hagr -offline -force ClusterService -any
```
- 5 Make the VCS configuration writable. On any node, type:

```
# haconf -makerw
```
- 6 Freeze all service groups except ClusterService. On any node, type:

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

You can ignore the failure to freeze the ClusterService group warning.
- 7 Save the configuration (main.cf) file with the groups frozen. On any node, type:

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

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

To perform a shutdown of VCS cluster and drivers

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

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

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

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


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

 The output resembles:


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

 Output for membership for port h does not appear.
 - On each node, run the command


```
# ps -ef | egrep "[ ](had|hashadow|CmdServer)"
```

 Terminate any instances of had, hashadow, or CmdServer that still run after 60 seconds.
- 4 Unconfigure vxfen if the VCS cluster uses the fencing option. On each node, type:


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


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


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

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


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


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


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

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


```
# /sbin/lltconfig -U
```
- 9 Unload the LLT portable streams driver. On each node, type:


```
# /usr/sbin/strload -u -d /usr/lib/drivers/pse/llt
```

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

```
*****  
installp PREVIEW: installation will not actually occur.  
*****  
+-----+  
                Pre-installation Verification...  
+-----+  
Verifying selections...done  
Verifying requisites...done  
Results...  
  
SUCSESSES  
-----  
  
Filesets listed in this section passed pre-installation verification  
and will be installed.  
  
Selected Filesets  
-----  
VRTSgab.rte 5.0.3.100          # Veritas Group Membership and...  
VRTSl1t.rte 5.0.3.100          # Veritas Low Latency Transpor...  
VRTSvc.s.rte 5.0.3.100         # Veritas Cluster Server 5.OMP...  
VRTSvc.sag.rte 5.0.3.100       # Veritas Cluster Server 5.OMP...  
VRTSVxfen.rte 5.0.3.100        # Veritas I/O Fencing 5.OMP3RP...  
  
End of the Success section.
```

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

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

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

(...omitted...)

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

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


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


```
# installp -agXv -d working_directory -f filesets.to.install
```
- Review the summaries at the end of each run and confirm that all of the intended patches were successfully applied.

To perform a service restart or cluster reboot

- 1 Reboot all nodes in the cluster.
- 2 After VCS has started, perform the following steps:
 - Verify all resources have been probed. On any node, type:


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


```
# haconf -makerw
# hagrps -list | sort -u +0b -1 | \
    while read grp sys ; do
        hagrps -unfreeze $grp -persistent
    done
# haconf -dump -makero
```

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

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

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

where system is the node name.

Upgrading to 5.0 MP3 RP1 on VCS cluster using a minimal downtime upgrade

When you perform a minimal downtime upgrade, you split your cluster roughly in half. Each half is then upgraded while services continue to run on the other half. You must have enough nodes in each half of the cluster to support all essential service groups for the minimal downtime upgrade to work.

Perform a minimal downtime upgrade in the following phases:

- Phase 0—back up and prepare the cluster. Divide the cluster nodes into two subsets. When you divide the cluster, make sure each subset can run essential service groups. Switch all service groups to nodes in subset 2.
- Phase 1—freeze all nodes in subset 1. Upgrade subset 1.
- Phase 2—stop VCS and all service applications on subset 2. Stop and unconfigure vxfen, gab and llt. Manually seed the nodes in subset 1 to establish cluster membership. Start VCS on subset 1 and bring online all essential service groups as needed.
- Phase 3—freeze all nodes in subset 2. Upgrade subset 2.

When both subsets have been upgraded, services can be redistributed across the entire cluster.

To perform pre-upgrade tasks

- 1 Select one or more nodes in the cluster as the subset that you want to upgrade first (subset 1).
All other nodes belong to subset 2. The selection of nodes into subsets must be such that either subset is capable of running all essential service groups.
- 2 Log on as superuser on one of the cluster nodes.
- 3 Verify that /opt/VRTS/bin is set in your PATH environment variable to execute all product commands.
- 4 Back up the llttab, llthosts, gabtab, types.cf, and main.cf files.

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

- 5 Establish which essential service groups are online on nodes in subset 1. At the shell prompt, enter:

```
# hagrps -state -sys s1_nodename1 [ s1_nodename2 ...]
```

where *s1_nodename1*, *s1_nodename2*... are all of the nodes in subset1.
 For example, if subset1 contains nodes node01 and node02:

```
# hagrps -state -sys node01 node02
#Group Attribute System Value
sg1 State node01 |ONLINE|
sg1 State node02 |OFFLINE|
sg2 State node01 |OFFLINE|
sg2 State node02 |ONLINE|
```

- 6 For each essential service group, switch the service group to one of the nodes in subset 2.

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

For example, if subset 2 contains nodes node03 and node04:

```
# hagrps -switch sg1 -to node04
# hagrps -switch sg2 -to node03
```

You cannot migrate certain parallel service groups such as VxSS.

- 7 Verify that all essential service groups are offline on all subset 1 nodes.

