

# Veritas Storage Foundation<sup>™</sup> for Oracle RAC Release Notes

AIX

5.0

Maintenance Pack 1



# Veritas Storage Foundation for Oracle RAC Release Notes

Copyright © 2007 Symantec Corporation. All rights reserved.

## Storage Foundation for Oracle RAC

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](http://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. Third-party copyrights associated with this product are listed in these release notes.

AIX is a registered trademark of IBM Corporation.  
Oracle is a registered trademark of Oracle Corporation.

## Licensing and registration

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

## Technical support

For technical assistance, visit <http://support.veritas.com> 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.



# Veritas Storage Foundation for Oracle RAC Release Notes

This document provides important information regarding Veritas™ Storage Foundation for Oracle RAC (SF Oracle RAC) for AIX. Review this entire document before installing this product.

See the following Technical Support TechNote for the latest information on updates, patches, and software issues regarding this release:

<http://support.veritas.com/docs/282024>

Veritas Storage Foundation for Oracle RAC is an integrated set of software products. It enables administrators of Oracle Real Application Clusters (RAC) to operate a database in an environment of cluster systems running Veritas Cluster Server (VCS) and the cluster features of Veritas Volume Manager and Veritas File System, also known as CVM and CFS, respectively.

Acronyms used for Storage Foundation for Oracle RAC components include the following:

<b>CVM</b>	Veritas Volume Manager cluster functionality
<b>CFS</b>	Veritas File System cluster functionality
<b>VCS</b>	Veritas Cluster Server
<b>VEA</b>	Veritas Enterprise Administrator graphical user interface
<b>VxFS</b>	Veritas File System
<b>VxVM</b>	Veritas Volume Manager

SF Oracle RAC components are described in the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*. Refer to component-product documentation and release notes for additional information to assist you during installation and configuration of SF Oracle RAC.

## Changes in this release

The 5.0 Maintenance Pack 1 for SF Oracle RAC introduces the following changes:

- VCS FEN error messages are now VXFEN error messages
- For the latest information on support for Oracle database versions, see the Veritas Technical Support TechNote:  
<http://support.veritas.com/docs/280186>

## Features introduced in SF 5.0 for Oracle RAC

We offer the following new features in SF 5.0 for Oracle RAC. See the *Installation and Configuration Guide* for details.

- Support for up to 8 nodes<sup>1</sup>
- Fencing with DMP disks
- Support for Oracle 10g Release 2

## System requirements

Make sure that you have the correct equipment to install SF Oracle RAC.

Item	Description
SF Oracle RAC systems	Two to eight systems with two or more CPUs at 2GHz or higher.
RAM	Each SF Oracle RAC system requires 2 GB or more of physical memory.
Network links	Two or more private links and one public link. Links must be 100BaseT or Gigabit Ethernet directly linking each node to the other node to form a private network that handles direct inter-system communication.  Symantec recommends switches for the private network links. These links must be of the same type; you cannot mix 100BaseT and Gigabit.
DVD drive	One drive that is accessible to all nodes in the cluster.
Fibre channel or SCSI host bus adapters	SF Oracle RAC requires at least one built-in SCSI adapter per system to access the operating system disks, and at least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.

1.A typographical error in previous Release Notes indicated 9 nodes.

Item	Description
Disks	<p>Typical SF Oracle RAC configurations require that shared disks support applications that migrate between systems in the cluster.</p> <p>The SF Oracle RAC I/O fencing feature requires that all disks used as data disks or as coordinator disks must support SCSI-3 Persistent Reservations (PR).</p>
Disk space	<p>SF Oracle RAC space requirement:</p> <ul style="list-style-type: none"><li>■ total: 2.6 G</li><li>■ /opt: 1.4 G</li><li>■ /usr: 200 KB</li><li>■ /tmp: 512 MB</li><li>■ /var: 32 MB</li><li>■ /var/tmp: 700 MB</li></ul>
Swap space	Two times the main memory.

## Supported hardware

Review the current compatibility list to confirm compatibility of your hardware:  
<http://support.veritas.com/docs/283161>

## Supported software

Supported software includes:

- Operating system and patches
- Veritas software
- Oracle software

### Operating system and patches

Within a cluster, all nodes must use the same operating system version and patch level. Run SF Oracle RAC on the following operating systems at the required patch levels.

- AIX 5.2 TL8
- AIX 5.3 TL5 with SP 1

### Veritas software

Only versions of VCS, VxVM, and VxFS provided on the software disc are supported. Remove other versions before you install the software from the Veritas Storage Foundation for Oracle RAC software disc.

## Oracle software

SF Oracle RAC supports the following Oracle software:

- Oracle9i Release 2
- Oracle 10g Release 1
- Oracle 10g Release 2

To verify the latest information on support for Oracle database versions, see the Technical Support TechNote:

<http://support.veritas.com/docs/280186>

## Oracle software patches

- Oracle 10g Release 2, Patchset 1: requires Oracle Patch 4637591
- Oracle 10g Release 2: requires Oracle Patches 4435949, 4637591, 5082958

# Software limitations

The following limitations apply to this release.

## Cluster File System (CFS)

The following are limitations of Cluster File System:

- CFS supports only disk layout version 6 and 7.
- No support for HSM/DMAPI.
- VxFS does not distribute its own versions of file utilities such as `cpio` or `mv`.

## Volume Manager in a shared disk environment (CVM)

Volume Manager cluster functionality does not support software RAID 5.

## Space-optimized volume snapshots

For minimal performance impact Symantec recommends that the Space Optimized Snapshots (SOS) be created only of the Data Volumes. A mirror break of snapshot should be created of the Oracle Log volume. Creating a mirror breakoff snapshot of the Oracle Log volume does not have much space overhead since the log volumes are typically small in size.

## Volume location for SRVM, ocrvol, and votevol

The Oracle RAC wizard does not support the following condition:

- For Oracle9i, the SRVM volume cannot exist in the same shared disk group as that of Oracle datafiles.
- For Oracle 10g, the ocrvol volume and votevol volume cannot exist in the same shared disk group as that of Oracle datafiles.

However, you can allow for this condition when you manually configure Oracle service groups.

## Oracle Disk Manager (ODM)

Oracle Disk Manager uses the Quick I/O driver for asynchronous I/O. Do not turn off the Quick I/O mount option, which is the default.

## Storage Checkpoints and Database FlashSnap

The following are limitations of Storage Checkpoints and Database FlashSnap:

- You cannot create a clone database using a mounted Storage Checkpoint. [See Etrack Incident 32726]
- If you create an Oracle instance using the `spfile` option, you must run the `dbed_update` command before you can successfully perform any Storage Checkpoint or Database FlashSnap functions.
- Storage Checkpoints require file system layout version 6 or 7. Use the `vxupgrade(1M)` command to check the current layout version and to change the layout version if necessary. When upgrading a CFS file system, issue the command from the primary node.  
**Note:** After you upgrade a system to layout version 6 or 7, the file system is no longer compatible with older VxFS file systems.
- When cloning a database using Database FlashSnap, the Oracle database must have at least one mandatory archive destination. For more information about Oracle parameters for archiving redo logs, see your Oracle documentation. [See Etrack Incident 270905]
- For an Oracle RAC database, when using the `dbed_vmsnap`, `dbed_vmclonedb`, and `dbed_vmchecksnap` commands, only online snapshots are supported. [See Etrack Incident 605330]
- After running `dbed_vmsnap -o reverse_resync_commit`, your primary database is started using a `pfile`. If your original primary database used an `spfile`, you need to shut down the database and restart it using the `spfile`. Then, run `dbed_update` to update the repository.
- The Storage Checkpoint and Database FlashSnap features of SF Oracle RAC do not support the GUI of the Veritas Storage Foundation for Oracle product.

- The Database FlashSnap feature does not support RAID-5 volumes. [See Etrack Incident 34570]

## RDAC driver and FASTT array for coordinator disks

For multipathing connected storage when using the AIX RDAC driver for FASTT arrays, the following points must be considered:

- Since it is an active/passive array, only the current active path is exposed to clients.
- The I/O fencing driver, vxfen, can use only a single active path and has no fore-knowledge of the passive paths to the coordinator disks on an array.

Therefore, if the single active path fails, all nodes in the cluster lose access to the coordinator disks. The loss of the path to the coordinator disks can potentially go unnoticed until a reboot, split brain, or any other reason that leads to a cluster membership change occurs. In any of these conditions, the cluster cannot form, and all nodes panic to prevent data corruption. No data loss occurs.

## LLT low priority heartbeat links not recommended

Symantec does not recommend using the public network as LLT low priority links for SF Oracle RAC. Low priority links reduce network overhead by reducing the frequency of heartbeats, but they are not suited for the typical volume of SF Oracle RAC data.

## Installation notes

The following information includes guidelines, tips, and other considerations for installing the product.

## Storage Foundation for Oracle RAC packages

Storage Foundation for Oracle RAC is a licensed product. The *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* describes how to obtain a license that applies for Veritas Storage Foundation for Oracle RAC, Veritas Volume Manager, Veritas File System, and Veritas Cluster Server.

### Software installed for Storage Foundation for Oracle RAC

The software components installed by the Veritas Storage Foundation 5.0 for Oracle RAC installation utility, `installsfprac`, include:

- Veritas Storage Foundation for Oracle RAC

- Veritas Cluster Server (VCS)
- Veritas Volume Manager (VxVM)
- Veritas Group Lock Manager
- Veritas File System (VxFS)
- Veritas extension for Oracle Disk Manager (ODM)
- Veritas Cluster Server agent for Oracle

## Installing SF Oracle RAC version 5.0 MP1 for the first time

If you are installing Storage Foundation 5.0 MP1 for Oracle RAC software for the first time and not upgrading an existing system, for pre-installation instructions:

See the *Veritas Storage Foundation 5.0 for Oracle RAC Installation and Configuration Guide*, “Preparing to install SF Oracle RAC” on page 41.

Installing SF 5.0 MP1 for Oracle RAC involves the following tasks:

- Installing SF Oracle RAC version 5.0 using `installsfrac` script  
Note that to upgrade to version 5.0 MP1, you must have SF Oracle RAC version 5.0 installed.
- Installing SF 5.0 MP1 using `installmp` script
- Configuring SF Oracle RAC using `installsfrac` script

### To install the SF 5.0 for Oracle RAC software

- 1 Log in as superuser on one of the systems for installation.
- 2 Install the Veritas Storage Foundation 5.0 for Oracle RAC GA software from the release CD. Run the product installation script instead of the generic installer script.

