

# Veritas Storage Foundation™ for Oracle® RAC Installation and Configuration Guide

Solaris

5.1 Service Pack 1

# Veritas Storage Foundation™ for Oracle RAC Installation and Configuration Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 5.1 SP1

Document version: 5.1SP1.1

## Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

Symantec, the Symantec logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate 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.

THE 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 as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation  
350 Ellis Street  
Mountain View, CA 94043  
<http://www.symantec.com>

# Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our Web site at the following URL:

[www.symantec.com/business/support/index.jsp](http://www.symantec.com/business/support/index.jsp)

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

## Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

[www.symantec.com/business/support/contact\\_techsupp\\_static.jsp](http://www.symantec.com/business/support/contact_techsupp_static.jsp)

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

## Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

[www.symantec.com/business/support/](http://www.symantec.com/business/support/)

## Customer service

Customer service information is available at the following URL:

[www.symantec.com/business/support/](http://www.symantec.com/business/support/)

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

## Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

|                                 |                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------|
| Asia-Pacific and Japan          | <a href="mailto:customercare_apac@symantec.com">customercare_apac@symantec.com</a> |
| Europe, Middle-East, and Africa | <a href="mailto:semea@symantec.com">semea@symantec.com</a>                         |
| North America and Latin America | <a href="mailto:supportsolutions@symantec.com">supportsolutions@symantec.com</a>   |

## Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec Web site.

<https://sort.symantec.com/documents>

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

[docs@symantec.com](mailto:docs@symantec.com)

## About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

# Contents

|                                                                                       |    |
|---------------------------------------------------------------------------------------|----|
| Technical Support .....                                                               | 4  |
| Section 1      Installation overview and planning .....                               | 25 |
| Chapter 1      Introducing Veritas Storage Foundation for Oracle<br>RAC .....         | 27 |
| About Veritas Storage Foundation for Oracle RAC .....                                 | 27 |
| Benefits of SF Oracle RAC .....                                                       | 27 |
| About basic SF Oracle RAC components .....                                            | 29 |
| About SF Oracle RAC features .....                                                    | 30 |
| About I/O fencing .....                                                               | 31 |
| Symantec Operations Readiness Tools .....                                             | 34 |
| About VCS notifications .....                                                         | 35 |
| About campus clusters .....                                                           | 35 |
| About global clusters .....                                                           | 35 |
| About Veritas Volume Replicator .....                                                 | 35 |
| About database management using SF Oracle RAC .....                                   | 35 |
| About SF Oracle RAC optional components .....                                         | 36 |
| About Symantec Product Authentication Service (AT) .....                              | 36 |
| Veritas Operations Manager .....                                                      | 36 |
| About Cluster Manager (Java Console) .....                                            | 37 |
| SF Oracle RAC cluster setup models .....                                              | 37 |
| Typical configuration of four-node SF Oracle RAC cluster .....                        | 38 |
| Typical configuration of SF Oracle RAC clusters in secure<br>mode .....               | 39 |
| Typical configuration of VOM-managed SF Oracle RAC<br>clusters .....                  | 40 |
| Typical configuration of SF Oracle RAC campus clusters for<br>disaster recovery ..... | 41 |
| Typical configuration of SF Oracle RAC global clusters for<br>disaster recovery ..... | 43 |

|           |                                                                                    |    |
|-----------|------------------------------------------------------------------------------------|----|
| Chapter 2 | System requirements .....                                                          | 45 |
|           | Important preinstallation information .....                                        | 45 |
|           | Hardware requirements .....                                                        | 46 |
|           | I/O fencing requirements .....                                                     | 48 |
|           | Coordinator disk requirements for I/O fencing .....                                | 48 |
|           | CP server requirements .....                                                       | 48 |
|           | Supported Solaris operating systems .....                                          | 51 |
|           | Required operating system patches .....                                            | 52 |
|           | Supported database software .....                                                  | 53 |
|           | Veritas File System requirements .....                                             | 54 |
|           | Supported replication technologies for global clusters .....                       | 55 |
| Chapter 3 | Planning to install SF Oracle RAC .....                                            | 57 |
|           | Planning your network configuration .....                                          | 57 |
|           | Planning the public network configuration for Oracle RAC .....                     | 58 |
|           | Planning the private network configuration for Oracle RAC .....                    | 58 |
|           | Planning the storage .....                                                         | 61 |
|           | Planning the storage for SF Oracle RAC .....                                       | 62 |
|           | Planning the storage for Oracle RAC .....                                          | 62 |
|           | Planning volume layout .....                                                       | 66 |
|           | Planning file system design .....                                                  | 67 |
|           | About planning to configure I/O fencing .....                                      | 67 |
|           | Typical SF Oracle RAC cluster configuration with disk-based I/O<br>fencing .....   | 70 |
|           | Typical SF Oracle RAC cluster configuration with server-based<br>I/O fencing ..... | 71 |
|           | Recommended CP server configurations .....                                         | 71 |
|           | Planning for cluster management .....                                              | 74 |
|           | Planning for disaster recovery .....                                               | 75 |
|           | Planning a campus cluster setup .....                                              | 75 |
|           | Planning a global cluster setup .....                                              | 76 |
|           | Data replication considerations .....                                              | 76 |
|           | Planning for Oracle VM Server for SPARC environments .....                         | 77 |
| Chapter 4 | Licensing SF Oracle RAC .....                                                      | 79 |
|           | About Veritas product licensing .....                                              | 79 |
|           | About SF Oracle RAC licenses .....                                                 | 80 |
|           | Setting or changing the product level for keyless licensing .....                  | 82 |
|           | Installing Veritas product license keys .....                                      | 83 |



|           |                                                                                    |     |
|-----------|------------------------------------------------------------------------------------|-----|
| Section 2 | Installation and configuration of SF Oracle RAC .....                              | 85  |
| Chapter 5 | Preparing to install SF Oracle RAC .....                                           | 87  |
|           | Setting the umask before installation .....                                        | 87  |
|           | Synchronizing time settings on cluster nodes .....                                 | 87  |
|           | Mounting the product disc .....                                                    | 88  |
|           | Setting up shared storage .....                                                    | 88  |
|           | Setting the environment variables .....                                            | 89  |
|           | Optimizing LLT media speed settings on private NICs .....                          | 89  |
|           | Guidelines for setting the media speed of the LLT interconnects .....              | 89  |
|           | Verifying the systems before installation .....                                    | 90  |
| Chapter 6 | Installing SF Oracle RAC .....                                                     | 91  |
|           | About installing SF Oracle RAC .....                                               | 91  |
|           | About installation and configuration methods .....                                 | 92  |
|           | Installing SF Oracle RAC using the Veritas script-based installation program ..... | 93  |
|           | Installing SF Oracle RAC using the Veritas Web-based installation program .....    | 96  |
|           | Before using the Veritas Web-based installer .....                                 | 97  |
|           | Starting the Veritas Web-based installer .....                                     | 97  |
|           | Installing products with the Veritas Web-based installer .....                     | 98  |
|           | Installing SF Oracle RAC using Solaris JumpStart .....                             | 99  |
|           | Task overview for SF Oracle RAC installation using JumpStart .....                 | 100 |
|           | Preparing the JumpStart installation resources .....                               | 100 |
|           | Installing and configuring SF Oracle RAC using JumpStart .....                     | 108 |
|           | Sample JumpStart finish file (basic installation) .....                            | 108 |
|           | Sample JumpStart finish file (for root encapsulation) .....                        | 113 |
|           | Installing SF Oracle RAC on an alternate root .....                                | 115 |
| Chapter 7 | Preparing to configure SF Oracle RAC .....                                         | 119 |
|           | Preparing to configure the clusters in secure mode .....                           | 119 |
|           | Installing the root broker for the security infrastructure .....                   | 123 |
|           | Creating authentication broker accounts on root broker system .....                | 124 |
|           | Creating encrypted files for the security infrastructure .....                     | 125 |

|                                                                         |     |
|-------------------------------------------------------------------------|-----|
| Preparing the installation system for the security infrastructure ..... | 127 |
| Setting up the CP server .....                                          | 128 |
| Planning your CP server setup .....                                     | 129 |
| Installing the CP server using the installer .....                      | 130 |
| Configuring the CP server cluster in secure mode .....                  | 131 |
| Setting up shared storage for the CP server database .....              | 132 |
| Configuring the CP server using the configuration utility .....         | 133 |
| Configuring the CP server manually .....                                | 141 |
| Verifying the CP server configuration .....                             | 143 |