```
# hagrps -state -sys s1_nodename1 [ s1_nodename2 ...]
#Group Attribute System Value
sg1 State node01 |OFFLINE|
sg1 State node02 |OFFLINE|
sg2 State node01 |OFFLINE|
sg2 State node02 |OFFLINE|
```

- 8 Insert the software disc into the disc drive of one of the nodes. Mount the disc on a suitable mount point.

- 9 Navigate to the directory that contains the 5.0 MP3 RP1 patches.

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

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

- 11 Uncompress each of the bundled filesets.

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

To upgrade the cluster nodes in subset 1

- 1 Freeze the nodes in subset 1. For each node in subset 1, run the command:
hasys -freeze nodename
For example:
hasys -freeze node01
hasys -freeze node02
- 2 Stop VCS on subset 1. On each node in the subset, type:
hastop -local
- 3 Shut down CmdServer. On each node in the subset, type:
/opt/VRTSvcs/bin/CmdServer -stop
Verify that VCS has shut down. On each node in the subset, run the command
ps -ef | egrep "[](had|hashadow|CmdServer)"
Terminate any instances of had, hashadow, or CmdServer that still run after 60 seconds.
- 4 Unconfigure vxfen if the VCS cluster uses the fencing option. On each node in the subset, type:
/sbin/vxfenconfig -U
- 5 Halt and unload vxfen. On each node in the subset, perform the following steps:
 - Unload the vxfen driver:
/etc/methods/vxfenext -stop
 - Confirm that the vxfen driver is unloaded:
sleep 3; /etc/methods/vxfenext -status
vxfen: unloaded
- 6 Unconfigure GAB. On each node in the subset, type:
/sbin/gabconfig -U

- 7 Halt and unload the GAB driver. On each node in the subset, perform the following steps:
 - Unload the GAB kernel module:
`# /etc/methods/gabkext -stop`
 - Confirm that the GAB driver is unloaded:
`# sleep 3; /etc/methods/gabkext -status`
gab: unloaded
- 8 Unconfigure LLT. On each node in the subset, type:
`# /sbin/lltconfig -U`
- 9 Unload the LLT portable streams driver. On each node in the subset, type:
`# /usr/sbin/strload -u -d /usr/lib/drivers/pse/llt`
- 10 Preview the pre-installation verification testing by performing the following steps:
 - On one node, type:
`# installp -pagXv -d working_directory all`
Review the verification data and confirm that:
 - No unexpected failures occur, and that
 - All selected filesets are successfully verified.
 - Repeat the installp command on all other nodes in the subset:
`# installp -pagXv -d working_directory all`
Symantec recommends that the local directory containing the unpacked patches is accessible from all nodes.

- Copy the list of successful filesets (in bold below) to a text file `filesets.to.install`. Save this file for later use.

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

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

Selected Filesets
-----
VRTSgab.rte 5.0.3.100                # Veritas Group Membership and...
VRTSl1t.rte 5.0.3.100                # Veritas Low Latency Transpor...
VRTSvc.s.rte 5.0.3.100              # Veritas Cluster Server 5.0MP...
VRTSvc.sag.rte 5.0.3.100            # Veritas Cluster Server 5.0MP...
VRTSvc.xfen.rte 5.0.3.100           # Veritas I/O Fencing 5.0MP3RP...

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

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

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

                (...omitted...)

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

11 Apply the patches to all or selected cluster nodes.

- To apply all patches, on each node in the subset, type:

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

```
# installp -agXv -d working_directory -f filesets.to.install
```
- Review the summaries at the end of each run and confirm that all of the intended patches were successfully applied.

12 Reboot all nodes in the subset.

To switch services between subsets

- 1 Stop VCS on all nodes in the cluster. On any node in subset 2, run the command:

```
# hastop -all
```

 On all nodes, run the command:

```
# /etc/init.d/vcs.rc stop
```
- 2 Unconfigure vxfen on all nodes in subset 2. On each node in subset 2, run the command:

```
# /sbin/vxfenconfig -U
```
- 3 Unconfigure gab on all nodes in subset 2. On each node in subset 2, run the command:

```
# /sbin/gabconfig -U
```
- 4 Unconfigure llt on all nodes in subset 2. On each node in subset 2, run the command:

```
# /sbin/lltconfig -U
```
- 5 Seed the nodes in subset 1 to form a new cluster. On each node in subset 1, run the command:

```
# /sbin/gabconfig -c
```

 On one node in subset 1, run the following command:

```
# /sbin/gabconfig -c -x
```

 To verify the results, run the command on the same node as above:

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

 Output resembles the following:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 1ebf2c membership 01
```

 Each node in subset 1 should be shown in the membership for port a.
- 6 If fencing is enabled, start vxfen. On each node in subset 1, run the command:

```
# /etc/init.d/vxfen.rc start
```

- 7 Start VCS on subset 1. On each node in subset 1, run the command:
`# /etc/init.d/vcs.rc start`
- 8 If any essential services are not online, bring them online in subset 1 as appropriate.