See the *Storage Foundation for Oracle RAC Installation and Configuration Guide*.

Enter the following command from the top-level directory of the mounted CD:

```
# ./installsfrac [-rsh] -installonly galaxy nebula
```

The `-rsh` option is required if you are using the remote shell (RSH) rather than the secure shell (SSH) to install the software simultaneously on several systems.

- 3 Insert the disc containing the SF 5.0 MP1 for Oracle RAC software into the disc drive, and mount the disc on a suitable mount point.
- 4 To install SF 5.0 MP1 for Oracle RAC, you must invoke the `installmp` script from one of your cluster nodes using the option that corresponds to your configuration.

For example:

- To install on the local system only:  
# `./installmp`
- To install on more than one system using secure shell (SSH) utilities:  
# `./installmp galaxy nebula ...`
- To install on more than one system using remote shell (RSH) utilities:  
# `./installmp -rsh galaxy nebula`

A restart of the system will be required after the patches are installed.

- If CRS is not configured under VCS control, log in as superuser on *galaxy* and *nebula* in the cluster and enter:  
# `/etc/init.crs start`
  - If CRS is configured under VCS control, log in as superuser and enter the following command for each system in the cluster:  
# `/opt/VRTSvcs/bin/hares -online cssd_resource -sys galaxy`  
# `/opt/VRTSvcs/bin/hares -online cssd_resource -sys nebula`
- 5 After the initial system checks and the requirement checks are complete, press Return to start upgrading the packages.
  - 6 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.
  - 7 Restart all systems in the cluster.  
# `/usr/sbin/shutdown -r`
  - 8 Change to `/opt/VRTS/install` directory.
  - 9 Run the `installsfrac` script, this time specifying the `-configure` option to configure the software, for example:  
# `./installsfrac [-rsh] -configure galaxy nebula`  
See the *Veritas Storage Foundation 5.0 for Oracle RAC Installation and Configuration Guide*, “Installing and Configuring SF Oracle RAC Software” on page 83.
  - 10 Install and configure the correct version of Oracle.  
See Chapter 5, “Preparing to install Oracle9i RAC” in the *Veritas Storage Foundation 5.0 for Oracle RAC Installation and Configuration Guide*.  
See Chapter 10, “Preparing to install Oracle 10g RAC” in the *Veritas Storage Foundation 5.0 for Oracle RAC Installation and Configuration Guide*.
  - 11 If you want to use Checkpoints, FlashSnap, or Storage Mapping, configure the repository database.  
See Chapter 18, “Configuring the repository database for Oracle” in the *Veritas Storage Foundation 5.0 for Oracle RAC Installation and Configuration Guide*.

## Upgrading from SF Oracle RAC 5.0 to 5.0 MP1

You can upgrade to 5.0 MP1 only if SF Oracle RAC 5.0 is installed and configured. For complete procedures:

See the *Storage Foundation for Oracle RAC Installation and Configuration Guide*.

---

**Warning:** After running the `installmp` script, the filesets are in an applied state. Once the filesets are committed, you cannot roll back the 5.0 Maintenance Pack 1 product to the previously installed release.

---

### Overview of upgrade tasks

You must perform the following tasks:

- [Upgrading your operating system](#)
- [Stopping Oracle and the cluster resources](#)
- [Upgrading to SF Oracle RAC 5.0 MP1 using the `installmp` installer](#)
- [Performing post-upgrade tasks for SF Oracle RAC 5.0 MP1](#)

---

**Note:** When invoking the installer, use the `-rsh` option if you have not configured SSH communications between systems. SSH is the default.

---

---

**Note:** The upgrade procedure necessitates a restart after installing the SF 5.0 Oracle RAC filesets.

---

### Upgrading your operating system

If you are using an unsupported version of AIX, upgrade now:

- For AIX 5.3, SF Oracle RAC for 5.0 MP1 requires that the operating system be at 5.3 TL5 SP1.
- For AIX 5.2, SF Oracle RAC for 5.0 MP1 requires that the operating system be at 5.2 TL8.

### Stopping Oracle and the cluster resources

Before you install Maintenance Pack 1 using the `installmp` script, you must stop Oracle and the cluster resources.

### To stop Oracle and the cluster resources[

- 1 Log in as superuser to one of the nodes, galaxy for example, in the cluster.
- 2 Create a backup of the existing cluster configuration. Back up the main.cf and types.cf on all cluster nodes:
 

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/config/types.cf.save
```
- 3 For Oracle9i, stop all Oracle resources including the database on all nodes.
  - Stop gsd. On each node, log in as oracle user and enter:
 

```
$ $ORACLE_HOME/bin/gsdctl stop
```
  - Offline all VCS service groups that contain resources for managing Oracle listener process or the database instances. As root user, enter:
 

```
# hagrps -offline group_name -any
```
- 4 For Oracle 10g, stop all resources.
  - If the database instances are under VCS control, take the corresponding VCS service groups offline. As superuser, enter:
 

```
# hagrps -offline group_name -any
```
  - If the database instances are under CRS control, then run the following on one node:
 

```
$ srvctl stop database -d database_name
```
  - If the resources are under CRS control, stop all resources configured under CRS control. As oracle user, enter:
 

```
$ srvctl stop nodeapps -n sys
```

 Repeat the above command for each node in the cluster. Note that if the Oracle listener or database instances are managed by CRS, they are taken offline by the “srvctl” command shown above.
- 5 For Oracle 10g, stop CRS on each node in the cluster. Symantec highly recommends putting CRS under VCS control.
  - If CRS is controlled by VCS, log in as superuser on any system in the cluster and enter:
 

```
# hares -offline cssd-resource -sys galaxy
# hares -offline cssd-resource -sys nebula
```
  - If CRS is not controlled by VCS, log in as superuser on each system in the cluster and enter:
 

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

 This may take a minute or two, because the CRS process takes some time to stop in the background after the command is given.
- 6 Take offline all other VCS groups that depend on VxFS file systems or VxVM disk groups, whether local or CFS.

- 7 Stop all applications that use VxFS or VxVM but that are not under VCS control.
- 8 Make sure that no processes are running which make use of mounted shared file system or shared volumes.
- 9 Unmount all the VxFS file system which is not under VCS control.
 

```
# mount |grep vxfs
# fuser -c /mount_point
# umount /mount_point
```
- 10 Deport any CVM disk groups that are not under VCS control.
- 11 Set the VCS resource attribute, AutoStart, to 0 for all VCS resources that manage the Oracle 10g CRS or Oracle database instances. From any system in the cluster, enter:
 

```
# haconf -makerw
# hares -modify oracle_res AutoStart 0
# haconf -dump -makero
```

 For Oracle 10g:
 

```
# haconf -makerw
# hares -modify cssd_resource AutoStart 0
# haconf -dump -makero
```
- 12 With Storage Foundation for Oracle RAC 5.0 installed and configured, run the following command on all nodes in the SF Oracle RAC cluster:
 

```
# hastop -all
```
- 13 For Oracle 10g, comment out the CRS-specific lines in the /etc/inittab file to prevent the system from starting the CRS daemons during system startup. On each system, prefix the CRS lines with a colon (":") to comment them out. For example:
 

```
:h1:2:respawn:/etc/init.evmd run >/dev/null 2>&1 </dev/null
:h2:2:respawn:/etc/init.cssd fatal >/dev/null 2>&1 </dev/null
:h3:2:respawn:/etc/init.crsd run >/dev/null 2>&1 </dev/null
```
- 14 On each system, invoke "slibclean" to unload all unused modules:
 

```
# /usr/sbin/slibclean
```
- 15 For an optional port verification step, make sure only ports a, b, d, and o are open:
 

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 6b5901 membership 01
Port b gen 6b5904 membership 01
Port d gen 6b5907 membership 01
Port o gen 6b5905 membership 01
```
- 16 For an optional disk group verification step, make sure no disk groups are imported:

```
# vxvg list
NAME          STATE          ID
#
^^^ NOTE: no diskgroups imported
```

## Upgrading to SF Oracle RAC 5.0 MP1 using the installmp installer

Start the installer and add the patches using the following procedure.

### To upgrade to SF Oracle RAC 5.0 MP1

- 1 Insert the disc containing the Veritas SF Oracle RAC 5.0 MP1 software in a disc drive connected to one of the nodes for installation.
- 2 Log in as superuser on one of the systems for installation.
- 3 Mount the software (5.0 MP1) disc.
 

```
# mkdir /cdrom
# mount -V cdrfs -o ro /dev/cd0 /cdrom
```
- 4 Browse to the directory containing installer program.
 

```
# cd /cdrom
```
- 5 Start the SF Oracle RAC installer utility.
 

Invoke installmp from one of your cluster nodes using the option that corresponds to your configuration:

  - To install using ssh (default):
 

```
# ./installmp galaxy nebula
```
  - To install using rsh:
 

```
# ./installmp -rsh galaxy nebula
```
- 6 Review the output as the installer checks that the systems are ready for installation. At the conclusion of these initial checks, and before installing any software, the installer does the following:
  - Creates a log file (installmp-xxxxxx) on each system in the directory /var/tmp.
  - Specifies the utilities it uses to communicate with the remote systems; typically this is ssh -x and scp.
  - Lists the filesets that it will install on each system
 

Storage Foundation Cluster File System and Veritas Enterprise Administrator processes must be stopped before adding patches. When you are prompted:

```
installmp is now ready to install MP1.
Required 5.0 processes that are currently running will be
stopped.
Are you sure you want to install MP1? [y,n,q] (y)
```

Press **Enter** to continue.

- 7 Note the location of the summary and log files for future reference.
- 8 On all nodes, restore the configuration files from the backups created in an earlier procedure.  
See [“To stop Oracle and the cluster resources\[”](#) on page 14.
  - Restore the original main.cf file:
 

```
# cp /etc/VRTSvcs/conf/config/main.cf.save \
/etc/VRTSvcs/conf/config/main.cf
```
  - Merge any changes you made to the original types.cf file and backed up in a previous procedure at:
 

```
/etc/VRTSvcs/conf/config/types.cf.save
```

 with the new types.cf file at:
 

```
/etc/VRTSvcs/conf/types.cf
```
- 9 Restart each system in the cluster when prompted.
 

```
# /usr/sbin/shutdown -r
```
- 10 Make sure that all eight GAB ports are up after the restart.
- 11 Make sure that all the resources configured under VCS except the CSSD and Oracle resource come online.