## Chapter 8      Configuring SF Oracle RAC ..... 145

|                                                                                 |     |
|---------------------------------------------------------------------------------|-----|
| About configuring SF Oracle RAC .....                                           | 145 |
| Configuring the SF Oracle RAC components using the script-based installer ..... | 146 |
| Configuring the SF Oracle RAC cluster .....                                     | 148 |
| Creation of SF Oracle RAC configuration files .....                             | 161 |
| Stopping and starting SF Oracle RAC processes .....                             | 161 |
| Configuring SF Oracle RAC using the Web-based installer .....                   | 162 |

## Chapter 9      Configuring SF Oracle RAC clusters for data integrity ..... 169

|                                                                                                             |     |
|-------------------------------------------------------------------------------------------------------------|-----|
| Setting up disk-based I/O fencing using installsfrac program .....                                          | 169 |
| Initializing disks as VxVM disks .....                                                                      | 169 |
| Identifying disks to use as coordinator disks .....                                                         | 170 |
| Checking shared disks for I/O fencing .....                                                                 | 170 |
| Configuring disk-based I/O fencing using installsfrac program .....                                         | 175 |
| Setting up disk-based I/O fencing manually .....                                                            | 178 |
| Setting up coordinator disk groups .....                                                                    | 178 |
| Creating I/O fencing configuration files .....                                                              | 179 |
| Modifying VCS configuration to use I/O fencing .....                                                        | 180 |
| Starting SF Oracle RAC on all nodes .....                                                                   | 181 |
| Verifying I/O fencing configuration .....                                                                   | 183 |
| Setting up server-based I/O fencing using installsfrac program .....                                        | 184 |
| Verifying the security configuration on the SF Oracle RAC cluster to use CP server coordination point ..... | 184 |
| Configuring server-based I/O fencing using the installsfrac program .....                                   | 186 |
| Setting up server-based I/O fencing manually .....                                                          | 195 |
| Preparing the CP servers manually for use by the SF Oracle RAC cluster .....                                | 195 |

|            |                                                                                   |     |
|------------|-----------------------------------------------------------------------------------|-----|
|            | Configuring server-based fencing on the SF Oracle RAC cluster manually .....      | 199 |
|            | Configuring Coordination Point agent to monitor coordination points .....         | 203 |
|            | Verifying server-based I/O fencing configuration .....                            | 205 |
| Chapter 10 | Performing post-installation and configuration tasks .....                        | 207 |
|            | Installing language packages .....                                                | 207 |
|            | About enabling LDAP authentication for clusters that run in secure mode .....     | 208 |
|            | Enabling LDAP authentication for clusters that run in secure mode .....           | 210 |
|            | Configuring Veritas Volume Replicator .....                                       | 216 |
|            | Verifying SF Oracle RAC installation using VCS configuration file .....           | 217 |
|            | Verifying LLT, GAB, and cluster operation .....                                   | 217 |
|            | Verifying LLT .....                                                               | 218 |
|            | Verifying GAB .....                                                               | 220 |
|            | Verifying the cluster .....                                                       | 221 |
|            | Verifying the cluster nodes .....                                                 | 222 |
|            | Running VRTSexplorer to diagnose issues in the cluster .....                      | 225 |
| Section 3  | Upgrade of SF Oracle RAC .....                                                    | 227 |
| Chapter 11 | Preparing to upgrade to SF Oracle RAC 5.1 SP1 .....                               | 229 |
|            | About upgrading SF Oracle RAC .....                                               | 229 |
|            | About types of upgrade .....                                                      | 230 |
|            | Supported upgrade paths .....                                                     | 230 |
|            | Preparing to upgrade to SF Oracle RAC 5.1 SP1 .....                               | 231 |
|            | Pre-upgrade tasks for migrating the SFDB repository database .....                | 233 |
| Chapter 12 | Performing a full upgrade to SF Oracle RAC 5.1 SP1 .....                          | 235 |
|            | About full upgrades .....                                                         | 235 |
|            | Upgrading to SF Oracle RAC 5.1 SP1 .....                                          | 236 |
|            | Upgrading SF Oracle RAC using the Veritas script-based installation program ..... | 238 |
|            | Upgrading SF Oracle RAC using a response file .....                               | 240 |

|            |                                                                               |     |
|------------|-------------------------------------------------------------------------------|-----|
| Chapter 13 | Performing a phased upgrade to SF Oracle RAC 5.1                              |     |
|            | SP1 .....                                                                     | 245 |
|            | About phased upgrade .....                                                    | 245 |
|            | Performing phased upgrade of SF Oracle RAC from version 5.0 .....             | 246 |
|            | Step 1: Performing pre-upgrade tasks on the first half of the cluster .....   | 247 |
|            | Step 2: Upgrading the first half of the cluster .....                         | 250 |
|            | Step 3: Performing pre-upgrade tasks on the second half of the cluster .....  | 251 |
|            | Step 4: Performing post-upgrade tasks on the first half of the cluster .....  | 252 |
|            | Step 5: Upgrading the second half of the cluster .....                        | 254 |
|            | Step 6: Performing post-upgrade tasks on the second half of the cluster ..... | 254 |
| Chapter 14 | Performing a rolling upgrade to SF Oracle RAC 5.1                             |     |
|            | SP1 .....                                                                     | 257 |
|            | About rolling upgrades .....                                                  | 257 |
|            | Prerequisites for a rolling upgrade .....                                     | 257 |
|            | Preparing to perform a rolling upgrade to SF Oracle RAC 5.1                   |     |
|            | SP1 .....                                                                     | 258 |
|            | Performing a rolling upgrade using the installer .....                        | 260 |
|            | Performing a rolling upgrade on kernel : phase 1 .....                        | 260 |
|            | Performing a rolling upgrade on non-kernel : phase 2 .....                    | 262 |
| Chapter 15 | Upgrading to SF Oracle RAC 5.1 SP1 using Live Upgrade                         |     |
|            | About Live Upgrade .....                                                      | 263 |
|            | About Live Upgrade in a Veritas Volume Replicator (VVR) environment .....     | 264 |
|            | Supported upgrade paths for Live Upgrade .....                                | 265 |
|            | Before you upgrade SF Oracle RAC using Solaris Live Upgrade .....             | 266 |
|            | Upgrading the operating system and SF Oracle RAC using Live Upgrade .....     | 267 |
|            | Upgrading SF Oracle RAC only using Live Upgrade .....                         | 268 |
|            | Upgrading Solaris only using Live Upgrade .....                               | 268 |
|            | Removing and reinstalling SF Oracle RAC using the installer .....             | 269 |
|            | Creating a new boot environment on the alternate boot disk .....              | 270 |
|            | Upgrading SF Oracle RAC manually .....                                        | 272 |
|            | Upgrading SF Oracle RAC using the installer for a Live Upgrade .....          | 277 |