To upgrade the cluster nodes in subset 2

- 1 Freeze the nodes in subset 2. On any node in subset 1, run the command for each node in subset 2:
`# hasys -freeze nodename`
For example:
`# hasys -freeze node03`
`# hasys -freeze node04`
- 2 Unconfigure LLT. On each node in the subset, type:
`# /sbin/lltconfig -U`
- 3 Halt and unload vxfen. On each node in the subset, perform the following steps:
 - Unload the vxfen driver:
`# /etc/methods/vxfenext -stop`
 - Confirm that the vxfen driver is unloaded:
`# sleep 3; /etc/methods/vxfenext -status`
vxfen: unloaded
- 4 Halt and unload the GAB driver. On each node in the subset, perform the following steps:
 - Unload the GAB kernel module:
`# /etc/methods/gabkext -stop`
 - Confirm that the GAB driver is unloaded:
`# sleep 3; /etc/methods/gabkext -status`
gab: unloaded
- 5 Unload the LLT portable streams driver. On each node in the subset, type:
`# /usr/sbin/strload -u -d /usr/lib/drivers/pse/llt`
- 6 Preview the pre-installation verification testing if desired.
See [“Preview the pre-installation verification testing by performing the following steps:”](#) on page 44.

- 7 Apply the patches to all nodes in subset 2.
 - To apply all patches, on each node in the subset, type:


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


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

Review the summaries at the end of each run and confirm that all of the intended patches were successfully applied.
- 8 Reboot all nodes in the subset.

To complete the minimal downtime upgrade

- 1 Unfreeze all nodes in subset 2. On any node where VCS is running, execute the command for each node in subset 2:


```
# hasys -unfreeze nodename
```
- 2 Migrate service groups onto nodes in subset 2 as desired for performance. For example:


```
# hagr -switch sg1 -to node01
# hagr -switch sg2 -to node02
```
- 3 Save the configuration.


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

Upgrading to 5.0 MP3 RP1 on a Storage Foundation HA cluster

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

To upgrade to 5.0 MP3 RP1 on a Storage Foundation HA cluster

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


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


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


```
# 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 23](#) for these nodes.
- 9 Stop VCS. Enter the following command on each node in the group that is being upgraded:

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

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

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

```
# /etc/rc.d/rc2.d/S97vxfen stop  
# /etc/rc.d/rc2.d/S92gab stop  
# /etc/rc.d/rc2.d/S701lt stop
```
- 12 If required, you can upgrade the nodes 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 started, if any, by the init scripts.
- 14 Copy the contents of the *product/patches* directory to local or shared storage.

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

where *path-to-media* is the path to the media.
where *product* is the product directory name such as *storage_foundation_for_db2* or *storage_foundation_for_oracle*.
where *working_directory* is your working directory.
- 15 Uncompress each of the bundled filesets.

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

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

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

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

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

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

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

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

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

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

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

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

17 Apply the patches to all or selected cluster nodes.

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

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

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

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

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

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

19 Run the following commands to start VCS:

```
# /etc/rc.d/rc2.d/S701lt start
# /etc/rc.d/rc2.d/S92gab start
# /etc/rc.d/rc2.d/S97vxfen start
# /etc/rc.d/rc2.d/vcs start
```

20 Make the VCS configuration writable again from any node in the upgraded group:

```
# haconf -makerw
```

21 Unfreeze the service group operations. Perform this task on each node if you had upgraded a group of nodes:

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

22 Make the VCS configuration read-only:

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

23 Switch the service group to the original node:

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

24 Repeat [step 9](#) through [step 23](#) for the second group of nodes.

25 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 RP1 installation.

For more information see the [“Software limitations”](#) on page 19 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 12 or [“Storage Foundation for DB2 fixed issues”](#) on page 13 for incident 1425261.

Upgrading to 5.0 MP3 RP1 on a Storage Foundation Cluster File System cluster

The following procedure describes upgrading on a SFCFS cluster.

To upgrading to 5.0 MP3 RP1 on a SFCFS cluster

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

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

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

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

```
# haconf -dump -makero
```
- 7 Select the group of nodes that are to be upgraded first, and follow [step 8](#) through [step 32](#) for these nodes.
- 8 Stop VCS by entering the following command on each node in the group being upgraded:

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

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

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

