

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

HP-UX

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.

HP-UX is a registered trademark of Hewlett-Packard Development Company, L.P.  
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

Visit [http://www.symantec.com/enterprise/support/assistance\\_care.jsp](http://www.symantec.com/enterprise/support/assistance_care.jsp) for technical assistance. 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. If you encounter an error when using a product, include the error number preceding the message when contacting Technical Services. You can also use the error number to search for information in TechNotes or documents on the website.



# Veritas Storage Foundation for Oracle RAC Release Notes

This document provides important information regarding Veritas™ Storage Foundation for Oracle RAC (SF Oracle RAC) for HP-UX. 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/281875>

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 SFRAC 5.0

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

- 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.

Item	Description
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.
Disks	Typical SF Oracle RAC configurations require that shared disks support applications that migrate between systems in the cluster.  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).
Disk space	SF Oracle RAC space needed for required depots: <ul style="list-style-type: none"><li>■ total: 4.1 GB</li><li>■ /opt: 2.3 GB</li><li>■ /usr: 425 MB</li><li>■ /var: 60 MB</li><li>■ /var/tmp: 600 MB</li><li>■ /tmp: 512 MB</li><li>■ /: 240 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 system at the required patch levels.

- September 2004 HP-UX 11i version 2.0 or later

## HP-UX Operating Environments

HP-UX 11i Operating Environments (OEs) are bundles of key applications available to users to help them install a complete operating system in a single-pass process. Symantec supports SF Oracle RAC for HP-UX on the following OEs:

HPUX11i-OE	HP-UX 11i v2 Foundation OE
HPUX11i-OE-Ent	HP-UX 11i v2 Enterprise OE
HPUX11i-OE-MC	HP-UX 11i v2 Mission Critical OE

### To identify the OE currently installed on your system

- ◆ Run the `swlist` command:

```
# swlist | grep HPUX
```

Output resembles:

```
HPUX11i-OE-MC          B.11.23.0409      HP-UX Mission
                    Critical Operating Environment Component
HPUXBaseAux           B.11.23.0409      HP-UX Base OS
                    Auxiliary
HPUXBaseOS            B.11.23           HP-UX Base OS
OnlineDiag            B.11.23.03.16     HPUX 11.23
                    Support Tools Bundle, Sep 2004
```

The output of this command includes a line that identifies the installed OE. For example:

```
HPUX11i-OE-MC          B.11.23.0409      HP-UX Mission
                    Critical Operating Environment Component
```

## HP-UX operating system patches

Install the following HP-UX depots and patches on each node before installing SF Oracle RAC:

DiskQuota-Enh	EnableVxFS	FSCmdsEnh
FSLibEnh	PHCO_32385	PHCO_32387
PHCO_32388	PHCO_32389	PHCO_32390
PHCO_32391	PHCO_32392	PHCO_32393
PHCO_32488	PHCO_32523	PHCO_32524
PHCO_32551	PHCO_32552	PHCO_32596
PHCO_32608	PHCO_32609	PHCO_32610

PHCO_32611	PHCO_33238	PHCO_34036
PHKL_32272	PHKL_32425	PHKL_32430
PHKL_32431	PHKL_33312	PHKL_34010

### To identify the patches currently installed on your system

- ◆ Use the `swlist` command. For example, to check for patch PHKL\_31500, type:

```
# swlist -l patch -a supersedes | egrep PHKL_31500
```

## 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>

# 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 `sfile` 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]
- SF Oracle RAC does not support the Veritas FlashSnap agent for Symmetrix (EMC TimeFinder) mapping functionality (package: `VRTSfas`).

## 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, `installsfrac`, 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 and configure SF Oracle RAC” on page 71.