|            |                                                                                           |     |
|------------|-------------------------------------------------------------------------------------------|-----|
|            | Completing the Live Upgrade .....                                                         | 278 |
|            | Verifying Live Upgrade of SF Oracle RAC .....                                             | 281 |
|            | Reverting to the primary boot environment .....                                           | 282 |
| Chapter 16 | Upgrading from Storage Foundation products to<br>SF Oracle RAC 5.1 SP1 .....              | 283 |
|            | About upgrading from Storage Foundation products to SF Oracle RAC<br>5.1 SP1 .....        | 283 |
|            | Upgrading from Storage Foundation products to SF Oracle RAC 5.1<br>SP1 .....              | 284 |
| Chapter 17 | Migrating from single instance Storage Foundation<br>for Oracle HA to SF Oracle RAC ..... | 287 |
|            | Migration overview .....                                                                  | 287 |
|            | Sample configuration before and after migration .....                                     | 288 |
|            | Migration requirements .....                                                              | 289 |
|            | Before you migrate .....                                                                  | 289 |
|            | Migrating to SF Oracle RAC 5.1 SP1 .....                                                  | 290 |
|            | Migrating Storage Foundation for Oracle HA to SF Oracle<br>RAC .....                      | 290 |
|            | Migrating a single instance Oracle database to Oracle RAC<br>database .....               | 291 |
|            | Completing post-migration tasks .....                                                     | 294 |
|            | Sample configuration files .....                                                          | 297 |
|            | VCS configuration file for Storage Foundation for Oracle<br>HA .....                      | 297 |
|            | Oracle initialization parameter file for Storage Foundation for<br>Oracle HA .....        | 298 |
|            | tnsnames.ora for Storage Foundation for Oracle HA .....                                   | 299 |
|            | listener.ora for Storage Foundation for Oracle HA .....                                   | 299 |
| Chapter 18 | Performing post-upgrade tasks .....                                                       | 301 |
|            | Relinking Oracle RAC libraries with the SF Oracle RAC libraries .....                     | 301 |
|            | Setting or changing the product license level .....                                       | 304 |
|            | Changing permissions for Storage Foundation for Databases .....                           | 304 |
|            | Upgrading disk layout versions .....                                                      | 305 |
|            | Upgrading VxVM disk group version .....                                                   | 305 |
|            | Post upgrade tasks for migrating the SFDB repository database .....                       | 305 |
|            | Migrating from a 5.0 repository database to 5.1 Service Pack<br>1 .....                   | 306 |

|                                                                                      |     |
|--------------------------------------------------------------------------------------|-----|
| Migrating from a 4.x repository database to 5.1 Service Pack 1 .....                 | 309 |
| After upgrading from 5.0.x and before migrating SFDB .....                           | 311 |
| Upgrading the Array Support Library .....                                            | 311 |
| Adding JBOD support for storage arrays for which there is not an ASL available ..... | 312 |
| Unsuppressing DMP for EMC PowerPath disks .....                                      | 313 |
| Converting from QuickLog to Multi-Volume support .....                               | 321 |
| Verifying the cluster .....                                                          | 323 |

## Section 4 Installation and upgrade of Oracle RAC ..... 325

|                                                                                                     |     |
|-----------------------------------------------------------------------------------------------------|-----|
| Chapter 19 Before installing Oracle RAC .....                                                       | 327 |
| About preparing to install Oracle RAC .....                                                         | 327 |
| Preparing to install Oracle RAC using the SF Oracle RAC installer or manually .....                 | 328 |
| Identifying the public virtual IP addresses for use by Oracle .....                                 | 329 |
| Setting the kernel parameters .....                                                                 | 330 |
| Verifying the user "nobody" exists .....                                                            | 330 |
| Launching the SF Oracle RAC installer .....                                                         | 330 |
| Creating users and groups for Oracle RAC .....                                                      | 331 |
| Creating storage for OCR and voting disk .....                                                      | 334 |
| Configuring private IP addresses for Oracle RAC .....                                               | 347 |
| Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually ..... | 360 |
| Setting up user equivalence .....                                                                   | 376 |
| Verifying whether the Veritas Membership library is linked to Oracle libraries .....                | 376 |
| Verifying the systems for Oracle RAC installation .....                                             | 377 |

|                                                                                              |     |
|----------------------------------------------------------------------------------------------|-----|
| Chapter 20 Installing Oracle RAC .....                                                       | 381 |
| About installing Oracle RAC .....                                                            | 381 |
| Installing the Oracle Clusterware/Grid Infrastructure software .....                         | 382 |
| Installing Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC installer .....    | 383 |
| Installing Oracle Clusterware/Grid Infrastructure using the Oracle Universal Installer ..... | 386 |
| Configuring LLT links in the GPnP profile .....                                              | 387 |
| Installing the Oracle RAC database software .....                                            | 388 |

|            |                                                                                               |     |
|------------|-----------------------------------------------------------------------------------------------|-----|
|            | Installing the Oracle RAC database using the SF Oracle RAC installer .....                    | 389 |
|            | Installing the Oracle RAC database using the Oracle Universal Installer .....                 | 391 |
|            | Verifying the Oracle Clusterware/Grid Infrastructure and database installation .....          | 392 |
|            | Node numbering discrepancies in Oracle RAC 11g Release 2 .....                                | 393 |
| Chapter 21 | Performing Oracle RAC post-installation tasks .....                                           | 395 |
|            | Adding Oracle RAC patches or patchsets .....                                                  | 395 |
|            | Configuring the CSSD resource .....                                                           | 396 |
|            | Configuring the CSSD resource using the SF Oracle RAC installer .....                         | 396 |
|            | Configuring the CSSD resource manually .....                                                  | 398 |
|            | Preventing automatic startup of Oracle Clusterware/Grid Infrastructure .....                  | 400 |
|            | Relinking the SF Oracle RAC libraries with Oracle RAC .....                                   | 400 |
|            | Relinking the SF Oracle RAC libraries with Oracle RAC using the SF Oracle RAC installer ..... | 401 |
|            | Relinking SF Oracle RAC libraries with Oracle RAC manually .....                              | 402 |
|            | Creating the Oracle RAC database .....                                                        | 406 |
|            | Adding Oracle UDP IPC private IP addresses to the Oracle initialization parameter file .....  | 407 |
|            | Preventing automatic database startup .....                                                   | 409 |
|            | Configuring VCS service groups for Oracle RAC .....                                           | 410 |
|            | Using the CLI to configure VCS service groups for Oracle RAC .....                            | 413 |
|            | Location of VCS log files .....                                                               | 415 |
|            | Removing permissions for communication .....                                                  | 415 |
|            | Configuring the SFDB repository database after installation .....                             | 415 |
| Chapter 22 | Upgrading Oracle RAC .....                                                                    | 417 |
|            | Supported upgrade paths .....                                                                 | 417 |
|            | Preparing to upgrade Oracle RAC .....                                                         | 418 |
|            | Preparing to upgrade from Oracle RAC 10g or Oracle RAC 11g .....                              | 418 |
|            | Preparing to upgrade from Oracle RAC 9i .....                                                 | 419 |
|            | Upgrading Oracle RAC binaries .....                                                           | 420 |
|            | Migrating the Oracle RAC database .....                                                       | 421 |
|            | Performing post-upgrade tasks .....                                                           | 421 |