## Performing post-upgrade tasks for SF Oracle RAC 5.0 MP1

To create the new skgxn archive library before relinking Oracle 10g

- 1 If you are running Oracle 10g, perform the following setup steps on each system in the cluster before bringing up CRS:
  - Unload any unused modules:
 

```
# /usr/sbin/slibclean
```
  - Create the 32-bit SKGXM archive library under /opt/VRTSvcs/rac/lib:
 

```
# cd /opt/VRTSvcs/rac/lib
# rm libskgxm2.a
# cp libvcsmm.so shr_skgxm2.o
# ar -X32 -qv libskgxm2.a shr_skgxm2.o
# rm shr_skgxm2.o
```
  - Create the 64-bit SKGXM archive library under /opt/VRTSvcs/rac/lib64
 

```
# cd /opt/VRTSvcs/rac/lib64
# rm libskgxm2.a
# cp libvcsmm.so shr_skgxm2.o
# ar -X64 -qv libskgxm2.a shr_skgxm2.o
# rm shr_skgxm2.o
```
- 2 If you are running Oracle 10g R1, copy these SF Oracle RAC SKGXM libraries to \$CRS\_HOME
 

```
# cp /opt/VRTSvcs/rac/lib/libskgxm2.a \
```

```

$CRS_HOME/lib32/libskgxn2.a
# cp /opt/VRTSvcs/rac/lib/libskgxn2.a \
$CRS_HOME/lib32/libskgxn2.a
# cp /opt/VRTSvcs/rac/lib64/libskgxn2.a \
$CRS_HOME/lib64/libskgxn2.a
# cp /opt/VRTSvcs/rac/lib64/libskgxn2.a \
$CRS_HOME/lib64/libskgxn2.a

```

Run the commands on *each* system in the cluster.

- 3 If you're running Oracle 10g R2, perform the following commands on each system in the cluster.

- Remove the following files from under /opt/ORCLcluster:

```

# rm /opt/ORCLcluster/lib/libskgxn2.a
# rm /opt/ORCLcluster/lib/libskgxn2.so
# rm /opt/ORCLcluster/lib32/libskgxn2.a
# rm /opt/ORCLcluster/lib32/libskgxn2.so

```

- Copy these SF Oracle RAC SKGXM libraries to /opt/ORCLcluster.

```

# cp /opt/VRTSvcs/rac/lib/libskgxn2.a \
/opt/ORCLcluster/lib32/libskgxn2.a
# cp /opt/VRTSvcs/rac/lib/libvcsmm.so \
/opt/ORCLcluster/lib32/libskgxn2.so
# cp /opt/VRTSvcs/rac/lib64/libskgxn2.a \
/opt/ORCLcluster/lib64/libskgxn2.a
# cp /opt/VRTSvcs/rac/lib64/libvcsmm.so \
/opt/ORCLcluster/lib64/libskgxn2.so

```

Run the commands on *each* system in the cluster.

- Copy the SF Oracle RAC SKGXP library to /opt/ORCLcluster:

```

# cd /opt/ORCLcluster/lib32
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_32.so .
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_32.a .

```

```

# cd /opt/ORCLcluster/lib64
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.so .
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.a .

```

- Copy the clsinfo script to /opt/ORCLcluster.

```

# cd /opt/ORCLcluster/bin
# cp /opt/VRTSvcs/rac/bin/scripts/clsinfo .

```

### To relink Oracle 10g R1 or R2 using the installer

- 1 Invoke installsfrac.

```

# cd /opt/VRTS/install
# ./installsfrac -configure

```

- 2 Enter the system names when prompted:

Enter the system names separated by spaces on which to configure  
SFRAC: galaxy nebula

- 3 Navigate to the “Install or Relink Oracle” menu.

- Select the appropriate Oracle 10g version:
    - 1) Oracle 9iR2
    - 2) Oracle 10gR1
    - 3) Oracle 10gR2
  - Select “Relink Oracle” from the menu:
    - 1) Install Oracle Clusterware (CRS)
    - 2) Install Oracle RDBMS server
    - 3) Relink Oracle
    - b) [Go to previous menu]
- 4 In the “Install or Relink Oracle” menu, enter the required information. For example:
- ```
Enter Oracle UNIX user name: [b] (oracle) oracle
Enter Oracle UNIX group name: [b] (oinstall) oinstall
Enter Oracle base directory: [b] /app/oracle
Enter absolute path of CRS home directory: [b] /app/crshome
Enter absolute path of Database Home directory: [b] /app/oracle/orahome
Enter Oracle Patch Level: [b] (0.1) 0.1
```
- 5 Confirm your responses in the verification screen. The installer copies the SF 5.0 Oracle RAC libraries to \$ORACLE\_HOME/lib.
- ```
Oracle environment information verification

Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Clusterware (CRS) Home: /app/crshome
Oracle Release: 10.2
Oracle Patch Level: 0.1
Oracle Base: /app/oracle
Oracle Home: /app/oracle/orahome

Is this information correct? [y,n,q] (y)
```
- 6 As root user, remove the prefix “:” to uncomment the CRS-specific lines in the /etc/inittab file.
- 7 Start CRS on each node in the cluster.
- If CRS is not controlled by VCS, log in as superuser on each system in the cluster and enter:
 

```
# /etc/init.crs start
```
  - If CRS is controlled by VCS, log in as superuser on any system in the cluster and enter:
 

```
# hares -online cssd-resource -sys galaxy
# hares -online cssd-resource -sys nebula
```
- 8 Start all resources.
- If the Oracle database instances are under VCS control, bring the corresponding VCS service groups online. As superuser, enter:

```
# hagrps -online oracle_group -sys galaxy
# hagrps -online oracle_group -sys nebula
```

- If the database instances are under CRS control, then run the following on one node:

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

- If the resources are under CRS control, start all resources configured under CRS control. As oracle user, enter:

```
$ srvctl start nodeapps -n sys
```

Repeat the above command for each node in the cluster. Note that if the Oracle listener or database instances are managed by CRS, they are taken offline by the “srvctl” command shown above.

- 9 Verify that the Oracle instances are up and running.
- 10 Reset the VCS resource attribute AutoStart to 1 for all resources that manage Oracle CRS and database instances. As root user on any system in the cluster, enter:

```
# haconf -makerw
# hares -modify cssd_res AutoStart 1
# hares -modify oracle_res AutoStart 1
# haconf -dump -makero
```

### To relink Oracle9i

- 1 Invoke installsfrac.
 

```
# cd /opt/VRTS/install
# ./installsfrac -configure
```
- 2 Enter the system names when prompted:
 

```
Enter the system names separated by spaces on which to configure
SFRAC: galaxy nebula
```
- 3 Navigate to the “Install or Relink Oracle” menu.
  - Select the Oracle 9i version:
    - 1) Oracle 9iR2
    - 2) Oracle 10gR1
    - 3) Oracle 10gR2
  - Select “Relink Oracle” from the menu:
    - 1) Install Oracle RDBMS server
    - 2) Relink Oracle
    - b) [Go to previous menu]
- 4 In the “Install or Relink Oracle” menu, enter the required information. For example:
 

```
Enter Oracle UNIX user name: [b] (oracle) oracle
Enter Oracle UNIX group name: [b] (oinstall) oinstall
Enter Oracle base directory: [b] /app/oracle
Enter absolute path of Database Home directory: [b] /app/oracle/
orahome
```

```
Enter Oracle Patch Level: [b] (0.1) 0.7
```

- 5 Confirm your responses in the verification screen. The installer copies the SF 5.0 Oracle RAC libraries to /opt/ORCLcluster, where it expects libskgxn.

```
Oracle environment information verification
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Release: 9.2
Oracle Patch Level: 0.7
Oracle Base: /app/oracle
Oracle Home: /app/oracle/orahome
Is this information correct? [y,n,q] (y)
```

- 6 As Oracle user, enter the command:

```
$ $ORACLE_HOME/bin/gsdctl start
```

- 7 Bring online the Oracle resources configured under VCS.

```
# hares -online listener_res -sys galaxy
# hares -online listener_res -sys nebula
# hagrps -online group_name -sys galaxy
# hagrps -online group_name -sys nebula
```

- 8 Verify that the Oracle instances are up and running.

- 9 Reset the VCS resource attribute AutoStart to 1 for all resources that manage Oracle and database instances. As root user on any system in the cluster, enter:

```
# haconf -makerw
# hares -modify oracle_res AutoStart 1
# haconf -dump -makero
```

### To commit the updated filesets

```
Use: installp -C filset_name
```

The updated filesets can still be removed manually after running `installmp`, and before committing. After committing, the updated filesets can not be rolled back to the previous release.

## Upgrading the repository database

After you have completed the SF Oracle RAC upgrade, if you plan to use the commands for Checkpoints, FlashSnap, or Storage Mapping, you must:

- Change permissions
- Upgrade the repository database

---

**Note:** Do not recursively change permissions, groups, or owners.

---

### To change permissions

- 1 Change permissions for the following directory:  
For Veritas Storage Foundation for Oracle:  

```
# chmod 550 /opt/VRTSdbed
```
- 2 Reset owner and group settings to the appropriate owner and group for the database administrators on your system.  
For example, to change owner to the user `oracle` and the group `dba`, run the following command:

```
# chown oracle:dba /opt/VRTSdbed
```

### To upgrade the repository in a cluster environment

- 1 Save the repository disk group name, mount point, and other attributes.
- 2 Unconfigure the SFUA repository from the VCS configuration:  

```
# /opt/VRTSdbcom/bin/sfua_db_config -o unconfig_cluster
```

  
This command:
  - removes the repository information from the VCS configuration
  - unmounts the repository
  - deports the disk group
  - takes the volumes offline
- 3 Reconfigure the repository:  

```
# vxdg import sfua_rep  
# vxvol -g dbed_rep startall
```
- 4 Mount the repository file system manually.  

```
# mount -V vxfs /dev/vx/dsk/sfua_rep/dbed_rep /sfua_rep
```
- 5 Run the repository upgrade command again with no options:  

```
# /opt/VRTSdbcom/bin/sfua_db_config
```

  
The following is required to configure the repository for SF Oracle RAC:
  - A mount point of already mounted Veritas Volume on a shared storage, with Veritas File system
  - A public NIC used by each system in the cluster
  - A Virtual IP address and netmask

## Minimum downtime upgrade from SF Oracle RAC 5.0 to 5.0MP1

This procedure is valid only for upgrading SF Oracle RAC 5.0 to 5.0 MP1. For upgrades from previous versions of SF Oracle RAC:

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

## Upgrading clusters with Oracle 10g

---

**Note:** There will be some downtime involved. Review the entire procedure and carefully plan your downtime before proceeding with any steps.