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 Stop all applications that use CFS if they are not under VCS control.
  - Make sure that no processes are using the CFS mount point.  
`# fuser -c mount_point`
  - Stop any processes using a CFS mount point.  
`# fuser -ck mount_point`
- 5 Disable the LMX helper thread on all nodes in the cluster. Change the lmx tunable.  
`kctune -B lmx_update_enabled=0`
- 6 Install SF 5.0 MP1 for Oracle RAC:
  - Make sure that communication between systems exists.
  - Invoke the installmp script from one of your cluster nodes using the following command:  
`# ./installmp [-rsh]`
- 7 After the initial system checks and the requirement checks are complete, press Return to start upgrading the packages.
- 8 When installation is complete, note the locations of the summary, log, and response files indicated by the installer.
- 9 Restart all systems in the cluster.  
`# /usr/sbin/shutdown -r now`
- 10 Change to /opt/VRTS/install directory.
- 11 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 components” on page 83.
- 12 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*.
- 13 If you want to use Checkpoints, FlashSnap, or Storage Mapping, configure the repository database.  
See Chapter 20, “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*.

### Overview of upgrade tasks

You must perform the following tasks:

- [Preparing to upgrade SF Oracle RAC to 5.0 MP1\[HP\]](#)
- [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.

---

### Preparing to upgrade SF Oracle RAC to 5.0 MP1

Depending on the version of Oracle, do one of the following:

- See [“To upgrade SF Oracle RAC from a cluster having Oracle 9i instances”](#) on page 23.
- See [“To upgrade SF Oracle RAC from a cluster having Oracle 10g instances”](#) on page 26.

#### To upgrade SF Oracle RAC from a cluster having Oracle 9i instances

- 1 Perform the following tasks on each node in the cluster to stop the gsd processes:
  - Determine if the gsd processes are running.  

```
$ $ORACLE_HOME/bin/gsdctl stat
```
  - If the processes are running, stop the processes.  

```
$ $ORACLE_HOME/bin/gsdctl stop
```
- 2 If *only* the Oracle9i instances are under VCS control, you must take the Veritas High Availability Agent for Oracle resources offline. The Veritas agent for Oracle include Oracle and Netlsnr resources.
  - Log in as superuser.
  - On each node, take the Oracle and Netlsnr resources in the VCS configuration file (main.cf) offline.  

```
# hares -offline oracle_resource -sys node_name  
# hares -offline netlsnr_resource -sys node_name
```

For example:

```
# hares -offline VRT -sys galaxy
# hares -offline rac -sys galaxy
# hares -offline VRT -sys nebula
# hares -offline rac -sys nebula
# hares -offline LISTENER -sys galaxy
# hares -offline LISTENER -sys nebula
```

These commands stop the Oracle and Netlsnr resources running on the specified nodes.

- 3 Stop all applications using the CFS mounts not under VCS control.
  - Ensure that no processes are using the CFS mount point.

```
# fuser -c mount_point
```
  - Stop any processes using a CFS mount point.

```
# fuser -ck mount_point
```
- 4 Unmount any CFS file systems that are not under VCS control on all nodes.
  - Determine the file systems to unmount by checking the `/etc/mnttab` file.

For example:

```
# cat /etc/mnttab | grep vxfs | grep cluster
```

The output shows each line of the `/etc/mnttab` file that contains an entry for a VxFS file system mounted in the cluster mode.
  - By specifying the mount point for the file system, unmount each file system listed in the output:

```
# umount mount_point
```
- 5 Stop VCS to take the service groups on all nodes offline.
  - Log in as superuser on one cluster node.
  - Stop VCS on all nodes.

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

### To upgrade SF Oracle RAC from a cluster having Oracle 10g instances

- 1 If only the Oracle 10g instances are under VCS control, you must take the Oracle resources offline.
  - Log in as superuser.
  - On each node, take the Oracle resources in the VCS configuration file (`main.cf`) offline.

```
# hagrps -offline oracle_group -sys node_name
```

For example:

```
# /opt/VRTSvcs/bin/hagrps -offline Oracle1 -sys galaxy
# /opt/VRTSvcs/bin/hagrps -offline Oracle1 -sys nebula
```

These commands stop the Oracle resources under VCS control running on the specified nodes.