|            |                                                                          |     |
|------------|--------------------------------------------------------------------------|-----|
| Section 5  | Installation of SF Oracle RAC and Oracle RAC using a response file ..... | 423 |
| Chapter 23 | About response files .....                                               | 425 |
|            | About response files .....                                               | 425 |
|            | Modular deployments using response files .....                           | 427 |
|            | About the installation simulator .....                                   | 428 |
|            | Response file syntax .....                                               | 428 |
|            | Guidelines for creating the SF Oracle RAC response file .....            | 429 |
|            | Installation scenarios for response files .....                          | 434 |
| Chapter 24 | Installing and configuring SF Oracle RAC using a response file .....     | 435 |
|            | Installing and configuring SF Oracle RAC .....                           | 435 |
|            | Sample response file for installing and configuring SF Oracle RAC .....  | 437 |
| Chapter 25 | Configuring I/O fencing using a response file .....                      | 439 |
|            | Configuring I/O fencing using response files .....                       | 439 |
|            | Sample response file for configuring disk-based I/O fencing .....        | 440 |
|            | Sample response file for configuring server-based I/O fencing .....      | 440 |
| Chapter 26 | Installing Oracle RAC using a response file .....                        | 443 |
|            | About installing Oracle RAC using response files .....                   | 443 |
|            | Before you install .....                                                 | 443 |
|            | Installing Oracle RAC .....                                              | 445 |
|            | Sample response file for installing Oracle RAC .....                     | 446 |
| Chapter 27 | Installing SF Oracle RAC and Oracle RAC using a response file .....      | 449 |
|            | About installing SF Oracle RAC and Oracle RAC using response files ..... | 449 |
|            | Information required in the SF Oracle RAC response file .....            | 450 |
|            | Before you install .....                                                 | 451 |
|            | Installing SF Oracle RAC and Oracle RAC .....                            | 452 |
|            | Sample response file for installing SF Oracle RAC and Oracle RAC .....   | 454 |



|            |                                                                                                |     |
|------------|------------------------------------------------------------------------------------------------|-----|
| Chapter 28 | Response file variable definitions .....                                                       | 457 |
|            | Response file variables for installing SF Oracle RAC .....                                     | 457 |
|            | Response file variables to configure SF Oracle RAC .....                                       | 460 |
|            | Response file variables to configure disk-based I/O fencing .....                              | 469 |
|            | Response file variables to configure server-based I/O fencing .....                            | 471 |
|            | Response file variable definitions for Oracle RAC .....                                        | 473 |
| Section 6  | Adding or removing nodes from an SF<br>Oracle RAC cluster .....                                | 491 |
| Chapter 29 | Adding a node to SF Oracle RAC clusters .....                                                  | 493 |
|            | About adding a node to a cluster .....                                                         | 493 |
|            | Before adding a node to a cluster .....                                                        | 494 |
|            | Meeting hardware and software requirements .....                                               | 494 |
|            | Setting up the hardware .....                                                                  | 494 |
|            | Preparing to add a node to a cluster .....                                                     | 496 |
|            | Adding a node to a cluster .....                                                               | 496 |
|            | Adding a node to a cluster using the SF Oracle RAC<br>installer .....                          | 496 |
|            | Adding the node to a cluster manually .....                                                    | 500 |
|            | Preparing the new node for installing Oracle RAC .....                                         | 513 |
|            | Preparing the new nodes for installing Oracle RAC using the SF<br>Oracle RAC installer .....   | 513 |
|            | Preparing the new node manually for installing Oracle RAC .....                                | 521 |
|            | Adding the new node to Oracle RAC .....                                                        | 533 |
|            | Updating the Storage Foundation for Databases (SFDB) repository<br>after adding a node .....   | 534 |
|            | Sample configuration file for adding a node to the cluster .....                               | 534 |
| Chapter 30 | Removing a node from SF Oracle RAC clusters .....                                              | 539 |
|            | About removing a node from a cluster .....                                                     | 539 |
|            | Removing a node from a cluster .....                                                           | 540 |
|            | Modifying the VCS configuration files on existing nodes .....                                  | 541 |
|            | Editing the /etc/llthosts file .....                                                           | 541 |
|            | Editing the /etc/gabtab file .....                                                             | 542 |
|            | Modifying the VCS configuration to remove the node .....                                       | 542 |
|            | Removing the node configuration from the CP server .....                                       | 543 |
|            | Removing security credentials from the leaving node .....                                      | 544 |
|            | Updating the Storage Foundation for Databases (SFDB) repository<br>after removing a node ..... | 545 |

|            |                                                                         |     |
|------------|-------------------------------------------------------------------------|-----|
|            | Sample configuration file for removing a node from the cluster .....    | 545 |
| Section 7  | Configuration of disaster recovery environments .....                   | 551 |
| Chapter 31 | Configuring a campus cluster setup for disaster recovery .....          | 553 |
|            | Requirements for an SF Oracle RAC campus cluster .....                  | 553 |
|            | Setting up an SF Oracle RAC campus cluster for disaster recovery .....  | 554 |
|            | Preparing to set up a campus cluster .....                              | 555 |
|            | Configuring VxVM disk groups for Oracle RAC on a campus cluster .....   | 555 |
|            | Tuning guidelines for campus clusters .....                             | 558 |
|            | Best practices for a campus cluster .....                               | 558 |
| Chapter 32 | Setting up a replicated global cluster .....                            | 561 |
|            | Replication in the SF Oracle RAC environment .....                      | 561 |
|            | Requirements for SF Oracle RAC global clusters .....                    | 562 |
|            | Supported software and hardware for SF Oracle RAC .....                 | 562 |
|            | Supported replication technologies for SF Oracle RAC .....              | 562 |
|            | About setting up a global cluster in an SF Oracle RAC environment ..... | 564 |
|            | Configuring a cluster at the primary site .....                         | 565 |
|            | Configuring a cluster at the secondary site .....                       | 567 |
|            | Setting up the cluster on the secondary site .....                      | 567 |
|            | Setting up the database for the secondary site .....                    | 568 |
|            | Configuring replication on clusters at both sites .....                 | 570 |
|            | Modifying the ClusterService group for global clusters .....            | 571 |
|            | Modifying the global clustering configuration using the wizard .....    | 571 |
|            | Defining the remote cluster and heartbeat objects .....                 | 572 |
|            | Configuring the VCS service groups for global clusters .....            | 575 |
| Chapter 33 | Configuring a global cluster using VVR .....                            | 577 |
|            | About configuring global clustering using VVR .....                     | 577 |
|            | Setting up replication using VVR on the primary site .....              | 578 |
|            | Creating the SRL volume on the primary site .....                       | 578 |
|            | Setting up the Replicated Volume Group (RVG) on the primary site .....  | 579 |