---

Assumptions:

- SF Oracle RAC 5.0 is installed
- Oracle 10g is installed
- Oracle binaries are installed on separate nodes for each node in the cluster

The example for this procedure assumes a four-node SF Oracle RAC 5.0 cluster with nodes *galaxy*, *nebula*, *jupiter* and *mercury*.

Upgrade the half of the nodes in the cluster at a time. The example nodes are *galaxy* and *nebula*.

### To prepare the first half of the Oracle 10g cluster for minimum downtime upgrade

On the first half of the cluster nodes, back up the configuration files:

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

### To upgrade the first half of the Oracle 10g cluster

- 1 Login to *galaxy* as root user, set Oracle group and cssd resource AutoStart 0.

```
# /opt/VRTSvcs/bin/haconf -makerw
# /opt/VRTSvcs/bin/hagrp -modify oracle_resource AutoStart 0
# /opt/VRTSvcs/bin/hares -modify cssd_resource AutoStart 0
# /opt/VRTSvcs/bin/haconf -dump -makero
```

- 2 Offline Oracle and all parallel groups on *galaxy* and *nebula* except CVM.

- If the Oracle instance is under VCS control:
 

```
# /opt/VRTSvcs/bin/hagrp -offline oracle_group -sys galaxy
# /opt/VRTSvcs/bin/hagrp -offline oracle_group -sys nebula
```
- If the Oracle instance is not under VCS control, log in as Oracle user and shutdown all the instances on *galaxy* and *nebula*.

For example:

```
$ srvctl stop instance -d database_name -I instance_name -n
galaxy
$ srvctl stop instance -d database_name -I instance_name -n
nebula
$ srvctl stop nodeapps -n galaxy
$ srvctl stop nodeapps -n nebula
```

- 3 Stop CRS daemon on node *galaxy* and *nebula*. If the CRS daemon is under VCS control, offline it.

```
# /opt/VRTSvcs/bin/hares -offline cssd_resource -sys galaxy
# /opt/VRTSvcs/bin/hares -offline cssd_resource -sys nebula
```

- 4 Switch Failover groups from *galaxy* to *jupiter* and *mercury*.

```
# /opt/VRTSvcs/bin/hagrp -switch failover_group -to jupiter/mercury
```

- 5 Unmount all the VxFS file system which is not under VCS control.

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

- 6 Offline CVM group:

```
# /opt/VRTSvcs/bin/hagrp -offline cvm-group -sys galaxy
# /opt/VRTSvcs/bin/hagrp -offline cvm-group -sys nebula
```

- 7 Log in to *galaxy* and *nebula* as root user and stop VCS.

```
# hastop -local
```

- 8 Comment out CRS-specific lines in */etc/inittab* file to prevent the system from starting the CRS daemons during system startup. On each system, prefix the CRS lines with a colon (":") to comment them out.

For example:

```
:h1:2:respawn:/etc/init/evmd run >/dev/null 2>&1 </dev/null
:h2:2:respawn:/etc/init/cssd fatal >/dev/null 2>&1 </dev/null
:h3:2:respawn:/etc/init.crsd run >/dev/null 2>&1 </dev/null
```

- 9 Login node *galaxy* and *nebula*, invoke `slibclean` to unload all unused modules:

```
# /usr/sbin/slibclean
```

- 10 Log in to *galaxy* and *nebula*, stop VEA if it is running.

```
# /opt/VRTSob/bin/vxsvctr1 stop
```

- 11 Change to CD-ROM directory.

```
# cd /cdrom
```

- 12 Start the installer.

```
# ./installmp [-rsh] galaxy nebula
```

- 13 Restart each system in the cluster when prompted.

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

- 14 When *galaxy* and *nebula* come up, make sure *galaxy* and *nebula* join *jupiter* and *mercury* at GAB/LLT/VXFEN level, however CVM/CFS/HAD will not come up due to mismatch versions. Verify the ports a, b, d, o all have 0123 membership.

For example, enter on node *galaxy*:

```
# /sbin/gabconfig -a
GAB Port Memberships
```

```

=====
Port a gen 5c3d0b membership 0123
Port b gen 5c3d10 membership 0123
Port d gen 5c3d0c membership 0123
Port o gen 5c3d0f membership 0123
Enter on jupiter:
# /sbin/gabconfig -a
GAB Port Memberships
=====
Port a gen 5c3d0b membership 0123
Port b gen 5c3d10 membership 0123
Port d gen 5c3d0c membership 0123
Port f gen 5c3d15 membership ; 23
Port f gen 5c3d15 visible 01
Port h gen 5c3d2a membership ; 23
Port h gen 5c3d2a visible 01
Port o gen 5c3d0f membership 0123
Port v gen 5c3d12 membership ; 23
Port v gen 5c3d12 visible 01
Port w gen 5c3d13 membership ; 23
Port w gen 5c3d13 visible 01

```

**To prepare for relinking after Oracle 10g upgrade**

1 Perform the following setup steps on *galaxy* and *nebula* in the cluster before bringing up CRS:

- Unload any unused modules:
 

```
# /usr/sbin/slibclean
```
- Create the 32-bit SKGXN archive library under /opt/VRTSvcs/rac/lib:
 

```
# cd /opt/VRTSvcs/rac/lib
# rm libskgx2.a
# cp libvcsmm.so shr_skgxn2.o
# ar -X32 -qv libskgx2.a shr_skgxn2.o
# rm shr_skgxn2.o
```
- Create the 64-bit SKGXN archive library under /opt/VRTSvcs/rac/lib64:
 

```
# cd /opt/VRTSvcs/rac/lib64
# rm libskgx2.a
# cp libvcsmm.so shr_skgxn2.o
# ar -X64 -qv libskgx2.a shr_skgxn2.o
# rm shr_skgxn2.o
```

2 If you are running Oracle 10gR1, copy these SF Oracle RAC SKGXN libraries to \$CRS\_HOME on *galaxy* and *nebula*.

```

# /usr/sbin/slibclean
# cp /opt/VRTSvcs/rac/lib/libskgx2.a \
$CRS_HOME/lib32/libskgx2.a
# cp /opt/VRTSvcs/rac/lib/libskgx2.a \
$CRS_HOME/lib32/libskgx2r.a

```

```
# cp /opt/VRTSvcs/rac/lib64/libskgxn2.a \
$CRS_HOME/lib/libskgxn2.a
# cp /opt/VRTSvcs/rac/lib64/libskgxn2.a \
$CRS_HOME/lib/libskgxn2.a
```

- 3 If you're running Oracle 10gR2, perform the following commands on node *galaxy* and *nebula*.

- Remove the following files from under /opt/ORCLcluster:

```
# /usr/sbin/slibclean
# rm /opt/ORCLcluster/lib/libskgxn2.a
# rm /opt/ORCLcluster/lib/libskgxn2.so
# rm /opt/ORCLcluster/lib32/libskgxn2.a
# rm /opt/ORCLcluster/lib32/libskgxn2.so
```

- Copy these SF Oracle RAC SKGXN libraries to /opt/ORCLcluster:

```
# cp /opt/VRTSvcs/rac/lib/libskgxn2.a \
/opt/ORCLcluster/lib32/libskgxn2.a
# cp /opt/VRTSvcs/rac/lib/libvcsmm.so \
/opt/ORCLcluster/lib32/libskgxn2.so
# cp /opt/VRTSvcs/rac/lib64/libskgxn2.a \
/opt/ORCLcluster/lib/libskgxn2.a
# cp /opt/VRTSvcs/rac/lib64/libvcsmm.so \
/opt/ORCLcluster/lib/libskgxn2.so
```

Run the commands on each system in the cluster.

- Copy the SF Oracle RAC SKGXP library to /opt/ORCLcluster:

```
# cd /opt/ORCLcluster/lib32
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_32.so .
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_32.a .
# cd /opt/ORCLcluster/lib
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.so .
# cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.a .
```

- Copy the clsinfo script to /opt/ORCLcluster/bin.

```
# cd /opt/ORCLcluster/bin
# cp /opt/VRTSvcs/rac/bin/clsinfo .
```

### To relink Oracle libraries on the first half of the Oracle 10g cluster

- 1 Invoke installsfrac:

```
# cd /opt/VRTS/install
# ./installsfrac [-rsh] -configure galaxy nebula
```

- 2 Navigate to the "Install or Relink Oracle" menu.

- Select the appropriate Oracle 10g version:

```
1) Oracle 9iR2
2) Oracle 10gR1
3) Oracle 10gR2
```

- Select "Relink Oracle" from the menu.

```
1) Install Oracle Clusterware (CRS)
2) Install Oracle RDBMS server
```

- 3) Relink Oracle
- b) [Go to previous menu]

- 3 In the "Install or Relink Oracle" menu, enter the required information. For example:

```
Enter Oracle UNIX user name: [b] (oracle) oracle
Enter Oracle UNIX group name: [b] (oinstall) oinstall
Enter Oracle base directory: [b] /app/oracle
Enter absolute path of Database Home directory: [b] /app/oracle/orahome
Enter Oracle Patch Level: [b] (0.1) 0.1
```

- 4 Confirm your responses in the verification screen. The installer copies the SF 5.0 Oracle RAC libraries to /opt/ORCLcluster, where it expects libskgxn.

```
Oracle environment information verification
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Clusterware (CRS) Home: /app/crshome
Oracle Release: 10.2
Oracle Patch Level: 0.1
Oracle Base: /app/oracle
Switch Applications to 5.0MP1 cluster
```

After relinking the Oracle libraries for the first half of your cluster, you can begin to upgrade the second half of your cluster.