- On each node, take the CSSD resources in the VCS configuration file (main.cf) offline.

```
# hares -offline cssd_resource -sys node_name
```

For example:

```
# hares -offline cssd-resource1 -sys galaxy
```

```
# hares -offline cssd-resource1 -sys nebula
```

These commands stop the CSSD resources under VCS control running on the specified nodes.

- Verify the state of the Oracle and CVM service groups.

```
# /opt/VRTSvcs/bin/hagrp -state
```

#Group	Attribute	System	Value
Oracle1	State	galaxy	OFFLINE
Oracle1	State	nebula	OFFLINE
cvm	State	galaxy	ONLINE
cvm	State	nebula	ONLINE

## 2 Stop all applications using the CFS mounts not under VCS control.

- Ensure that no processes are using the CFS mount point.

```
# fuser -c mount_point
```

- Stop any processes using a CFS mount point.

```
# fuser -ck mount_point
```

## 3 Unmount any CFS file systems that are not under VCS control on all nodes.

- Determine the file systems to unmount by checking the /etc/mnttab file.

For example:

```
# cat /etc/mnttab | grep vxfs | grep cluster
```

The output shows each line of the /etc/mnttab file that contains an entry for a VxFS file system mounted in the cluster mode.

- By specifying the mount point for the file system, unmount each file system listed in the output:

```
# umount mount_point
```

## 4 Stop VCS to take the service groups on all nodes offline.

- Log in as superuser on one cluster node.

- Stop VCS on all nodes.

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

## 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 -F cdfs /dev/dsk/c0t6d0 /cdrom
```
- 4 Browse to the directory containing installer program.
- 5 Disable the LMX helper thread on all nodes in the cluster. Change the lmx tunable.  

```
kctune -B lmx_update_enabled=0
```
- 6 Install SF 5.0 MP1 for Oracle RAC:
  - Make sure that communication between systems exists.
  - Invoke the installmp script from one of your cluster nodes using the following command:  

```
# ./installmp [-rsh]
```
- 7 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 patches that it will install on each system

```
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.
```
- 8 Note the location of the summary and log files for future reference.
- 9 Restart each system in the cluster when prompted.  

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

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

If you want to upgrade to Oracle, refer to the Oracle documentation. You must relink Oracle to the Veritas libraries after upgrading to SF Oracle RAC 5.0 MP1.

Refer to the *Veritas Storage Foundation 5.0 for Oracle RAC Installation and Configuration guide* for instructions on relinking Oracle to the Veritas libraries.

## 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 750 /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 -F 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 Log in to *galaxy* and *nebula*, stop VEA if it is running.
 

```
# /opt/VRTSob/bin/vxsvctrl stop
```
- 10 Change to CD-ROM directory.
 

```
# cd /cdrom
```
- 11 Disable the LMX helper thread on all nodes in the cluster. Change the *lmx* tunable.

```
kctune -B lmx_update_enabled=0
```

**12 Start the installer.**

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

**13 Restart each system in the cluster when prompted.**

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

**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 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
```
- 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 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 41.
- 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*:
- ```
# /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*.
- 4 Stop VEA if it is running:
 

```

# /opt/VRTSob/bin/vxsvcctl 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 now

```
- 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 Verify CRS is offline before relinking.

- 9 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 45.

### 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
```

## 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  

```
$ gsctl 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
 

```
# hstop -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 now
```
  - 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
```
- 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 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 41.
  - 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/hagrps -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/vxsvcctrl status`  
# `/opt/VRTSob/bin/vxsvcctrl 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 now`
- 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 52.
- 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
```

## 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 750 /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
```

## Rolling back to SF Oracle RAC 5.0 from 5.0 MP1

You must remove the SF Oracle RAC 5.0 MP1 patch set if you want to roll back to SF Oracle RAC 5.0.