```
# /etc/rc.d/rc2.d/S97vxfs stop  
# /etc/rc.d/rc2.d/S92gab stop  
# /etc/rc.d/rc2.d/S70llt stop
```
- 11 If required, you can upgrade the nodes at this stage, and patch them to a supported kernel version.
See “[System requirements](#)” on page 8.

- 12 On each node, use the following command to check if any Storage Checkpoints are mounted:
- ```
mount | grep vxfs
```
- If any Storage Checkpoints are mounted, on each node in the cluster unmount all Storage Checkpoints.
- ```
# umount /checkpoint_name
```
- 13 On each node, use the following command to check if any VxFS file systems are mounted:
- ```
mount | grep vxfs
```
- a If any VxFS file systems are present, on each node in the cluster unmount all the VxFS file systems:
- ```
# umount /filesystem
```
- 14 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 `vxrvvg stop` command to stop each RVG individually:
- ```
vxrvvg -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.
-
- 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, stop all VxVM volumes by entering the following command for each disk group:
- ```
vxvol -g diskgroup stopall
```
- To verify that no volumes remain open, use the following command:
- ```
# vxprint -Aht -e v_open
```
- 17 Check if the VEA service is running:
- ```
/opt/VRTS/bin/vxsvcctl status
```
- If the VEA service is running, stop it:
- ```
# /opt/VRTS/bin/vxsvcctl stop
```

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

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

where *path-to-media* is the path to the media.
 where *product* is the product directory name such as *sfcfs*.
 where *working_directory* is your working directory.

- 19 Uncompress each of the bundled filesets.

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

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

- On one node, type:

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

Review the verification data and confirm that:

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

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

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

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

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

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

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

- 21 Apply the patches to all or selected cluster nodes.
 - To apply all patches, on each node, type:
`# installp -agXv -d working_directory all`
 - To apply only patches listed in the file `filesets.to.install`, on each node, type:
`# installp -agXv -d working_directory -f filesets.to.install`

Review the summaries at the end of each run and confirm that all of the intended patches were successfully applied.
- 22 After all the nodes in the cluster are upgraded, shut down and reboot each of the upgraded nodes.
- 23 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node.
- 24 Make the VCS configuration writable again from any node in the upgraded group:
`# haconf -makerw`
- 25 Enter the following command on each node in the upgraded group to unfreeze HA service group operations:
`# hasys -unfreeze -persistent nodename`
- 26 Make the configuration read-only:
`# haconf -dump -makero`

- 27 Switch the service group to the original node:

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

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

```
# vxvol -g diskgroup startall
```
- 30 If you stopped any RVGs in [step 14](#), restart each RVG:

```
# vxrvgs -g diskgroup start rvg_name
```
- 31 Remount all VxFS file systems on all nodes:

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

```
# mount /checkpoint_name
```
- 33 Repeat [step 8](#) through [step 32](#) for the second group of nodes.

Upgrading to 5.0 MP3 RP1 on a Storage Foundation for Oracle RAC cluster

The following procedure describes upgrading on a SF for Oracle RAC cluster.

To upgrading to 5.0 MP3 RP1 on a SFRAC cluster

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
- 3 Select the group of nodes that are to be upgraded first.
- 4 Switch the failover service group, if any, to a node in the second group that is upgraded later in the procedure.

```
# hagrps -switch service_group -to nodename
```
- 5 Follow [step 6](#) through [step 33](#) for the first group of nodes to be upgraded.
- 6 If CRS is not controlled by VCS, use the following command on each node to stop CRS.

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

On stopping CRS if any gsd relevant process remains active, you must stop that process manually.

- 7 All Oracle users must stop all applications using the CFS mounts not under VCS control.
 - a To verify that no processes use the CFS mount point, enter the following command:

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

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

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

```
# umount mount_point
```
- 9 Stop VCS by entering the following command on each node in the group being upgraded:

```
# 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 VCSMM and LMX if they are running.

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

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

```
# fsclustadm cfsdeinit
```
- 13 Stop cluster fencing, ODM, GAB, and LLT.

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

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

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

```
# /etc/rc.d/rc2.d/S701lt stop
```
- 14 If required, you can upgrade the nodes at this stage, and patch them to a supported kernel version.

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

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

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


```
# mount | grep vxfs
```

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


```
# umount /mount-point
```
 - b On each node, verify that all file systems have been cleanly unmounted:


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

A *clean_value* value of 0x5a indicates the file system is clean, 0x3c indicates the file system is dirty, and 0x69 indicates the file system is dusty. A dusty file system has pending extended operations.
 - c If a file system is not clean, enter the following commands for that file system:


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

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

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

```
file system device busy
```

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

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


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

- 16 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.
- 17 On each node, stop all VxVM volumes by entering the following command for each disk group:

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


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

```
# vxprint -Aht -e v_open
```
- 18 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
```
- 19 Copy the contents of the *product/patches* directory to local or shared storage.