### To switch applications to the upgraded Oracle 10g cluster nodes

---

**Note:** Your downtime starts now.

---

- 1 On the second half of the cluster nodes, back up the configuration files:
 

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.save
```
- 2 Offline Oracle and all parallel groups on *jupiter* and *mercury* except CVM.
  - If the Oracle instance is under VCS control:
 

```
# /opt/VRTSvcs/bin/hagrp -offline oracle_group -sys jupiter
# /opt/VRTSvcs/bin/hagrp -offline oracle_group -sys mercury
```
  - If the Oracle instance is not under VCS control, log in as Oracle user and shutdown all the instances on *jupiter* and *mercury*.
 For example:
 

```
$ srvctl stop instance -d database_name -I instance_name \
-n jupiter
$ srvctl stop instance -d database_name -I instance_name \
-n mercury
$ srvctl stop nodeapps -n jupiter
$ srvctl stop nodeapps -n mercury
```
- 3 Stop CRS daemon on *jupiter* and *mercury*. If the CRS daemon is under VCS control, offline it.

- ```
# /opt/VRTSvcs/bin/hares -offline cssd-resource -sys jupiter
# /opt/VRTSvcs/bin/hares -offline cssd-resource -sys mercury
```
- 4 Offline Failover groups on *jupiter* and *mercury*.
 

```
# /opt/VRTSvcs/bin/hagrp -offline failover_group -sys jupiter
# /opt/VRTSvcs/bin/hagrp -offline failover_group -sys mercury
```
  - 5 Unmount all the VxFS file system which is not under VCS control.
 

```
# mount |grep vxfs
# fuser -c /mount_point
# umount /mount_point
```
  - 6 Offline CVM group:
 

```
# /opt/VRTSvcs/bin/hagrp -offline cvm-group -sys galaxy
# /opt/VRTSvcs/bin/hagrp -offline cvm-group -sys nebula
```
  - 7 Log in *galaxy* as root user, stop VCS
 

```
# /opt/VRTSvcs/bin/hastop -all
```
  - 8 On *galaxy* and *nebula*, as root user, removing the prefix ":" to uncomment the CRS-specific lines in the `/etc/inittab` file.
 

```
h1:2:respawn:/etc/init/evmd run >/dev/null 2>&1 </dev/null
h2:2:respawn:/etc/init/cssd fatal >/dev/null 2>&1 </dev/null
h3:2:respawn:/etc/init.crsd run >/dev/null 2>&1 </dev/null
```
  - 9 On all nodes, restore the configuration files from the backups created in an earlier procedure.
 

See [“To prepare the first half of the Oracle 10g cluster for minimum downtime upgrade”](#) on page 23.

    - Restore the original main.cf file:
 

```
# cp /etc/VRTSvcs/conf/config/main.cf.save \
/etc/VRTSvcs/conf/config/main.cf
```
    - Merge any changes you made to the original types.cf file and backed up in a previous procedure at:
 

```
/etc/VRTSvcs/conf/config/types.cf.save
```

 with the new types.cf file at:
 

```
/etc/VRTSvcs/conf/types.cf
```
  - 10 Start VCS on *galaxy* and *nebula*.
 

```
# /opt/VRTSvcs/bin/hastart
```
  - 11 Start CRS daemon on node *galaxy* and *nebula*:
    - If CRS is not configured under VCS control, log in as superuser on *galaxy* and *nebula* in the cluster and enter:
 

```
# /etc/init.crs start
```
    - If CRS is configured under VCS control, log in as superuser and enter the following command for each system in the cluster:
 

```
# /opt/VRTSvcs/bin/hares -online cssd_resource -sys galaxy
# /opt/VRTSvcs/bin/hares -online cssd_resource -sys nebula
```

## 12 Bring Oracle online on *galaxy* and *nebula*.

- If oracle instance is under VCS control:
 

```
# /opt/VRTSvcs/bin/hagrp -online oracle_group -sys galaxy
# /opt/VRTSvcs/bin/hagrp -online oracle_group -sys nebula
```
- If Oracle instance is not under VCS control, log in as Oracle user and start all the instances on *galaxy* and *nebula*. For example:
 

```
$ srvctl start instance -d database_name -I instance_name \
  -n galaxy
$ srvctl start instance -d database_name -I instance_name \
  -n nebula
```

## 13 Bring online all the applications on *galaxy* and *nebula* which were offline before upgrade.

---

**Note:** Your downtime ends now. You have one Oracle instance running on *galaxy* and one instance running on node *nebula*.

---

## To upgrade the second half of the Oracle 10g cluster

- 1 Log into node *jupiter* and *mercury*.
- 2 Comment out CRS-specific lines in `/etc/inittab` file to prevent the system from starting the CRS daemons during system startup. On each system, prefix the CRS lines with a colon (":") to comment them out. For example:
 

```
:h1:2:respawn:/etc/init/evmd run >/dev/null 2>&1 </dev/null
:h2:2:respawn:/etc/init/cssd fatal >/dev/null 2>&1 </dev/null
:h3:2:respawn:/etc/init.crsd run >/dev/null 2>&1 </dev/null
```
- 3 Log in to *jupiter* and *mercury*, invoke `slibclean` to unload all unused modules:
 

```
# /usr/sbin/slibclean
```
- 4 Stop VEA if it is running:
 

```
# /opt/VRTSob/bin/vxsvcctrl stop
```
- 5 Upgrade SFRAC on *jupiter* and *mercury*.
 

```
# cd /cdrom
For ssh
# ./installmp jupiter mercury
For rsh
# ./installmp -rsh jupiter mercury
```
- 6 Restart each system in the cluster when prompted.
 

```
# /usr/sbin/shutdown -r
```
- 7 When *jupiter* and *mercury* come back online, they will join *galaxy* and *nebula* to form a four-node SFRAC 5.0 MP1 cluster.
 

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

```

Port a gen 5c3d0d membership 0123
Port b gen 5c3d13 membership 0123
Port d gen 5c3d0e membership 0123
Port f gen 5c3d33 membership 0123
Port h gen 5c3d2d membership 0123
Port o gen 5c3d12 membership 0123
Port v gen 5c3d2f membership 0123
Port w gen 5c3d31 membership 0123

```

- 8 To prepare for relinking after Oracle 10g upgrade  
See [“To prepare for relinking after Oracle 10g upgrade”](#) on page 25.
- 9 Verify CRS is offline before relinking.
- 10 Relink the Oracle libraries on the second half of the Oracle 10g cluster:  
See [“To relink Oracle libraries on the first half of the Oracle 10g cluster”](#) on page 26.

#### To bring up applications on the second half of the Oracle 10g cluster

- 1 On *jupiter* and *mercury*, as root user, removing the prefix ":" to uncomment the CRS-specific lines in the `/etc/inittab` file.

```

h1:2:respawn:/etc/init/evmd run >/dev/null 2>&1 </dev/null
h2:2:respawn:/etc/init/cssd fatal >/dev/null 2>&1 </dev/null
h3:2:respawn:/etc/init.crsd run >/dev/null 2>&1 </dev/null

```

- 2 Start the CRS daemon on *jupiter* and *mercury*:
  - If CRS is not configured under VCS control, log in as superuser on *jupiter* and *mercury* in the cluster and enter:
 

```
# /etc/init.crs start
```
  - If CRS is configured under VCS control, log in as superuser and enter the following command for each system in the cluster:
 

```
# /opt/VRTSvcs/bin/hares -online cssd_resource -sys jupiter
# /opt/VRTSvcs/bin/hares -online cssd_resource -sys mercury
```
- 3 Bring Oracle online on *jupiter* and *mercury*.
  - If oracle instance is under VCS control
 

```
# /opt/VRTSvcs/bin/hagrp -online oracle_resource \
  -sys jupiter
# /opt/VRTSvcs/bin/hagrp -online oracle_resource \
  -sys mercury
```
  - If Oracle instance is not under VCS control, Login as Oracle user and start all the instances on *jupiter* and *mercury*. For example:
 

```
$ srvctl start instance -d database_name -I instance_name \
  -n jupiter
$ srvctl start instance -d database_name -I instance_name \
  -n mercury
```
- 4 Bring up all the applications which were offline before upgrade on *jupiter* and *mercury*.

## 5 Change oracle group and cssd AutoStart attribute to 1

```
# /opt/VRTSvcs/bin/haconf -makerw
# /opt/VRTSvcs/bin/hares -modify oracle_resource AutoStart 1
# /opt/VRTSvcs/bin/hares -modify cssd_resource AutoStart 1
# /opt/VRTSvcs/bin/haconf -dump -makero
```

### To commit the updated filesets

Use: `installp -C filset_name`

The updated filesets can still be removed manually after running `installmp`, and before committing. After committing, the updated filesets can not be rolled back to the previous release.

## Upgrading clusters with Oracle9i

---

**Note:** There will be some downtime involved. Review the entire procedure and carefully plan your downtime before proceeding with any steps.

---

Assumptions:

- SF Oracle RAC 5.0 is installed
- Oracle9i R2 is installed
- Oracle binaries are installed on separate nodes for each node in the cluster

The example for this procedure assumes a four-node SF Oracle RAC 5.0 cluster with nodes *galaxy*, *nebula*, *jupiter* and *mercury*.

Upgrade the half of the nodes in the cluster at a time. The example nodes are *galaxy* and *nebula*.

### To prepare the first half of the Oracle9i cluster for minimum downtime upgrade

On the first half of the cluster nodes, back up the configuration files:

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

### To upgrade the first half of the Oracle9i cluster

#### 1 Set Oracle group AutoStart 0.

```
# haconf -makerw
# /opt/VRTSvcs/bin/hares -modify oracle_resource AutoStart 0
# haconf -dump -makero
```

#### 2 Offline all Parallel groups on *galaxy* and *nebula* except CVM. If Oracle instance is under VCS control, offline it.

- ```
# /opt/VRTSvcs/bin/hagrp -offline group_name -sys galaxy
# /opt/VRTSvcs/bin/hagrp -offline group_name -sys nebula
```
- 3 Stop the GSD daemon

```
$ gsdctl stop
```
  - 4 Switch Failover groups from *galaxy* to any other node *jupiter* and *mercury*.

```
# /opt/VRTSvcs/bin/hagrp -switch failover_group -to jupiter/mercury
```
  - 5 Unmount all the vxfs file system which is not under VCS control.