|                                                                   |         |
|-------------------------------------------------------------------|---------|
| Setting up replication using VVR on the secondary site .....      | 581     |
| Creating the data and SRL volumes on the secondary site .....     | 581     |
| Editing the /etc/vx/vras/.rdg files .....                         | 582     |
| Setting up IP addresses for RLINKs on each cluster .....          | 582     |
| Setting up the disk group on secondary site for replication ..... | 583     |
| Starting replication of Oracle RAC database volume .....          | 585     |
| Starting replication using automatic synchronization .....        | 585     |
| Starting replication using full synchronization with              |         |
| Checkpoint .....                                                  | 586     |
| Verifying replication status .....                                | 587     |
| Configuring VCS to replicate the database volume using VVR .....  | 587     |
| About modifying the VCS configuration for replication .....       | 588     |
| Configuration examples before and after modification .....        | 589     |
| Modifying the VCS Configuration on the Primary Site .....         | 592     |
| Modifying the VCS Configuration on the Secondary Site .....       | 595     |
| Using VCS commands on SF Oracle RAC global clusters .....         | 600     |
| Using VVR commands on SF Oracle RAC global clusters .....         | 600     |
| About migration and takeover of the primary site role .....       | 600     |
| Migrating the role of primary site to the secondary site .....    | 601     |
| Migrating the role of new primary site back to the original       |         |
| primary site .....                                                | 602     |
| Taking over the primary role by the remote cluster .....          | 602     |
| VCS agents to manage wide-area failover .....                     | 606     |
| <br>Section 8      Uninstallation of SF Oracle RAC .....          | <br>609 |
| <br>Chapter 34    Preparing to uninstall SF Oracle RAC from a     |         |
| cluster .....                                                     | 611     |
| About uninstalling SF Oracle RAC from a cluster .....             | 611     |
| Options for uninstalling SF Oracle RAC .....                      | 612     |
| Preparing to uninstall SF Oracle RAC from a cluster .....         | 613     |
| Stopping Oracle instances .....                                   | 614     |
| Backing up the Oracle database .....                              | 615     |
| Unlinking the SF Oracle RAC libraries from Oracle RAC .....       | 615     |
| Uninstalling Oracle RAC (optional) .....                          | 617     |
| Removing root disk encapsulation .....                            | 617     |
| Stopping the applications that use CVM or CFS (outside of VCS     |         |
| control) .....                                                    | 618     |
| Unmounting CFS file systems (outside of VCS control) .....        | 618     |
| Stopping VCS .....                                                | 619     |
| Stopping the applications that use VxVM or VxFS (outside of       |         |
| VCS control) .....                                                | 620     |

|            |                                                                                                  |     |
|------------|--------------------------------------------------------------------------------------------------|-----|
|            | Unmounting VxFS file systems (outside of VCS control) .....                                      | 620 |
| Chapter 35 | Uninstalling SF Oracle RAC from a cluster .....                                                  | 621 |
|            | Uninstalling SF Oracle RAC from a cluster .....                                                  | 621 |
|            | Removing the CP server configuration using the removal script .....                              | 622 |
|            | Removing the SF Oracle RAC packages .....                                                        | 626 |
|            | Removing other configuration files (optional) .....                                              | 627 |
|            | Removing the Storage Foundation for Databases (SFDB) repository after removing the product ..... | 628 |
|            | Uninstalling SF Oracle RAC with the Veritas Web-based installer .....                            | 629 |
|            | Uninstalling SF Oracle RAC using a response file .....                                           | 631 |
|            | Response file variables to uninstall SF Oracle RAC .....                                         | 631 |
|            | Sample response file for uninstalling SF Oracle RAC .....                                        | 632 |
| Section 9  | Installation reference .....                                                                     | 635 |
| Appendix A | Installation scripts .....                                                                       | 637 |
|            | About installation scripts .....                                                                 | 637 |
|            | Installation script options .....                                                                | 638 |
| Appendix B | SF Oracle RAC installation and configuration sample values .....                                 | 645 |
|            | About the installation and configuration worksheets .....                                        | 645 |
|            | SF Oracle RAC worksheet .....                                                                    | 646 |
|            | Veritas Cluster Server component information .....                                               | 649 |
|            | I/O fencing information .....                                                                    | 650 |
|            | SF Oracle RAC secure mode cluster information .....                                              | 651 |
|            | SF Oracle RAC add user information .....                                                         | 651 |
|            | Global cluster information .....                                                                 | 652 |
|            | Oracle RAC worksheet .....                                                                       | 653 |
|            | Replicated cluster using VVR worksheet .....                                                     | 659 |
|            | Replicated cluster using SRDF worksheet .....                                                    | 660 |
|            | Required installation information for Oracle Clusterware/Grid Infrastructure .....               | 661 |
|            | Required installation information for Oracle database .....                                      | 666 |

|            |                                                                       |     |
|------------|-----------------------------------------------------------------------|-----|
| Appendix C | Sample configuration files .....                                      | 669 |
|            | About VCS configuration file .....                                    | 669 |
|            | About the LLT and GAB configuration files .....                       | 670 |
|            | About I/O fencing configuration files .....                           | 672 |
|            | Sample configuration files .....                                      | 675 |
|            | sfrac01_main.cf file .....                                            | 676 |
|            | sfrac02_main.cf file .....                                            | 677 |
|            | sfrac03_main.cf file .....                                            | 678 |
|            | sfrac04_main.cf file .....                                            | 679 |
|            | sfrac05_main.cf file .....                                            | 681 |
|            | sfrac06_main.cf file .....                                            | 682 |
|            | sfrac07_main.cf and sample08_main.cf files .....                      | 683 |
|            | sfrac09_main.cf and sample10_main.cf files .....                      | 685 |
|            | sfrac11_main.cf file .....                                            | 688 |
|            | sfrac12_main.cf and sample13_main.cf files .....                      | 689 |
|            | sfrac14_main.cf file .....                                            | 692 |
|            | Sample configuration files for CP server .....                        | 693 |
| Appendix D | Setting up inter-system communication .....                           | 701 |
|            | About configuring ssh or rsh using the Veritas installer .....        | 701 |
|            | Setting up inter-system communication .....                           | 702 |
|            | Setting up ssh on cluster systems .....                               | 702 |
| Appendix E | Automatic Storage Management .....                                    | 707 |
|            | About ASM in SF Oracle RAC environments .....                         | 707 |
|            | ASM configuration with SF Oracle RAC .....                            | 708 |
|            | Configuring ASM in SF Oracle RAC environments .....                   | 710 |
|            | Setting up Automatic Storage Management .....                         | 711 |
|            | Copying the VCS IPC libraries .....                                   | 712 |
|            | Creating database storage on ASM .....                                | 713 |
|            | Creating ASM disk groups and instances .....                          | 714 |
|            | Verify the ASM setup .....                                            | 715 |
|            | Configuring VCS service groups for database instances on<br>ASM ..... | 715 |
|            | Sample configuration file Veritas CVM and ASM main.cf file .....      | 718 |
| Appendix F | Creating a test database .....                                        | 719 |
|            | About creating a test database .....                                  | 719 |
|            | Creating a database for Oracle .....                                  | 719 |
|            | Creating the database storage on raw volumes .....                    | 720 |
|            | Creating the database storage on CFS .....                            | 721 |

|            |                                                                 |     |
|------------|-----------------------------------------------------------------|-----|
| Appendix G | High availability agent information .....                       | 723 |
|            | About agents .....                                              | 723 |
|            | VCS agents included within SF Oracle RAC .....                  | 724 |
|            | VCS agents for Oracle included within SF Oracle RAC .....       | 724 |
|            | CVMCluster agent .....                                          | 725 |
|            | Entry points for CVMCluster agent .....                         | 725 |
|            | Attribute definition for CVMCluster agent .....                 | 726 |
|            | CVMCluster agent type definition .....                          | 726 |
|            | CVMCluster agent sample configuration .....                     | 727 |
|            | CVMVxconfigd agent .....                                        | 727 |
|            | Entry points for CVMVxconfigd agent .....                       | 728 |
|            | Attribute definition for CVMVxconfigd agent .....               | 728 |
|            | CVMVxconfigd agent type definition .....                        | 732 |
|            | CVMVxconfigd agent sample configuration .....                   | 732 |
|            | CVMVolDg agent .....                                            | 732 |
|            | Entry points for CVMVolDg agent .....                           | 733 |
|            | Attribute definition for CVMVolDg agent .....                   | 734 |
|            | CVMVolDg agent type definition .....                            | 736 |
|            | CVMVolDg agent sample configuration .....                       | 736 |
|            | CFSMount agent .....                                            | 736 |
|            | Entry points for CFSMount agent .....                           | 737 |
|            | Attribute definition for CFSMount agent .....                   | 737 |
|            | CFSMount agent type definition .....                            | 740 |
|            | CFSMount agent sample configuration .....                       | 741 |
|            | CFSfsckd agent .....                                            | 741 |
|            | Entry points for CFSfsckd agent .....                           | 742 |
|            | Attribute definition for CFSfsckd agent .....                   | 742 |
|            | CFSfsckd agent type definition .....                            | 745 |
|            | CFSfsckd agent sample configuration .....                       | 745 |
|            | PrivNIC agent .....                                             | 745 |
|            | Functions of the PrivNIC agent .....                            | 746 |
|            | Attributes of the PrivNIC agent .....                           | 746 |
|            | States of the PrivNIC agent .....                               | 748 |
|            | Sample service group configuration with the PrivNIC agent ..... | 749 |
|            | Type definition of the PrivNIC resource .....                   | 749 |
|            | Sample configuration of the PrivNIC resource .....              | 750 |
|            | MultiPrivNIC agent .....                                        | 750 |
|            | Managing high availability of private interconnects .....       | 751 |
|            | Functions of the MultiPrivNIC agent .....                       | 751 |
|            | Attributes of the MultiPrivNIC agent .....                      | 751 |
|            | States of the MultiPrivNIC agent .....                          | 753 |