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


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

```
# for bff in working_directory/*.gz; do gunzip $bff; done
```
- 21 Apply all patches, on each node, type:

```
# installp -agXv -d working_directory all
```
- 22 After all the nodes in the cluster are upgraded, shut down and reboot each of the upgraded nodes.

```
# shutdown -r now
```
- 23 If necessary, reinstate any missing mount points in the */etc/filesystems* file on each node.
- 24 On the second group of nodes, stop the failover service group. Downtime starts for failover service groups.

```
# hagrps -offline failover_service_group
```
- 25 Make the VCS configuration writable again from any node second group:

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

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

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

- 28 On the second group of nodes, shutdown VCS. Execute following command on one of the nodes.

```
# hstop -all -force
```
- 29 On the first group of nodes, start VCS on each of the nodes:

```
# hstart
```

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

```
# /etc/init.crs start
```
- 31 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
```
- 32 After all the services are up on the first group of nodes, stop VCS. Execute following command on one of the nodes.

```
# hstop -all -force
```
- 33 On the second group of nodes, start VCS on each of the nodes

```
# hstart
```
- 34 Repeat [step 6](#) through [step 22](#) for the second group of nodes.
- 35 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node in second group.
- 36 If CRS is not controlled by VCS, use the following command on each node in second group to start CRS.

```
# /etc/init.crs start
```
- 37 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
```
- 38 Make the VCS configuration writable again from any node second group:

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

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

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

```
# hstart
```

- 42 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
- 43 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 RP1 installation.
- For more information see the [“Software limitations”](#) on page 19 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 12 or [“Storage Foundation for DB2 fixed issues”](#) on page 13 for incident 1425261.

Upgrading to 5.0 MP3 RP1 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 RP1 on a standalone system

- 1 Log in as superuser.
- 2 Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
- 3 If required, you can upgrade the system at this stage, and patch it to a supported kernel version.
- 4 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```
- 5 Unmount all Storage Checkpoints and file systems:

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

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

```
# vxrvvg -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, use the following command:

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

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

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```
- 10 Copy the contents of the `product/patches` directory to local or shared storage.

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

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

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

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

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

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

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

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

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

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

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

13 Apply the patches to all or selected cluster nodes.

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

14 Shut down and restart the system.

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

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

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

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

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

18 Remount all VxFS file systems and Storage Checkpoints:

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


- 19 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```
- 20 If you are currently using backup and restore for the DBED repository. Perform a full backup of the DBED repository database after completing the 5.0 MP3 RP1 installation.
For more information see the [“Software limitations”](#) on page 19 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 12 or [“Storage Foundation for DB2 fixed issues”](#) on page 13 for incident 1425261.

Verifying software versions

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

```
# lspp -L VRTS\*
```

Removing 5.0 MP3 RP1

Roll back of the 5.0 MP3 RP1 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 RP1 to the release 5.0 MP3 version for Veritas Cluster Server.

- [Removing 5.0 MP3 RP1 from Veritas Cluster Server](#)
- [Removing 5.0 MP3 RP1 on Storage Foundation or Storage Foundation Cluster File System](#)
- [Removing 5.0 MP3 RP1 on Storage Foundation for Oracle RAC](#)

Removing 5.0 MP3 RP1 from Veritas Cluster Server

Use the following procedure to manually remove VCS 5.0 MP3 RP1 from your cluster.

To manually remove 5.0 MP3 RP1 from VCS

- 1 Verify that all of the VCS 5.0MP3RP1 patches are in the APPLIED state.
Create a text file called `filesets.to.reject` that contains the name and version of each fileset, one per line, exactly as shown below.

```
VRTSgab.rte      5.0.3.100
VRTSl1t.rte      5.0.3.100
VRTSvcs.rte      5.0.3.100
VRTSvcsag.rte    5.0.3.100
VRTSvxfen.rte    5.0.3.100
```

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

```
# nohdr='^Z$'
# while read pkg ver; do
    lsipp -l $pkg | egrep -v "$nohdr"
    nohdr='^ Fileset +Level State '
done < fileset.to.reject
```

3 Review the output and confirm that all of the updated filesets are in the APPLIED state.

Example output follows:

Fileset	Level	State	Description

Path: /usr/lib/objrepos VRTSgab.rte	5.0.3.100	APPLIED	Veritas Group Membership and Atomic Broadcast 5.0MP3RP1 by Symantec 12/15/2008-11:49:23 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505
Path: /etc/objrepos VRTSgab.rte	5.0.3.100	APPLIED	Veritas Group Membership and Atomic Broadcast 5.0MP3RP1 by Symantec 12/15/2008-11:49:23 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505

Path: /usr/lib/objrepos VRTSl1t.rte	5.0.3.100	APPLIED	Veritas Low Latency Transport 5.0MP3RP1 by Symantec 12/15/2008-11:49:11 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505
Path: /etc/objrepos VRTSl1t.rte	5.0.3.100	APPLIED	Veritas Low Latency Transport 5.0MP3RP1 by Symantec 12/15/2008-11:49:11 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505

Path: /usr/lib/objrepos VRTSvc.s.rte	5.0.3.100	APPLIED	Veritas Cluster Server 5.0MP3RP1 by Symantec 12/15/2008-11:48:10 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505
Path: /etc/objrepos VRTSvc.s.rte	5.0.3.100	APPLIED	Veritas Cluster Server 5.0MP3RP1 by Symantec 12/15/2008-11:48:10 Build Env:AIX ogma 3 5 0001D4CAD300 IBM,9115-505

Path: /usr/lib/objrepos VRTSvc.sag.rte	5.0.3.100	APPLIED	Veritas Cluster Server 5.0MP3RP1 Bundled Agents by Symantec