```
# mount | grep vxfs
# fuser -c /mount_point
# umount /mount_point
```
  - 6 Offline CVM group

```
# /opt/VRTSvcs/bin/hagrp -offline cvm_group -sys galaxy
# /opt/VRTSvcs/bin/hagrp -offline cvm_group -sys nebula
```
  - 7 Stop VCS on node galaxy and nebula

```
# hastop -local
```
  - 8 Stop VEA if it is running

```
# /opt/VRTSob/bin/vxsvcctrl status
# /opt/VRTSob/bin/vxsvcctrl stop
```
  - 9 Change to CD-ROM directory, start the installer

```
# cd /cdrom
# ./installmp [-rsh] galaxy nebula
```
  - 10 Restart each system in the cluster when prompted.

```
# /usr/sbin/shutdown -r
```
  - 11 Make sure *galaxy* and *nebula* join *jupiter* and *mercury* at GAB/LLT/VXFEN level, however CVM/CFS/HAD will not come up due to mismatch versions. Verify the ports a, b, d, o all have 0123 membership. For example:  
Enter on node *galaxy*,

```
# /sbin/gabconfig -a
GAB Port Memberships
=====
Port a gen 5c3d0b membership 0123
Port b gen 5c3d10 membership 0123
Port d gen 5c3d0c membership 0123
Port o gen 5c3d0f membership 0123
```

Enter on node *jupiter*,

```
# /sbin/gabconfig -a
GAB Port Memberships
=====
Port a gen 5c3d0b membership 0123
Port b gen 5c3d10 membership 0123
Port d gen 5c3d0c membership 0123
Port f gen 5c3d15 membership ; 23
Port f gen 5c3d15 visible 01
```

```

Port h gen 5c3d2a membership ; 23
Port h gen 5c3d2a visible 01
Port o gen 5c3d0f membership 0123
Port v gen 5c3d12 membership ; 23
Port v gen 5c3d12 visible 01
Port w gen 5c3d13 membership ; 23
Port w gen 5c3d13 visible 01

```

### To relink Oracle 9iR2 library

#### 1 Invoke installsfrac:

```

# cd /opt/VRTS/install
# ./installsfrac [-rsh] -configure galaxy nebula

```

#### 2 Navigate to the “Install or Relink Oracle” menu.

- Select the Oracle 9i version:
  - 1) Oracle 9iR2
  - 2) Oracle 10gR1
  - 3) Oracle 10gR2
- Select “Relink Oracle” from the menu:
  - 1) Install Oracle RDBMS server
  - 2) Relink Oracle
  - b) [Go to previous menu]

#### 3 In the “Install or Relink Oracle” menu, enter the required information. For example:

```

Enter Oracle UNIX user name: [b] (oracle) oracle
Enter Oracle UNIX group name: [b] (oinstall) oinstall
Enter Oracle base directory: [b] /app/oracle
Enter absolute path of Database Home directory: [b] /app/oracle/
orahome
Enter Oracle Patch Level: [b] (0.1) 0.7

```

#### 4 Confirm your responses in the verification screen. The installer copies the SF 5.0 Oracle RAC libraries to /opt/ORCLcluster, where it expects libskgxn.

```

Oracle environment information verification
Oracle Unix User: oracle
Oracle Unix Group: oinstall
Oracle Release: 9.2
Oracle Patch Level: 0.7
Oracle Base: /app/oracle
Oracle Home: /app/oracle/orahome
Is this information correct? [y,n,q] (y)

```

---

**Note:** Your downtime starts now.

---

### To switch applications to the SF Oracle RAC 5.0 MP1 cluster

- 1 Offline all Parallel groups on *jupiter* and *mercury* except CVM. If Oracle instance is under VCS control, offline it.

- ```
# /opt/VRTSvcs/bin/hagrp -offline group_name -sys jupiter
# /opt/VRTSvcs/bin/hagrp -offline group_name -sys mercury
```
- 2 Stop the GSD daemon.
 

```
$ gsdctl stop
```
  - 3 Offline Failover groups *jupiter* and *mercury*.
 

```
# /opt/VRTSvcs/bin/hagrp -offline failover_group -sys jupiter
# /opt/VRTSvcs/bin/hagrp -offline failover_group -sys mercury
```
  - 4 Unmount all the vxfs file system which is not under VCS control.
 

```
# mount | grep vxfs
# fuser -c /mount_point
# umount /mount_point
```
  - 5 Offline the CVM group.
 

```
# /opt/VRTSvcs/bin/hagrp -offline cvm_group -sys galaxy
# /opt/VRTSvcs/bin/hagrp -offline cvm_group -sys nebula
```
  - 6 Stop VCS on node *jupiter* and *mercury*.
 

```
# /opt/VRTSvcs/bin/hastop -all
```
  - 7 On all nodes, restore the configuration files from the backups created in an earlier procedure.
 

See [“To prepare the first half of the Oracle 10g cluster for minimum downtime upgrade”](#) on page 23.

    - Restore the original main.cf file:
 

```
# cp /etc/VRTSvcs/conf/config/main.cf.save \
/etc/VRTSvcs/conf/config/main.cf
```
    - Merge any changes you made to the original types.cf file and backed up in a previous procedure at:
 

```
/etc/VRTSvcs/conf/config/types.cf.save
```

 with the new types.cf file at:
 

```
/etc/VRTSvcs/conf/types.cf
```
  - 8 Start VCS on *galaxy* and *nebula*.
 

```
# /opt/VRTSvcs/bin/hastart
```
  - 9 Start gsd. Bring Oracle online on *galaxy* and *nebula*.
 

```
$ gsdctl start
# /opt/VRTSvcs/bin/hagrp -online oracle group -sys galaxy/nebula
```

---

**Note:** Your downtime ends now. You have one oracle instance running on node *galaxy* and one instance running on node *nebula*.

---

### To upgrade the second half of the Oracle9i cluster

- 1 Stop VEA if it's running
 

```
# /opt/VRTSob/bin/vxsvectl status
```

- ```
# /opt/VRTSob/bin/vxsvcctl stop
```
- 2 Change to CD-ROM directory, start the installer
 

```
# cd /cdrom
# ./installmp [-rsh] jupiter mercury
```
  - 3 Restart each system in the cluster when prompted.
 

```
# /usr/sbin/shutdown -r
```
  - 4 When *jupiter* and *mercury* come back online, they will join *galaxy* and *nebula* to form a four-node SFRAC 5.0MP1 cluster.
  - 5 Relink the Oracle libraries on the second half of the Oracle 10g cluster: See [“To relink Oracle 9iR2 library”](#) on page 33.
  - 6 Bring online the Oracle resources configured under VCS.
 

```
# hares -online listener_res -sys galaxy
# hares -online listener_res -sys nebula
# hagrps -online group_name -sys galaxy
# hagrps -online group_name -sys nebula
```
  - 7 Verify that the Oracle instances are up and running.
  - 8 Reset the VCS resource attribute AutoStart to 1 for all resources that manage Oracle and database instances. As root user on any system in the cluster, enter:
 

```
# haconf -makerw
# hares -modify oracle_res AutoStart 1
# haconf -dump -makero
```

### To commit the updated filesets

Use: `installmp -C fileset_name`

The updated filesets can still be removed manually after running `installmp`, and before committing. After committing, the updated filesets can not be rolled back to the previous release.

## Upgrading the SFDB repository

After you have completed the SF Oracle RAC upgrade, if you plan to use the commands for Checkpoints, FlashSnap, or Storage Mapping, you must:

- Change permissions
- Upgrade the repository

---

**Note:** Do not recursively change permissions, groups, or owners.

---

### To change permissions

- 1 Change permissions for the following directory:

For Veritas Storage Foundation for Oracle:

```
# chmod 550 /opt/VRTSdbed
```

- 2 Reset owner and group settings to the appropriate owner and group for the database administrators on your system.

For example, to change owner to the user `oracle` and the group `dba`, run the following command:

```
# chown oracle:dba /opt/VRTSdbed
```

#### To upgrade the repository in a cluster environment

- 1 Unconfigure the SFUA repository from the VCS configuration:
 

```
# /opt/VRTSdbcom/bin/sfua_db_config -o unconfig_cluster
```
- 2 Mount the repository file system manually.
- 3 Run the repository upgrade command again with no options:
 

```
# /opt/VRTSdbcom/bin/sfua_db_config
```

## Operation notes

The following information includes guidelines, tips, and other considerations for operating the product. See the *Storage Foundation for Oracle RAC Installation and Configuration Guide* for complete procedures.

### Enabling fencing

When the product installer completes installation of SF Oracle RAC, I/O fencing is disabled on all cluster nodes. Symantec does not support using SF Oracle RAC with I/O fencing disabled. See the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* for details on configuring I/O fencing in a cluster environment.

### Setting LIBPATH in the EnvFile attribute of Veritas agent for Oracle

If you do not specify the use of the EnvFile attribute in the Oracle resource definition in the `main.cf` file, Oracle starts with the root user's LIBPATH. This setting may not be appropriate for Oracle.

Edit the file (or files for multiple instances) that correspond to the EnvFile attribute as specified in the Oracle resource definition.

For example, to set the variable, edit:

```
LIBPATH=$ORACLE_HOME/lib:${LIB_PATH}
export LIB_PATH
```

Set `$ORACLE_HOME` to the home directory of Oracle in your environment.

## Determining if disk arrays have array support library (ASL)

Each disk array you use for shared storage with Storage Foundation for Oracle RAC requires an Array Support Library. Use the following command to list all ASLs currently available on your system:

```
# vxddladm listsupport
```

Use the following command to list arrays on your system and the ASLs they currently use:

```
# vxdmpadm listenclosure all
```

## Preventing automatic database startup for Oracle 10g R2

CRS and Oracle agent might attempt to start the instance at the same time if the database mount is available. To prevent automatic database starting, you must change the Management policy for the database (automatic, manual) to MANUAL using SRVCTL command. The command will change AUTO\_START attribute of CRS database and instance resources. [See Etrack Incident 838841]

### To prevent automatic database starting

#### ◆ Depending on the status of the database:

- If the database is already registered and running:

To change it to manual execute the following commands as Oracle user:

```
$ srvctl stop database -d db-name
```

```
$ srvctl modify database -d db-name -y manual
```

- If the database is not registered, use the following command to register database:

```
$ srvctl add database -d db-name -p \  
location-of-parameter-file -y manual
```

## Open issues

The following issues were reported for this release:

### Issues affecting Oracle 10g Release 2

#### Oracle VIP Configuration Assistant

During Oracle 10g Release 2 installation, the VIP Configuration Assistant may fail with the following error message:

"The given interface(s), "en0" is not public. Public interfaces should be used to configure virtual IPs."

**Workaround:** Invoke the vipca utility manually.