|                                                                                        |            |
|----------------------------------------------------------------------------------------|------------|
| Sample service group configuration with the MultiPrivNIC agent .....                   | 753        |
| Type definition of the MultiPrivNIC resource .....                                     | 754        |
| Sample configuration of the MultiPrivNIC resource .....                                | 755        |
| CSSD agent .....                                                                       | 755        |
| Functions of the CSSD agent .....                                                      | 756        |
| Attributes of the CSSD agent .....                                                     | 756        |
| States of the CSSD agent .....                                                         | 757        |
| Sample service group configurations with the CSSD agent .....                          | 757        |
| Type definition of the CSSD resource .....                                             | 758        |
| Sample configuration of the CSSD resource .....                                        | 758        |
| VCS agents for Oracle .....                                                            | 758        |
| Oracle agent functions .....                                                           | 759        |
| Resource type definition for the Oracle agent .....                                    | 764        |
| Netlsnr agent functions .....                                                          | 771        |
| Resource type definition for the Netlsnr agent .....                                   | 772        |
| ASMDG agent functions .....                                                            | 776        |
| Resource type definition for the ASMDG agent .....                                     | 777        |
| CRSResource agent .....                                                                | 779        |
| Functions of the CRSResource agent .....                                               | 780        |
| States of the CRSResource agent .....                                                  | 780        |
| Attributes of the CRSResource agent .....                                              | 781        |
| VCS service group dependencies with the CRSResource agent .....                        | 781        |
| Resource type definition for the CRSResource agent .....                               | 786        |
| Sample configuration for the CRSResource agent .....                                   | 786        |
| <b>Appendix H SF Oracle RAC deployment scenarios .....</b>                             | <b>789</b> |
| SF Oracle RAC cluster with VCS IPC and PrivNIC agent .....                             | 789        |
| SF Oracle RAC cluster with UDP IPC and PrivNIC agent .....                             | 792        |
| SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent ..... | 793        |
| SF Oracle RAC cluster with isolated Oracle traffic and MultiPrivNIC agent .....        | 796        |
| SF Oracle RAC cluster with NIC bonding, VCS IPC and PrivNIC agent .....                | 798        |
| SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent .....               | 800        |
| Configuration diagrams for setting up server-based I/O fencing .....                   | 802        |
| Two unique client clusters served by 3 CP servers .....                                | 802        |
| Client cluster served by highly available CPS and 2 SCSI-3 disks .....                 | 803        |

- Two node campus cluster served by remote CP server and 2 SCSI-3 disks ..... 805
- Multiple client clusters served by highly available CP server and 2 SCSI-3 disks ..... 807
- SF Oracle RAC deployment scenarios in Oracle VM Server for SPARC environments ..... 809
- Sample configuration scenarios ..... 809
- Preparing to deploy SF Oracle RAC in Logical Domain environments ..... 810
- SF Oracle RAC with Oracle RAC database on I/O domains of two hosts ..... 811
- SF Oracle RAC with Oracle RAC database on guest domains of two hosts ..... 812
- SF Oracle RAC with Oracle RAC database on guest domains of single host ..... 814
- SF Oracle RAC with Oracle RAC database on I/O domain and guest domain of single host ..... 816
- Deploying Storage Foundation for Databases (SFDB) tools in a Storage Foundation for Oracle RAC environment ..... 818
- Glossary ..... 821
- Index ..... 825



## Installation overview and planning

- [Chapter 1. Introducing Veritas Storage Foundation for Oracle RAC](#)
- [Chapter 2. System requirements](#)
- [Chapter 3. Planning to install SF Oracle RAC](#)
- [Chapter 4. Licensing SF Oracle RAC](#)



# Introducing Veritas Storage Foundation for Oracle RAC

This chapter includes the following topics:

- [About Veritas Storage Foundation for Oracle RAC](#)
- [About basic SF Oracle RAC components](#)
- [About SF Oracle RAC features](#)
- [About SF Oracle RAC optional components](#)
- [SF Oracle RAC cluster setup models](#)

## About Veritas Storage Foundation for Oracle RAC

Veritas Storage Foundation™ for Oracle® RAC (SF Oracle RAC) leverages proprietary storage management and high availability technologies to enable robust, manageable, and scalable deployment of Oracle RAC on UNIX platforms. The solution uses Veritas Cluster File System technology that provides the dual advantage of easy file system management as well as the use of familiar operating system tools and utilities in managing databases.

The solution stack comprises the Veritas Cluster Server (VCS), Veritas Cluster Volume Manager (CVM), Veritas Oracle Real Application Cluster Support (VRTSdbac), Veritas Oracle Disk Manager (VRTSodm), Veritas Cluster File System (CFS), and Veritas Storage Foundation, which includes the base Veritas Volume Manager (VxVM) and Veritas File System (VxFS).

## Benefits of SF Oracle RAC

SF Oracle RAC provides the following benefits:

- Support for file system-based management. SF Oracle RAC provides a generic clustered file system technology for storing and managing Oracle data files as well as other application data.
- Support for high-availability of cluster interconnects.  
For Oracle RAC 10g Release 2:  
The combination of LMX/LLT protocols and the PrivNIC/MultiPrivNIC agents provides maximum bandwidth as well as high availability of the cluster interconnects, including switch redundancy.  
For Oracle RAC 11g Release 1/Oracle RAC 11g Release 2:  
The PrivNIC/MultiPrivNIC agents provide maximum bandwidth as well as high availability of the cluster interconnects, including switch redundancy.
- Use of clustered file system and volume management technologies for placement of Oracle Cluster Registry (OCR) and voting disks. These technologies provide robust shared block and raw interfaces for placement of OCR and voting disks. In the absence of SF Oracle RAC, separate LUNs need to be configured for OCR and voting disks.
- Support for a standardized approach toward application and database management. A single-vendor solution for the complete SF Oracle RAC software stack lets you devise a standardized approach toward application and database management. Further, administrators can apply existing expertise of Veritas technologies toward SF Oracle RAC.
- Increased availability and performance using dynamic multi-pathing (DMP). DMP provides wide storage array support for protection from failures and performance bottlenecks in the HBAs, SAN switches, and storage arrays.
- Easy administration and monitoring of SF Oracle RAC clusters from a single web console.
- Support for many types of applications and databases.
- Improved file system access times using Oracle Disk Manager (ODM).
- Ability to configure ASM disk groups over CVM volumes to take advantage of dynamic multi-pathing (DMP).
- Enhanced scalability and availability with access to multiple Oracle RAC instances per database in a cluster.
- Support for backup and recovery solutions using volume-level and file system-level snapshot technologies. SF Oracle RAC enables full volume-level snapshots for off-host processing and file system-level snapshots for efficient backup and rollback.
- Ability to failover applications without downtime using clustered file system technology.