---

**Note:** Before uninstalling, verify the supported Oracle versions for 5.0.

---

The tasks involve:

- [Preparing to remove SF Oracle RAC to 5.0 MP1](#)
- [Removing SF Oracle RAC 5.0 MP1 patches](#)
- [Performing the post-uninstallation tasks](#)

## Preparing to remove SF Oracle RAC to 5.0 MP1

Depending on the version of Oracle, do one of the following:

- See [“To remove SF Oracle RAC 5.0 MP1 patches from a cluster having Oracle 9i instances”](#) on page 56.
- See [“To remove SF Oracle RAC 5.0 MP1 patches from a cluster having Oracle 10g instances”](#) on page 57.

### To remove SF Oracle RAC 5.0 MP1 patches from a cluster having Oracle 9i instances

- 1 Perform the following tasks on each node in the cluster to stop the gsd processes:
  - Determine if the gsd processes are running.  

```
$ $ORACLE_HOME/bin/gsdctl stat
```
  - If the processes are running, stop the processes.  

```
$ $ORACLE_HOME/bin/gsdctl stop
```
- 2 If *only* the Oracle9i instances are under VCS control, you must take the Veritas High Availability Agent for Oracle resources offline. The Veritas agent for Oracle include Oracle and Netlsnr resources.
  - Log in as superuser.
  - On each node, take the Oracle and Netlsnr resources in the VCS configuration file (main.cf) offline.  

```
# hares -offline oracle_resource -sys node_name  
# hares -offline netlsnr_resource -sys node_name
```

For example:  

```
# hares -offline VRT -sys galaxy  
# hares -offline rac -sys galaxy  
# hares -offline VRT -sys nebula  
# hares -offline rac -sys nebula  
# hares -offline LISTENER -sys galaxy  
# hares -offline LISTENER -sys nebula
```

These commands stop the Oracle and Netlsnr resources running on the specified nodes.
- 3 Because Oracle must be set up to run with the SF Oracle RAC components after the update, you must prevent the Oracle resources from coming online after the systems reboot. Use the following commands to make the cluster writable and to disable the Oracle resource:  

```
# haconf -makerw  
# hares -modify oracle_resource Enabled 0
```

Verify you have made this change:  

```
# hares -display oracle_resource | grep Enabled
```

If the change is made successfully, save the configuration:

- ```
# haconf -dump -makero
```
- 4 Stop all applications using the CFS mounts not under VCS control.
    - Ensure that no processes are using the CFS mount point.

```
# fuser -c mount_point
```
    - Stop any processes using a CFS mount point.

```
# fuser -ck mount_point
```
  - 5 Unmount any CFS file systems that are not under VCS control on all nodes.
    - Determine the file systems to unmount by checking the /etc/mnttab file.  
For example:

```
# cat /etc/mnttab | grep vxfs | grep cluster
```

The output shows each line of the /etc/mnttab file that contains an entry for a VxFS file system mounted in the cluster mode.
    - By specifying the mount point for the file system, unmount each file system listed in the output:

```
# umount mount_point
```
  - 6 Stop VCS to take the service groups on all nodes offline.
    - Log in as superuser on one cluster node.
    - Stop VCS on all nodes.

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

#### To remove SF Oracle RAC 5.0 MP1 patches from a cluster having Oracle 10g instances

- 1 If only the Oracle 10g instances are under VCS control, you must take the Oracle resources offline.
  - Log in as superuser.
  - On each node, take the Oracle resources in the VCS configuration file (main.cf) offline.

```
# hagr -offline oracle_group -sys node_name
```

For example:

```
# /opt/VRTSvcs/bin/hagr -offline Oracle1 -sys galaxy  
# /opt/VRTSvcs/bin/hagr -offline Oracle1 -sys nebula
```

These commands stop the Oracle resources under VCS control running on the specified nodes.
  - On each node, take the CSSD resources in the VCS configuration file (main.cf) offline.

```
# hares -offline cssd_resource -sys node_name
```

For example:

```
# hares -offline cssd-resource1 -sys galaxy  
# hares -offline cssd-resource1 -sys nebula
```

These commands stop the CSSD resources under VCS control running on the specified nodes.