```
-----
Path: /usr/lib/objrepos
VRTSvxfen.rte          5.0.3.100  APPLIED  Veritas I/O Fencing 5.0MP3RP1
                        by Symantec
                        12/15/2008-11:49:40 Build
                        Env:AIX ogma 3 5 0001D4CAD300
                        IBM,9115-505
```

```
Path: /etc/objrepos
VRTSvxfen.rte          5.0.3.100  APPLIED  Veritas I/O Fencing 5.0MP3RP1
                        by Symantec
                        12/15/2008-11:49:40 Build
                        Env:AIX ogma 3 5 0001D4CAD300
                        IBM,9115-505
```

- 4 Any updates that are in COMMITTED state cannot be rejected (undone). You must remove each one and then re-install it.
- 5 List the service groups in your cluster and their status. On any node, type:
hagrps -state
- 6 Take the ClusterService service group offline if it is running. On any node, type:
hagrps -offline -force ClusterService -any
- 7 Make the VCS configuration writable. On any node, type:
haconf -makerw
- 8 Freeze all service groups except the ClusterService service group. On any node, type:

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

You can safely ignore the warning about the failure to freeze the ClusterService group.
- 9 Save the configuration (main.cf) file with the groups frozen. On any node, type:
haconf -dump -makero
- 10 Make a backup copy of the current main.cf and all types.cf configuration files. For example, on one node in the cluster, type:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
  /etc/VRTSvcs/conf/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
  /etc/VRTSvcs/conf/types.cf.save
```
- 11 Shut down VCS. On any node, type:
/opt/VRTSvcs/bin/hastop -all -force

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

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

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

The output resembles:

```
GAB Port Memberships
Port a gen 23dc0001 membership 01
Output for membership for port h does not appear.
```
 - On each node, run the command

```
# ps -ef | egrep "[ ](had|hashadow|CmdServer) "
```

Terminate any instances of had, hashadow, or CmdServer that still run after 60 seconds.
- 14 Unconfigure vxfen if the VCS cluster uses the fencing option. On each node, type:

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

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

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

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

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

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

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

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

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

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

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

- 20 Preview the patch removal selection and validity tests. On each node, type:

```
# installp -pr -gXv -f filesets.to.reject
```

 Confirm that the patches to be removed are exactly the same as those listed in the filesets.to.reject file that you created in [step 1](#).
- 21 Perform the patch removal. On each node, type:

```
# installp -r -gXv -f filesets.to.reject
```

 Review the summaries at the end of each run and confirm that all of the intended patches removed successfully.
- 22 Reboot all nodes in the cluster.
- 23 After VCS has started, perform the following steps:
 - Verify all resources have been probed. On any node, type:

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

```
# haconf -makerw
# hagr -list | sort -u +0b -1 | \
  while read grp sys ; do
    hagr -unfreeze $grp -persistent
  done
# haconf -dump -makero
```

You can safely ignore the warning about the failure to unfreeze the ClusterService group.
- 24 Bring the ClusterService service group online, if necessary. On any node, type:

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

 where system is the node name.

Removing 5.0 MP3 RP1 on Storage Foundation or Storage Foundation Cluster File System

You can use the following procedure to uninstall 5.0 MP3 RP1 on Storage Foundation or Storage Foundation Cluster File System (SFCFS).

To uninstall 5.0 MP3 RP1 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 Stop VCS along with all the resources. Then, stop the remaining resources manually:

```
# /etc/rc.d/rc2.d/vcs stop
```
- 4 Stop the VCS command server:

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

where `pid_of_CmdServer` is the process ID of CmdServer.
- 5 Uninstall VCS:

```
# cd /opt/VRTS/install  
# ./uninstallvcs [-usersh]
```
- 6 If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

```
# rm /etc/vxfenmode
```
- 7 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# mount | grep vxfs
```
- 8 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```
- 9 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:
 - 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 rvrg_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.

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

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

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

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

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

- 12 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
```

- 13 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, use the following commands:

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

```
# ./uninstallsf [-usersh]
```

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 `-usersh` 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 cluster nodes.