- Prevention of data corruption in split-brain scenarios with robust SCSI-3 Persistent Group Reservation (PGR) based I/O fencing or Coordination Point Server-based I/O fencing. The preferred fencing feature also enables you to specify how the fencing driver determines the surviving subcluster.
- Support for sharing all types of files, in addition to Oracle database files, across nodes.
- Fast disaster recovery with minimal downtime and interruption to users. Users can transition from a local high availability site to a wide-area disaster recovery environment with primary and secondary sites. If a node fails, clients that are attached to the failed node can reconnect to a surviving node and resume access to the shared database. Recovery after failure in the SF Oracle RAC environment is far quicker than recovery for a failover database.
- Verification of disaster recovery configuration using fire drill technology without affecting production systems.
- Support for a wide range of hardware replication technologies as well as block-level replication using VVR.
- Support for campus clusters with the following capabilities:
  - Consistent reattach with Site Awareness
  - Site aware reads with VxVM mirroring
  - Monitoring of Oracle resources
  - Protection against split-brain scenarios

## About basic SF Oracle RAC components

The basic components of SF Oracle RAC are as follows:

|                        |                                                                                                                                                 |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Veritas Cluster Server | Veritas Cluster Server (VCS) manages Oracle RAC databases and infrastructure components in a clustered environment.                             |
| Cluster Volume Manager | Cluster Volume Manager (CVM) enables simultaneous access to the shared volumes that are based on technology from Veritas Volume Manager (VxVM). |
| Cluster File System    | Cluster File System (CFS) enables simultaneous access to the shared file systems that are based on technology from Veritas File System (VxFS).  |

#### Oracle Disk Manager

Oracle Disk Manager (ODM) is a disk and file management interface that is provided by Oracle to improve disk I/O performance. ODM enables Oracle to allocate and release disk space, manage tablespaces, and read/write disk blocks directly. SF Oracle RAC uses a custom driver that enables applications to use ODM for enhanced file system performance and easy file administration.

#### RAC Extensions

RAC Extensions manage the cluster membership and communications between cluster nodes.

For a detailed understanding of each component and the architectural overview, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

## About SF Oracle RAC features

You can configure the following features in an SF Oracle RAC cluster:

- I/O fencing  
See [“About I/O fencing”](#) on page 31.
- Symantec Operations Readiness Tools  
See [“Symantec Operations Readiness Tools”](#) on page 34.
- VCS notifications  
See [“About VCS notifications”](#) on page 35.
- Campus clusters  
See [“About campus clusters”](#) on page 35.
- Global clusters  
See [“About global clusters”](#) on page 35.
- Storage Foundation Database Management tools  
See [“About database management using SF Oracle RAC”](#) on page 35.
- Veritas Volume Replicator  
See [“About Veritas Volume Replicator”](#) on page 35.

---

**Note:** I/O fencing is mandatory in SF Oracle RAC installations. All other features are optional and may be configured to suit your business needs.

---

## About I/O fencing

I/O fencing protects the data on shared disks when nodes in a cluster detect a change in the cluster membership that indicates a split-brain condition.

The fencing operation determines the following:

- The nodes that must retain access to the shared storage
- The nodes that must be ejected from the cluster

This decision prevents possible data corruption. The installer installs the I/O fencing driver, VRTSvxfen, when you install SF Oracle RAC. To protect data on shared disks, you must configure I/O fencing after you install and configure SF Oracle RAC.

I/O fencing technology uses coordination points for arbitration in the event of a network partition.

You can configure I/O fencing to use one or both of the following components as coordination points:

|                                       |                                                                                                                                                                              |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Coordinator disk                      | I/O fencing that uses coordinator disks is referred to as disk-based I/O fencing.<br><br>Disk-based I/O fencing ensures data integrity in a single cluster.                  |
| Coordination point server (CP server) | I/O fencing that uses at least one CP server system is referred to as server-based I/O fencing.<br><br>Server-based I/O fencing ensures data integrity in multiple clusters. |

See the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

### About server-based I/O fencing

In a disk-based I/O fencing implementation, the vxfen driver handles various SCSI-3 PR based arbitration operations completely within the driver. I/O fencing also provides a framework referred to as customized fencing wherein arbitration operations are implemented in custom scripts. The vxfen driver invokes the custom scripts.


The CP server-based coordination point uses a customized fencing framework. Note that SCSI-3 PR based fencing arbitration can also be enabled using customized fencing framework. This allows the user to specify a combination of SCSI-3 LUNs and CP servers as coordination points using customized fencing. Customized

fencing can be enabled by specifying `vxfen_mode=customized` and `vxfen_mechanism=cps` in the `/etc/vxfenmode` file.

Moreover, both `/etc/vxfenmode` and `/etc/vxfentab` files contain additional parameter "security" which indicates if communication between CP server and SF Oracle RAC cluster nodes is secure.

Figure 1-1 displays a schematic of the customized fencing options.

**Figure 1-1** Customized fencing



A user level daemon `vxfend` interacts with the `vxfen` driver, which in turn interacts with `GAB` to get the node membership update. Upon receiving membership updates, `vxfend` invokes various scripts to race for the coordination point and fence off data disks. The `vxfend` daemon manages various fencing agents. The customized fencing scripts are located in the `/opt/VRTSvcs/vxfen/bin/customized/cps` directory.



The scripts that are involved include the following:

- `generate_snapshot.sh`: Retrieves the SCSI ID's of the coordinator disks and/or UUID ID's of the CP servers  
For information about the UUID (Universally Unique Identifier), see the *Veritas Cluster Server Administrator's Guide*.
- `join_local_node.sh`: Registers the keys with the coordinator disks or CP servers
- `race_for_coordination_point.sh`: Races to determine a winner after cluster reconfiguration
- `unjoin_local_node.sh`: Removes the keys that are registered in `join_local_node.sh`
- `fence_data_disks.sh`: Fences the data disks from access by the losing nodes.
- `local_info.sh`: Lists local node's configuration parameters and coordination points, which are used by the vxfen driver.

### About the CP server database

CP server requires a database for storing the registration keys of the SF Oracle RAC cluster nodes. CP server uses a SQLite database for its operations. By default, the database is located at `/etc/VRTScps/db`.

For a single node VCS cluster hosting a CP server, the database can be placed on a local file system. For an SFHA cluster hosting a CP server, the database must be placed on a shared file system. The file system must be shared among all nodes that are part of the SFHA cluster.

In an SFHA cluster hosting the CP server, the shared database is protected by setting up SCSI-3 PR based I/O fencing. SCSI-3 PR based I/O fencing protects against split-brain scenarios.

---

**Warning:** The CP server database must not be edited directly and should only be accessed using `cpsadm(1M)`. Manipulating the database manually may lead to undesirable results including system panics.

---

### About the CP server user types and privileges

The CP server supports the following user types, each with a different access level privilege:

- CP server administrator (admin)
- CP server operator

Different access level privileges permit the user to issue different commands. If a user is neither a CP server admin nor a CP server operator user, then the user has guest status and can issue limited commands.

The user types and their access level privileges are assigned to individual users during SF Oracle RAC cluster configuration for fencing. During the installation process, you are prompted for a user name, password, and access level privilege (CP server admin or CP server operator).

To administer and operate a CP server, there must be at least one CP server admin.

A root user on a CP server is given all the administrator privileges, and these administrator privileges can be used to perform all the CP server specific operations.