- Verify the state of the Oracle and CVM service groups.

```
# /opt/VRTSvcs/bin/hagrp -state
#Group          Attribute      System      Value
Oracle1         State         galaxy     |OFFLINE|
Oracle1         State         nebula     |OFFLINE|
cvm             State         galaxy     |ONLINE|
cvm             State         nebula     |ONLINE|
```

- 2 Stop all applications using the CFS mounts not under VCS control.

- Ensure that no processes are using the CFS mount point.

```
# fuser -c mount_point
```

- Stop any processes using a CFS mount point.

```
# fuser -ck mount_point
```

- 3 Unmount any CFS file systems that are not under VCS control on all nodes.

- Determine the file systems to unmount by checking the `/etc/mnttab` file.

For example:

```
# cat /etc/mnttab | grep vxfs | grep cluster
```

The output shows each line of the `/etc/mnttab` file that contains an entry for a VxFS file system mounted in the cluster mode.

- By specifying the mount point for the file system, unmount each file system listed in the output:

```
# umount mount_point
```

- 4 Because Oracle must be set up to run with the SF Oracle RAC components after the update, you must prevent the Oracle service group from coming online after the systems reboot. Use the following commands to make the cluster writable and to disable the Oracle service group:

```
# haconf -makerw
```

```
# hagrp -disable oracle_group -sys node_name
```

For example:

```
# /opt/VRTSvcs/bin/hagrp -disable Oracle1 -sys galaxy
```

If the change is made successfully, save the configuration:

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

- 5 Stop VCS to take the service groups on all nodes offline.

- Log in as superuser on one cluster node.

- Stop VCS on all nodes.

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

- On all the nodes, stop VxFEN and ODM processes:

```
# /sbin/init.d/vxfen stop
```

```
# /sbin/init.d/odm stop
```

## Removing SF Oracle RAC 5.0 MP1 patches

After preparing to remove the SF Oracle RAC 5.0 MP1 patches, you can now remove the patches.

---

**Note:** Using the `swremove` command to uninstall patches related to Cluster Management Console `PVCO_03686` (`VRTScmccc`) or `PVCO_03689` (`VRTScmcs`) does not restore those software selections to their original 5.0 versions. If the cluster is configured with either the Cluster Management Console or the cluster connector, neither component functions after the attempted reversion. The workaround is to uninstall and then reinstall the CMC cluster component packages (`VRTScmccc` and/or `VRTScmcs`) using the original version 5.0 installation media. [865666, 865627]

---

### To remove SF Oracle RAC 5.0 MP1 patches

- 1 Log in as superuser.
- 2 On all the nodes, remove all the patches using the `swremove` command.

```
# swremove -x autoreboot=true -x enforce_dependencies=false  
patch_name1, patch_name2 ...
```

---

**Caution:** You must always install or remove the patches `PHKL_35178`, `PHCO_35179`, and `PVCO_03696` together using a single `swremove` command. Not installing or removing them together can result in data loss.

---

For example, to remove all the 5.0 MP1 patches:

```
# swremove -x autoreboot=true \  
-x enforce_dependencies=false \  
PHCO_35127 PHCO_35179 PHCO_35124 PHCO_35126 \  
PHCO_35125 PHCO_35180 PHCO_35213 PHCO_35212 \  
PHCO_35214 PHCO_35215 PHCO_35216 PHCO_35217 \  
PHCO_35301 PHCO_35332 PHCO_35375 PHCO_35355 \  
PHCO_35354 PHCO_35425 PVCO_03671 PVCO_03672 \  
PVCO_03676 PVCO_03678 PVCO_03679 PVCO_03680 \  
PVCO_03686 PVCO_03689 PVCO_03690 PVCO_03694 \  
PVCO_03696 PVCO_03697 PVCO_03698 PVKL_03688 \  
PHKL_35178 PHKL_35284 PHKL_35305 PHKL_35312 \  
PHNE_35413 PHKL_35424 PHNE_35783
```

## Performing the post-uninstallation tasks

You must relink the Veritas libraries and enable Oracle to use Oracle with SF Oracle RAC 5.0.