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 RP1 on Storage Foundation for Oracle RAC

You can use the following procedure to uninstall the 5.0 MP3 RP1 on Storage Foundation for Oracle RAC systems.

To uninstall the 5.0 MP3 RP1 on SF Oracle RAC

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

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

- 2 Stop all other VCS resources on each cluster node:

```
# hstop -local
```

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

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

Sample output:

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

- 4 Uninstall Storage Foundation for Oracle RAC.

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

```
# ./uninstallsfrac galaxy nebula
```

See the *Veritas Storage Foundation for Oracle RAC 5.0 MP3 Installation and Configuration Guide* for more information.

After uninstalling the packages, refer to the Storage Foundation for Oracle RAC 5.0 MP3 Installation and Configuration Guide to reinstall the 5.0 MP3 software.

Documentation addendum

The following is an addition to the *Veritas Cluster Server Bundled Agents Reference Guide for AIX*.

MemCPUALlocator agent

Use the MemCPUALlocator agent to allocate CPU and memory to an IBM AIX dedicated partition. Set this resource's attribute values to specify the amount of CPU and memory that you want to allocate to a service group on a DLPAR. Configure this resource as a leaf node in the service group dependency tree.

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

“[MemCPUALlocator agent notes](#)” on page 77


Platform

AIX

Dependencies

Set the MemCPUALlocator resource as a leaf node in a resource dependency tree. Select the amount of CPU and memory that you want the DLPAR to have before it comes online.

Figure 1-1 Sample service group for a MemCPUALlocator resource, where the MemCPUA resource represents the MemCPUALlocator resource



Agent functions

Online	The MemCPUAllocator agent dynamically allocates the required amount of memory and CPU to the DLPAR from the Hardware Management Console (HMC).
Offline	The agent deallocates the amount of memory and CPU it acquired during the online agent function. It then returns the resources back to the pool.
Monitor	<p>Checks that the online agent function succeeded. If it succeeded, then the monitor agent function reports the resource state as ONLINE. If it did not succeed, then the monitor agent function reports the resource state as OFFLINE.</p> <p>If the agent is not able to allocate the required resources during the online agent function, the subsequent monitor reports OFFLINE and the resource faults. Because the resource is a leaf node, VCS engine stops bringing other resources online and marks the group as FAULTED. The VCS engine then tries to bring the group online on some other DLPAR. This check ensures that the agent can dynamically allocate the resources that the service group requires for the DLPAR.</p>

Attributes

Table 1-24 Required attributes

Required attribute	Description
ManagedSystem	<p>The name of the managed system that contains the partition.</p> <p>Type-dimension: string-scalar</p> <p>Example: mymachine</p>
HMC	<p>Name of the HMC</p> <p>The list of HMCs that control the managed systems. The agent tries to connect to any HMC on this list in the order that they are specified.</p> <p>Type-dimension: string-vector</p> <p>Example: HMC = { myhmc1, myhmc2 }</p>

Table 1-25 Optional attributes

Optional attribute	Description
MemoryRequired	Amount of RAM (in MB) that you want to allocate. Type-dimension: string-scalar Default: 0 Example: 256
MemoryCritical	Specifies whether the memory allocation is critical. A value of 0 indicates that the online agent function should go ahead even when the required memory was not successfully allocated. Type-dimension: boolean-scalar Default: 0 Example: 1
CPURequired	The number of dedicated CPUs that you want to allocate. Type-dimension: string-scalar Example: 2
CPUCritical	Specifies whether the CPU allocation is critical. A value of 0 indicates that the online agent function should proceed even when the required CPU was not successfully allocated. Type-dimension: boolean-scalar Default: 0 Example: 1

Resource type definition

```
type MemCPUAllocator (
  static str ArgList[] = { ManagedSystem, HMC, MemoryRequired,
    MemoryCritical, CPUCritical, CPURequired }
  str ManagedSystem
  str HMC[]
  str MemoryRequired
  str CPURequired
  boolean CPUCritical = 0
  boolean MemoryCritical = 0
  temp boolean IsOnline = 0
)
```

MemCPUAllocator agent notes

The MemCPUAllocator agent has the following notes:

- See [“Configuring password free SSH communication between VCS nodes and HMC”](#) on page 77.
- See [“Dynamic resource allocation scenarios”](#) on page 78.
- See [“Configuring MemCPUAllocator”](#) on page 82.

Configuring password free SSH communication between VCS nodes and HMC

To use remote command operations on the HMC, you must have SSH installed on the DLPAR nodes in the VCS cluster. You must configure the HMC to allow password free SSH access from these partitions. Refer to the appropriate IBM AIX documentation for information.