```
# export DISPLAY=nebula:0.0
# $CRS_HOME/bin/vipca
```

## Oracle Cluster Verification Utility Warning message

During the final stage of Oracle 10g Release 2 installation, you may receive a warning message with the Oracle Cluster Verification Utility. For example:

```
Utility
=====
OUI-25031:Some of the configuration assistants failed. It is
strongly recommended that you retry the configuration
assistants at this time. Not successfully running any
"Recommended" assistants means your system will not be correctly
configured.
1. Check the Details panel on the Configuration Assistant Screen
to see the errors resulting in the failures.
2. Fix the errors causing these failures.
3. Select the failed assistants and click the 'Retry' button to
retry them.

=====
```

**Workaround:** You can safely ignore this message.

## Deporting and importing of shared disk groups

If you manually deport a shared disk group, the CVMVolDg agent does not automatically reimport it as a shared group. You must manually reimport it as a shared group.

## Oracle 10.1.0.3 databases created by DBCA may fail on reboot

When a cluster runs an Oracle 10g (10.1.0.3) database created by the Oracle Database Creation Assistant (DBCA), shutting down or rebooting one node may kill the database instances on other nodes. This problem occurs when CRS control of the database is disabled. The 10.1.0.4 release addresses this issue.

**Workaround:** Refer to the following TechNote on the Technical Support website for the latest information on this problem:

<http://support.veritas.com/docs/275388>.

## Handling of RAW files with ODM enabled for Oracle 10g

Oracle can use RAW files for data files, control files and redo logs. Previously, the SQL REUSE clause was ignored if the associated file was a RAW file. From Oracle 10g onwards, the REUSE clause of SQL is mandatory when Oracle Disk Manager (ODM) is enabled. [See Etrack Incident 540429]

**Workaround:** To create database files on raw volumes, you must set the REUSE option in the SQL commands. For examples of error messages and SQL workarounds, see the following Technical Support TechNote:

<http://support.veritas.com/docs/283362>

## Stopping CRS before stopping VCS in Oracle 10g clusters

In SF Oracle RAC clusters with Oracle 10g, if CRS is not under VCS control, stop the CRS daemons before using `hastop -all` command to stop VCS.

To enable VCS to control CRS daemons, refer to the information on configuring VCS service groups for Oracle 10g in the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

## Avoid using lltconfig -u in an Oracle 10g RAC environment

Using `lltconfig -u` in an Oracle 10g RAC environment to unconfigure a link with Oracle's private IP address plumbed on it prevents the `PrivNIC` agent from failing over this IP address. This usage may result in a machine reboot by the Oracle CRS.

## Changing the Veritas agent for Oracle error handling

The Veritas agent for Oracle provides enhanced handling of Oracle errors encountered during detailed monitoring. The agent uses the reference file, `oraerror.dat`, which provides a list of Oracle errors and actions to address the errors.

Refer to the *Veritas High Availability Agent for Oracle Installation and Configuration Guide* for a description of the actions.

Currently, the file specifies the NOFAILOVER action for the following Oracle errors:

ORA-00061, ORA-02726, ORA-6108, ORA-06114

The NOFAILOVER action means that the agent sets the state of the resource to OFFLINE and freezes the service group.

You can stop the agent, edit `oraerror.dat`, and change the NOFAILOVER action to another action that is appropriate for your environment. The changes go into effect when you restart the agent.

## Stopping cluster nodes configured with I/O fencing

The I/O fencing feature protects against data corruption resulting from a failed cluster interconnect or "split brain." See "VCS Communications, Membership,

and I/O Fencing” in the *VCS User’s Guide* for details on the problems a failed interconnect can create and the protection I/O fencing provides.

I/O fencing uses SCSI-3 Persistent Reserve keys to implement data protection. The software places keys on I/O fencing coordinator and data disks. The administrator must be aware of several operational changes needed when working with clusters protected by I/O fencing. Specific shutdown procedures ensure keys are removed from coordinator disks and data disks to prevent possible difficulties with subsequent cluster startup.

Using the `reboot` command rather than the `shutdown` command bypasses shutdown scripts and can leave keys on the coordinator and data disks. Depending on the order of reboot and subsequent startup events, the cluster might warn of a possible split brain condition and fail to start up.

**Work round:** Use the `shutdown` command instead of the `reboot` command to perform a graceful reboot for systems.

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

## Number of links required for the PrivNIC agent

Interfaces specified in the PrivNIC configuration should be exactly the same in name and total number as those which have been used for LLT configuration. [See Etrack Incident 587136]

The attribute table on page 459 of the Storage Foundation for Oracle RAC Installation and Configuration Guide should be replaced with the following attribute definition table.

**Table 1-1** Required attributes for PrivNIC agent

Attribute	Dimension	Description
Device	string - association	<p>Specifies the network interface device as shown by the “ifconfig” command and the “network-id” associated with the interface. Network-ids of the interfaces connected to the same physical network must match. The interface with the lower network-id has the higher preference for failover. Interfaces specified in the PrivNIC configuration should be exactly the same in name and total number as those which have been used for LLT configuration.</p> <p>Example:</p> <pre>Device@galaxy = {en1=0, en2=1, en3=2} Device@nebula = {en1=0, en2=1, en3=2}</pre>

**Table 1-1** Required attributes for PrivNIC agent

Attribute	Dimension	Description
Address	string-scalar	The numerical private IP address. For example: Address = "192.168.12.1"
NetMask	string - association	The numerical netmask for the private IP address. For example: Address = "255.255.255.0"

## Stopping VCS does not unregister port f from GAB membership

In an SF Oracle RAC cluster with all the CFS resources under VCS control, when you stop VCS, all the CFS resources must go down cleanly and CFS must unregister port f from GAB membership.

Oracle 10g Cluster Ready Services (CRS) does not clean up all its processes when it is stopped. Now, when you stop VCS, all the CFS resources go down. However, due to the left over Oracle processes, CFS does not unregister port f from GAB membership.

**Workaround:** Perform the following steps to bring down port f.

- 1 Kill all the Oracle processes.  
# `kill -9 `ps -u oracle|awk '{print $1}'``
- 2 Verify that all CFS files systems are unmounted.  
# `mount | grep cluster`
- 3 Unregister port f from GAB membership.  
# `fsclustadm cfsdeinit`

## The gcoconfig command assigns priority 0 to all nodes

If you configure GCO using the `/opt/VRTSvcs/bin/gcoconfig` command, the gcoconfig utility assigns the same priority '0' to all the nodes that are in the SystemList of the ClusterService group. [See Etrack Incident 857159]

**Workaround:** Edit main.cf and give the proper node number for the cluster nodes in the SystemList of the ClusterService group.

Use one of the following approaches to edit the main.cf file:

- Veritas Cluster Server GUI
- VCS commands
- Stop VCS and manually edit the main.cf file  
Note that this approach has HA downtime.

## Fixed issues

The following issues were corrected for this release. Fixed issues and software enhancement requests are referenced by incident number.

**Table 1-2** Fixed issues in SF Oracle RAC version 5.0 MP1

Incident number	Description
626451	The installer sometimes does not display the global clustering configuration option.
780737	libskgxp causes oracle to dump core when lmx helper thread is enabled.
790019	Disable lmx helper thread by default.
850513	Do not return EINTR in lmxpoll.
830225	PrivNIC agent does not work with CRS.

**Table 1-3** Fixed issues in SF Oracle RAC version 5.0

Incident number	Description
496151	Infinite start check loop in init.cssd.

## Documentation

Product guides are available on the documentation disc in PDF and HTML formats. We recommend copying pertinent information, such as installation guides and release notes, from the disc to your system directory /opt/VRTS/docs for reference.

**Table 1-4** Storage Foundation for Oracle RAC Documentation Set

Document Title	File Name
<b>Veritas Storage Foundation for Oracle RAC</b>	
<i>Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide</i>	sfrac_install.pdf
<i>Veritas Storage Foundation for Oracle RAC Release Notes</i>	sfrac_notes.pdf
<b>Veritas Cluster Server</b>	
<i>Veritas Cluster Server Installation Guide</i>	vcs_install.pdf

**Table 1-4** Storage Foundation for Oracle RAC Documentation Set

Document Title	File Name
<i>Veritas Cluster Server Release Notes</i>	vcs_notes.pdf
<i>Veritas Cluster Server User's Guide</i>	vcs_users.pdf
<i>Veritas Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents.pdf
<i>Veritas Cluster Server Agent Developer's Guide</i>	vcs_agent_dev.pdf
<i>Veritas Cluster Server Enterprise Agent for Oracle Installation and Configuration Guide</i>	vcs_oracle_install.pdf
<b>Veritas Storage Foundation</b>	
<i>Veritas Storage Foundation Release Notes</i>	sf_notes.pdf
<i>Veritas Storage Foundation Installation Guide</i>	sf_install.pdf
<i>Veritas Storage Foundation Intelligent Storage Provisioning Solutions Guide</i>	sf_isp_solutions.pdf
<i>Veritas Volume Manager Cross-platform Data Sharing Administrator's Guide</i>	sf_cds_admin.pdf
<i>Veritas FlashSnap Point-in-Time Copy Solutions Administrator's Guide</i>	flashsnap_admin.pdf
<b>Veritas Volume Manager</b>	
<i>Veritas Volume Manager Administrator's Guide</i>	vxvm_admin.pdf
<i>Veritas Volume Manager Troubleshooting Guide</i>	vxvm_tshoot.pdf
<b>Veritas File System</b>	
<i>Veritas File System Administrator's Guide</i>	vxfs_admin.pdf

## Documentation errata

Updates and corrections to the *Storage Foundation for Oracle RAC 5.0 Installation and Configuration Guide*.

### Creating the coordinator disk group

The *Storage Foundation for Oracle RAC 5.0 Installation and Configuration Guide* contains text that is incorrect. [See Etrack Incident 794321]

On page 113, replace:

#### To create the coordinator disk group

- 1 On one node, create the disk group by specifying the device name of one of the disks; the option `coordinator=on` sets the coordinator attribute:

```
# vxdbg -o coordinator=on init vxfencoorddg c1t1d0
```

- 2 Add the other two disks to the disk group:

```
# vxdbg -g vxfencoorddg adddisk c2t1d0
```

```
# vxdbg -g vxfencoorddg adddisk c3t1d0
```

Refer to the Veritas Volume Manager documentation for details on creating disk groups.

with:

#### To create the coordinator disk group

- 1 On one node, create the disk group by specifying the device name of one of the disks:

```
# vxdbg init vxfencoorddg rhdisk75
```

- 2 Add the other two disks to the disk group:

```
# vxdbg -g vxfencoorddg adddisk rhdisk76
```

```
# vxdbg -g vxfencoorddg adddisk rhdisk77
```

Refer to the Veritas Volume Manager documentation for details on creating disk groups.

- 3 Set the option coordinator=on to set the coordinator attribute:

```
# vxdbg -g vxfencoorddg set coordinator=on
```

## Stopping the CRS daemon

The *Storage Foundation for Oracle RAC 5.0 Installation and Configuration Guide* contains text on page 460 that is not completely accurate

Replace:

---

**Note:** VCS will not actually stop the CRS daemon. Instead, it will report an error message to the VCS engine log file if an offline is issued. Refer to the Oracle documentation to understand how to safely stop the CRS daemon. VCS will not start the CRS daemon. It will wait for the daemon to start automatically upon system boot up. If CRS daemon is stopped, refer to the Oracle documentation to understand how to safely restart the daemon.

---

with:

---

**Note:** If you are using Oracle 10.1.0.3 and 10.1.0.2, VCS will not actually stop the CRS daemon. Instead, it will report an error message to the VCS engine log file if an offline is issued. Refer to the Oracle documentation to understand how to safely stop the CRS daemon. VCS will not start the CRS daemon. It will wait for the daemon to start automatically upon system boot up. If CRS daemon is stopped, refer to the Oracle documentation to understand how to safely restart the daemon. This behavior does not apply for Oracle 10.1.0.4 and above.

---

## Corrections to adding and removing SF Oracle RAC on systems with Oracle 10g

The “Adding and removing SF Oracle RAC on systems with Oracle 10g” chapter is missing a step from following procedure[803809].

Add the following at the beginning of “To run the uninstall utility” on page 258.

- 1 Offline all Oracle parallel groups  

```
# hagrps -offline oracle_group -sys nebula
```
- 2 Offline CRS  

```
# hares -offline cssd_resource -sys nebula
```

If CRS is not under VCS control

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

The following changes need to be made to the “Adding and removing SF Oracle RAC on systems with Oracle 10g.” [See Etrack Incident 614479]

- Change vxfen, vcsmm, lmx paths for AIX.

On page 253, Step 9, replace the following commands:

```
# rm /etc/vxfenmode
# /etc/init.d/vxfen start
# /etc/init.d/vcsmm start
# /etc/init.d/lmx start
```

with:

```
# rm /etc/vxfenmode
# /etc/init.d/vxfen.rc start
# /etc/init.d/vcsmm.rc start
# /etc/init.d/lmx.rc start
```

- Add rsh and rcp check for AIX.

On page 254, add the following step:

Check that the rcp or rsh command works for Oracle. For example, log in as user Oracle to the node saturn and test the command to nebula:

```
$ rsh nebula ls
file_a
file_b
$
```

Make sure that rsh and rcp work between all nodes.

## Corrections to Preparing for global clustering chapter

The *Storage Foundation for Oracle RAC 5.0 Installation and Configuration Guide* contains text in the “Preparing for global clustering” chapter that is not completely accurate.

On page 276, under Installing SF Oracle RAC on the secondary site, add a bullet:

- Make sure you use the same oracle group and user IDs for both primary and secondary clusters.

## Corrections to setting up the repository database for Oracle

The following information is missing from the chapter “Configuring the repository database for Oracle.” It should be inserted at the end of the chapter. This information is required only for Oracle9i, since Oracle 10g is configured during installation. [See Etrack Incident 565419]

### Setting up Oracle 9i RAC objects with srvctl

When configured within an Oracle RAC environment, you must set up the Oracle `srvctl` service and register the name of the RAC database with `srvctl`, so that Veritas Storage Foundation for Oracle can learn the status of remote database instances. Otherwise, commands such as `dbed_ckptcreate -o offline` may fail.

#### To set up Oracle 9i RAC objects

- 1 Look in `/var/opt/oracle/srvConfig.loc` to learn the pathname to the SRVM configuration file as defined by the variable `srvconfig_loc`. For example:
 

```
srvconfig_loc=/orasrvm/srvm.ora
```
- 2 List the details of the SRVM configuration file with `ls -l`:
 

```
# ls -l /orasrvm/srvm.ora
```
- 3 If the configuration file does not exist, create and initialize the file:
 

```
# touch /orasrvm/srvm.ora
# srvconfig -init
```
- 4 If the configuration file exists, note the size of the file shown by the output of `ls -l`.
 

```
-rw-r--r-- 1 oracle dba 10569216 Jan 20 14:29 /orasrvm/srvm.ora
```
- 5 If the configuration file size is greater than zero (as shown in the example above), the file is initialized. If the file size is zero, initialize it:

- ```
# srvconfig -init
```
- 6 Start the Oracle RAC Manageability daemon on each system:
 

```
$ gsdctl start
```
  - 7 Confirm the GSD daemon status:
 

```
$ gsdctl stat
```

 GSD is running on the local node
  - 8 Add the database to the srvctl configuration:
 

```
$ srvctl add database -d vrts -o /app/oracle/orahome
```

```
$ srvctl config database
```

 vrts
  - 9 Add each instance to the configuration. For example, in a two-instance configuration, add the first instance:
 

```
$ srvctl add instance -d vrts -i vrts1 -n galaxy
```

```
$ srvctl config database -d vrts
```

 galaxy vrts /app/oracle/orahome
 Then add the second instance:
 

```
$ srvctl add instance -d vrts -i vrts2 -n nebula
```

```
$ srvctl config database -d vrts
```

 galaxy vrts1 /app/oracle/orahome
 nebula vrts2 /app/oracle/orahome
  - 10 Check the status of the instances to confirm they are running:
 

```
$ srvctl status database -d vrts
```

 Instance vrts1 is running on node galaxy
 Instance vrts2 is running on node nebula

## Single-host and two-host terminology

References to “single-host” and “two-host” configurations in the `dbed_vmclonedb` sections are not correct in an SF Oracle RAC configuration. These terms were borrowed from the Storage Foundation for Oracle Database environment, but “single-host” is not valid for SF Oracle RAC, and “two-host” is not correct for three or more node SF Oracle RAC configurations. [See Etrack Incident 840198]

For SF Oracle RAC configurations, change these references:

- From “single-host” to “same-node”
- From “two-host” to “node in the cluster”

## Corrections to Checkpoint commands

The following instances of `dbed` commands in the “Using Checkpoints and Storage Rollback with Storage Foundation for Oracle RAC” chapter of the *SF Oracle RAC Installation and Configuration Guide* are incorrect and must be changed. [See Etrack Incident 819369]

- Change `db2ed_clone` to `dbed_clonedb`  
 On page 328, Replace:  
 “Caution: If the `-o cio` option was used with the mount command to mount your primary database file systems, Concurrent I/O will not be preserved when using `db2ed_clonedb` to create a clone database.”  
 with:  
 “Caution: If the `-o cio` option was used with the mount command to mount your primary database file systems, Concurrent I/O will not be preserved when using `dbed_clonedb` to create a clone database.”
- Add extra space between “`dbed_ckptcreate`” and “`-S PROD.`”  
 On page 325, Step 1, Replace:  

```

$ /opt/VRTS/bin/dbed_ckptcreate-S PROD -H \
/oracle/product -o online
Creating online Storage Checkpoint of database PROD.
Storage Checkpoint Checkpoint_903937870 created.
$ mkdir /tmp/ckpt_ro
$ /opt/VRTS/bin/dbed_ckptmount-S PROD \
-c Checkpoint_903937870 -m /tmp/ckpt_ro
with:
$ /opt/VRTS/bin/dbed_ckptcreate -S PROD -H \
/oracle/product -o online
Creating online Storage Checkpoint of database PROD.
Storage Checkpoint Checkpoint_903937870 created.
$ mkdir /tmp/ckpt_ro
$ /opt/VRTS/bin/dbed_ckptmount -S PROD \
-c Checkpoint_903937870 -m /tmp/ckpt_ro

```

## Correction to `dbed_vmchecksnap` example `-t PROD`

Corrections are needed to the “Using database FlashSnap for backup and off-host processing” chapter for an invalid `dbed_vmchecksnap` example: `-t PROD`; it should be: `-t host2[839056]`

On page 363, replace:

```

dbed_vmchecksnap -S $ORACLE_SID -H $ORACLE_HOME -O setdefaults \
-p my_tag -f snap1 -t PROD
Snapplan snap1 for PROD
=====
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=host1
SECONDARY_HOST=host2

```

with:

```

dbed_vmchecksnap -S PROD -H $ORACLE_HOME -O setdefaults \
-p my_tag -f snap1 -t host2
^^^ NOTE: -S PROD and -t host2
Snapplan snap1 for PROD
=====

```

```
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=host1
SECONDARY_HOST=host2
```

## Finding FlashSnap status information

For Database FlashSnap status information, see *Veritas Storage Foundation for Oracle Administrator's Guide*. [See Etrack Incident 855452]

The 5.0 documentation disc contains the `sf_ora_admin.pdf`.

## Miscellaneous errata

The following need to be updated for the SF Oracle RAC Installation and Configuration Guide [See Etrack Incident 622495]:

- Update `ifconfig` commands in the chapter “Preparing to Install Oracle 10g RAC” on page 208:

Replace:

```
# ifconfig en0:1 plumb
# ifconfig en0:1 inet 10.11.12.60 netmask 255.255.255.0
# ifconfig en0:1 up
```

with:

```
# ifconfig en0 inet 10.11.12.60 netmask 255.255.255.0 broadcast
10.11.95.255 alias
# ifconfig en0 up
```

- Update Figure 7.1 in the chapter “Configuring Oracle9i service groups” on page 149:

Replace “`listener_eth_0`” with “`listener_en_0`.”

- Update `init.crs` path in the chapter “Upgrading and migrating Oracle software” on page 230, and again in the chapter “Uninstalling SF Oracle RAC from Oracle 10g systems” on page 266:

Replace:

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

with:

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