### To perform the post-uninstallation tasks

- 1 Relink Veritas libraries to Oracle.

Refer to *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* for 5.0.

- 2 If Oracle is under VCS control, do the following depending on the version of Oracle.

#### For Oracle9i

- Enable the Oracle resource.

```
# haconf -makerw
# hares -modify oracle_resource Enabled 1
# haconf -dump -makero
```

- Bring the Oracle resource online.

```
# hares -online oracle_resource -sys node_name
```

#### For Oracle 10g

- Enable the Oracle service group.

```
# haconf -makerw
# hagrps -enable oracle_group -sys node_name
# haconf -dump -makero
```

- Bring the Oracle service group online.

```
# hares -online oracle_resource -sys node_name
```

## 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.

### 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), "lan0" 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 re commended that you retry the configuration
assistants at this time. Not successf ully running any
"Recommended" assistants means your system will not be correctl
y configured.
1. Check the Details panel on the Configuration Assistant Screen
to see the erro rs 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.

## The vxfsentsthdw command requires "+ +" in /.rhosts file

If you use remsh for communication between the nodes including the one on which you run the command, then the vxfsentsthdw requires that the /.rhosts file must contain a "+ +" character in the first line.

## 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>

## 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 now
```

## Accessing version 7 file system when removing VxFS

A disk layout version 7 file system created with VxFS 5.0 software is not accessible if you remove the VxFS 5.0 software and revert the system to VxFS 3.5 or 4.1. In addition, if a disk layout version 7 file system exists on a boot disk and you remove the VxFS 5.0 software, the host does not reboot successfully and remains at the `bcheckrc` prompt. To reboot the host successfully, first edit the `fstab` file and comment out entries for disk layout Version 7 file system.

## 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 537 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	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.  Example:  Device@galaxy = {lan0=0, lan1=1, lan2=2} Device@nebula = {lan0=0, lan1=1, lan2=2}
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`

## CVM disks get stuck in an error state

CVM disks might get into an error state in some scenarios. For instance, the disk might get into an error state after a reboot of the cluster nodes. It might also happen when a disk group is deported and then imported.

**Workaround:** Perform the following steps to bring the CVM disks online.

- 1 Run the following command on any of the cluster nodes to stop VCS.  

```
# /opt/VRTSvcs/bin/hastop -all
```
- 2 Run the following commands on each cluster node to clear the error state of the disks.  

```
# vxdisk -a online  
# vxdisk scandisks
```
- 3 Create a file that has all the VxVM disk names with the complete path. For example: `/dev/rdisk/c4t1d2`.
- 4 Check for SCSI keys on the VxVM devices.  

```
# vxfenadm -g all -f disk_filename
```

where *disk\_filename* is a file that contains all of the fully qualified path to each disk.
- 5 Register with the disks in the *disk\_filename*.  

```
# vxfenadm -m -k <key> -f disk_filename
```
- 6 Remove the SCSI keys on the VxVM devices:  

```
# vxfenadm -p -k <key> -v victim_node_key -f disk_filename
```

where *victim\_node\_key* is the SCSI key listed while running the command in [step 4](#).  
Repeat the step till all the SCSI keys are cleared.
- 7 Remove the key registered in [step 5](#).  

```
# vxfenadm -x -k <key> -f disk_filename
```
- 8 Run the following command on each of the cluster nodes to start VCS.  