To verify that you have password free SSH access

- ◆ From each DLPAR in the cluster, execute the following command to test if the password free access works.

```
Eagle> ssh -l hscroot hmc2.veritas.com
Last login:Thur Jun 16 22:46:51 2005 from 10.182.9.34
hscroot@hmc2:~>
```

Once each node can connect to the HMC using SSH without a password, you can start to use the MemCPUAllocator agent.

Dynamic resource allocation scenarios

This section describes different examples of the resource allocation scenarios that the MemCPUAllocator agent can handle. For ease of explanation, consider only the memory resource in these examples. CPU resource implementation is similar.

Consider two DLPARs named Eagle and Vulture. These DLPARs are configured with the following minimum and maximum values memory values.

Table 1-26 The minimum and maximum memory for the DLPARs Eagle and Vulture

DLPAR	Minimum	Maximum
Eagle	512 MB	2 GB
Vulture	512 MB	2 GB

Two service groups SG1 and SG2 have the following resource requirements.


Table 1-27 The memory that is required for service group SG1 and SG2

Service group	Required memory
SG1	512 MB
SG2	512 MB

Scenario 1: A DLPAR node has minimum resources

Assume that the DLPARs start with the minimum values for memory. When SG1 is brought online on Eagle, the online agent function for the agent attempts to allocate 512 MB to Eagle from the free pool. The agent retains the minimum resources for the DLPAR's overhead operations and allocates resources for the service group in addition to the existing memory. For SG1 to come online the agent allocates an additional 512 MB to Eagle. After this allocation the total current memory for eagle is 1 GB. If SG1 goes offline, the agent deallocates the 512 MB that it allocated when the service group came online. This deallocation brings back the current memory of Eagle to 512 MB.


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



Scenario 2: Bringing another service group online

In this scenario, the Eagle DLPAR starts with 512 MB, and has SG1 online on it. It uses a total of 1 GB of memory. If SG2 is brought up on Eagle, the agent allocates an additional 512 MB of memory to Eagle. This reallocation brings the total memory to 1.5 GB.


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



Scenario 3: DLPAR has required resources


Instead of starting with 512 MB, Eagle starts with 1 GB of initial memory. Eagle has 512 MB more than its minimum amount. If SG1 is brought online on Eagle, the agent determines that Eagle has an extra 512 MB more than its minimum. No service groups use this extra 512 MB. The agent does not allocate any additional memory to Eagle. SG1 is brought online on Eagle and the current memory for Eagle stays 1 GB.

Figure 1-4 DLPAR Eagle starting with 1 GB of initial memory



Scenario 4: Cannot allocate required resources

Figure 1-5 Exceeding the maximum amount of memory on a DLPAR



Consider the stage in Scenario 2, where SG1 and SG2 are both online on Eagle, which brings its current memory to 1.5 GB. An additional service group SG3 enters the picture and requires 1 GB memory. SG3 tries to come up on Eagle. The agent determines that allocating 1 GB more memory to Eagle exceeds its maximum limit of 2 GB. The agent therefore does not allocate the memory and the online agent function fails, which leads to a resource fault. This resource fault makes the VCS engine stop the online of SG3 on Eagle and try it on Vulture. If Vulture starts with 512 MB and the agent allocates an additional 1 GB to Vulture, its current memory is 1.5 GB. SG3 can fail over and come online on Vulture.

Scenario 5: Service group failover

As in Scenario 2, SG1 and SG2 are both online on Eagle, which brings its current memory to 1.5 GB. Vulture has a current memory configuration of 512 MB. If you switch the service groups from Eagle to Vulture:

- The MemCPUAllocator agent's offline agent function deallocates 1 GB from Eagle (512 MB for SG1 and 512 MB for SG2).
- The VCS engine migrates SG1 and SG2 to Vulture and the agent's online agent function allocates 1 GB to Vulture. This allocation brings Vulture's memory to 1.5 GB.

Configuring MemCPUAllocator

Before you can use the MemCPUAllocator agent, you need to set up SSH access between the HMC and the DLPAR nodes. You must also make sure to configure the MemCPUAllocator resource as a leaf node in the service group's dependency tree in the main.cf file.

See [Figure 1-1, “Sample service group for a MemCPUAllocator resource, where the MemCPUA resource represents the MemCPUAllocator resource,”](#) on page 74.

Provide values to the MemCPUAllocator resource to specify the resource requirements for that service group. For example, if a service group needs 512 MB memory and two CPUs to start with, the MemCPUAllocator resource definition resembles:

```
MemCPUAllocator mymem (  
    ManagedSystem @eagle = eagle-server  
    ManagedSystem @vulture = vulture-server  
    HMC = { testhmc }  
    RequiredMemory = 512  
    RequiredCPU = 2  
    MemoryCritical = 1  
    CPUCritical = 1  
)
```