```
# /opt/VRTSvcs/bin/hastart
```

## Simultaneous startup and shutdown of Oracle processes

Continuous startup and shutdown of the database incidents at the same time on all the nodes in an eight node SF Oracle RAC cluster can cause the cluster to hang.

**Workaround:** Do not startup and shutdown repeatedly at the same time on all the eight nodes in an SF Oracle RAC cluster. [See Etrack Incident 856071]

## 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 priority 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
644286	IPC errors while running OCE 10gr2
635646	Linkrac fails when you apply Oracle patch set and execute the patch
626451	installsfrac does not bring up GCO configuration
595453	IPC errors while running OCE 10gr2 on HP-UX PA
583747	In the PrivNIC agent section of the Installation and Configuration Guide, the number of links to be used is not mentioned correctly.

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

Incident number	Description
596047	vcsmmconfig -c hangs if the max slave parameter value differs from running cluster
595651	vcsmmdebug -s results in segmentation fault
587999	lmx unload may panic the machine after heavy lmx buffer transfers

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

Incident number	Description
579175	Oracle instance hangs during peak load
568762	IPC throws "NCON" messages in Oracle alert logs
544172	Utility to get the Oracle IPC version number at the time of configuration.
538897	vcsmm panics when process registers as slave with two primaries
535866	Port 'o' does not join after the reboot since module fails to load
534950	Fencing shuts down the larger cluster after a split brain
525049	vcsmm may panic the machine due to bug in its kernel logging
520819	vcsmm can miss a wake up and this may result in oracle hang
509099	vcsmm unconfiguration fails with "Device busy" error
506182	Add support for CRS online/offline through VCS

## 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
<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

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

Document Title	File Name
<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 are required for the following topics in 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:

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

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

```
# vxdg -g vxfencoorddg adddisk c2t1d0
# vxdg -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:

```
# vxvg init vxfencoorddg c1t1d0
```

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

```
# vxvg -g vxfencoorddg adddisk c2t1d0
```

```
# vxvg -g vxfencoorddg adddisk c3t1d0
```

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

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

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

## Setting permissions for database administration

The *Storage Foundation for Oracle RAC 5.0 Installation and Configuration Guide* contains text that is misleading in “Setting permissions for database administration,” on page 104. [See Etrack Incident 989674]

Replace:

After installing SF Oracle RAC, the default settings allows only the superuser to access the `/opt/VRTSdbed` folder. If you want the database administrators (DBAs) to access SF Oracle RAC components, then you must set the required permissions. You can skip setting the database administration permission and advance to configuring the cluster volume manager.

With:

After SF Oracle RAC is installed, the default settings allow only the superuser to access the `/opt/VRTSdbed` folder.

- If you already have an Oracle user and group, and you want database administrators (DBAs) to access SF Oracle RAC components, you must set the required permissions.
- If you do not have an Oracle user and group, defer setting the database administration permission and advance to configuring the cluster volume manager. You can configure the database permissions when setting up the database repository.

## 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 256.

- 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

## Preparing for global clustering

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. [See Etrack incident 609566]

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.

## 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
```

## Using 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”

## Using 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]

## Using the `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 396, 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`.

## Modifying the `csgnic` definition in `main.cf` file

The sample `main.cf` files for the clusters that use the GCO option and for the clusters that are centrally managed using Cluster Management Console contain the following `csgnic` definition.

```
NIC csgnic (  
    Device = lan0  
)
```

With this definition, the csgnic resource may fault intermittently. To overcome this issue, the NetworkHosts parameter must be provided. [See Etrack incident 838660]

This correction also applies to the section “Modifying the global clustering configuration using the main.cf” on page 304 of the *Storage Foundation for Oracle RAC 5.0 Installation and Configuration Guide*.

Modify the csgnic definition to add NetworkHosts attribute.

The sample configuration must be as follows:

```
NIC csgnic (  
    Device = lan0  
    NetworkHosts = { "10.10.12.2", "10.10.12.3" }  
)
```