## About preferred fencing

The I/O fencing driver uses coordination points to prevent split-brain in a VCS cluster. By default, the fencing driver favors the subcluster with maximum number of nodes during the race for coordination points. With the preferred fencing feature, you can specify how the fencing driver must determine the surviving subcluster.

You can configure the preferred fencing policy using the cluster-level attribute PreferredFencingPolicy as follows:

- Enable system-based preferred fencing policy to give preference to high capacity systems.
- Enable group-based preferred fencing policy to give preference to service groups for high priority applications.
- Disable preferred fencing policy to use the default node count-based race policy.

See the *Veritas Storage Foundation for Oracle RAC Administrator's Guide* for more details.

## Symantec Operations Readiness Tools

Symantec™ Operations Readiness Tools (SORT) is a set of Web-based tools that supports Symantec enterprise products. SORT increases operational efficiency and helps improve application availability.

Among its broad set of features, SORT provides patches, patch notifications, and documentation for Symantec enterprise products.

To access SORT, go to:

<http://sort.symantec.com>

## About VCS notifications

You can configure both SNMP and SMTP notifications for VCS. Symantec recommends you to configure at least one of these notifications. You have the following options:

- Configure SNMP trap notification of VCS events using the VCS Notifier component
- Configure SMTP email notification of VCS events using the VCS Notifier component.

See the *Veritas Cluster Server Administrator's Guide*.

## About campus clusters

A campus cluster has alternate nodes located in different data centers. Campus clusters are connected using a high speed cable that guarantees network access between the nodes. The campus cluster configuration provides local high availability and disaster recovery functionality in a single SF Oracle RAC cluster. This configuration uses data mirroring to duplicate data at different sites.

SF Oracle RAC supports campus clusters that employ shared disk groups mirrored with Veritas Volume Manager (VxVM).

## About global clusters

Global clusters provide the ability to fail over applications between geographically distributed clusters when disaster occurs. This type of clustering involves migrating applications between clusters over a considerable distance. You can set up HA/DR using hardware-based or software-based replication technologies.

## About Veritas Volume Replicator

Veritas Volume Replicator (VVR) is a software-based replication technology used in global cluster disaster recovery setups that replicates data to remote sites over any standard IP network. You can have up to 32 remote sites.

## About database management using SF Oracle RAC

You can leverage the database management capabilities of the Storage Foundation for Databases (SFDB) tools to simplify storage management and improve database performance.

For information on supported capabilities, see the *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases* guide.

## About SF Oracle RAC optional components

You can add the following optional components to SF Oracle RAC:

|                                         |                                                                                      |
|-----------------------------------------|--------------------------------------------------------------------------------------|
| Symantec Product Authentication Service | See <a href="#">“About Symantec Product Authentication Service (AT)”</a> on page 36. |
| Veritas Operations Manager              | See <a href="#">“Veritas Operations Manager”</a> on page 36.                         |
| Cluster Manager (Java console)          | See <a href="#">“About Cluster Manager (Java Console)”</a> on page 37.               |

### About Symantec Product Authentication Service (AT)

SF Oracle RAC uses Symantec Product Authentication Service (AT) to provide secure communication between cluster nodes and clients. It uses digital certificates for authentication and SSL to encrypt communication over the public network to secure communications.

AT uses the following brokers to establish trust relationship between the cluster components:

- Root broker

A root broker serves as the main registration and certification authority; it has a self-signed certificate and can authenticate other brokers. The root broker is only used during initial creation of an authentication broker.

A root broker on a stable external system can serve multiple clusters. Symantec recommends that you install a single root broker on a utility system. The utility system, such as an email server or domain controller, can be highly available. You can also configure one of the nodes in the SF Oracle RAC cluster to serve as a root and an authentication broker.

- Authentication brokers

Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have root-signed certificates. Each node in VCS serves as an authentication broker.

See Symantec Product Authentication Service documentation for more information.

See [“Preparing to configure the clusters in secure mode”](#) on page 119.

### Veritas Operations Manager

Symantec recommends use of Veritas Operations Manager to manage Storage Foundation and Cluster Server environments.

The Veritas Enterprise Administrator (VEA) console is no longer packaged with Storage Foundation products. If you wish to continue using VEA, a version is available for download from [http://go.symantec.com/vcsm\\_download](http://go.symantec.com/vcsm_download). Veritas Storage Foundation Management Server is no longer supported.

If you wish to manage a single cluster using Cluster Manager (Java Console), a version is available for download from [http://go.symantec.com/vcsm\\_download](http://go.symantec.com/vcsm_download). Veritas Cluster Server Management Console is no longer supported.

Veritas Operations Manager provides a centralized management console for Veritas Storage Foundation and High Availability products. You can use Veritas Operations Manager to monitor, visualize, and manage storage resources and generate reports. Veritas Operations Manager is not available on the Storage Foundation and High Availability Solutions release. You can download Veritas Operations Manager at no charge at <http://go.symantec.com/vom>.

Refer to the Veritas Operations Manager documentation for installation, upgrade, and configuration instructions.

## About Cluster Manager (Java Console)

Cluster Manager (Java Console) offers complete administration capabilities for your cluster. Use the different views in the Java Console to monitor clusters and VCS objects, including service groups, systems, resources, and resource types.

See *Veritas Cluster Server Administrator's Guide*.

You can download the console from [http://go.symantec.com/vcsm\\_download](http://go.symantec.com/vcsm_download).

## SF Oracle RAC cluster setup models

SF Oracle RAC supports a variety of cluster configurations.

Depending on your business needs, you may choose from the following setup models:

- Basic setup  
See “[Typical configuration of four-node SF Oracle RAC cluster](#)” on page 38.
- Secure setup  
See “[Typical configuration of SF Oracle RAC clusters in secure mode](#)” on page 39.
- Central management setup  
See “[Typical configuration of VOM-managed SF Oracle RAC clusters](#)” on page 40.
- Campus cluster setup

See “Typical configuration of SF Oracle RAC campus clusters for disaster recovery” on page 41.


■ Global cluster setup

See “Typical configuration of SF Oracle RAC global clusters for disaster recovery” on page 43.

## Typical configuration of four-node SF Oracle RAC cluster

Figure 1-2 depicts a high-level view of a basic SF Oracle RAC configuration for a four-node cluster.

**Figure 1-2** Sample four-node SF Oracle RAC cluster



A basic topology has the following layout and characteristics:

- Multiple client applications that access nodes in the cluster over a public network.
- Nodes that are connected by at least two private network links (also called cluster interconnects) using 100BaseT or gigabit Ethernet controllers on each system.  
If the private links are on a single switch, isolate them using VLAN.
- Nodes that are connected to iSCSI or Fibre Channel shared storage devices over SAN.  
All shared storage must support SCSI-3 PR.
- Nodes that are connected with private network links using similar network devices and matching port numbers.  
For example, if you use bge1 on one end of a link, it is recommended that the other end also use bge1.
- Every system must have a local disk.  
The root disk can be encapsulated using VxVM. The root disk group for the storage must not be shared between systems.
- The Oracle Cluster Registry and vote disks configured on the shared storage that is available to each node.  
For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1: The shared storage for Oracle Cluster Registry and vote disks can be a cluster file system or raw VxVM volumes.  
For Oracle RAC 11g Release 2: The shared storage for Oracle Cluster Registry and vote disks can be a cluster file system or ASM disk groups created using raw VxVM volumes.
- Three or more odd number of disks or LUNs used as coordinator disks for I/O fencing.
- VCS manages the resources that are required by Oracle RAC. The resources must run in parallel on each node.


## Typical configuration of SF Oracle RAC clusters in secure mode

SF Oracle RAC uses Symantec Product Authentication Service (AT) to provide secure communication between cluster nodes and clients.

See [“About Symantec Product Authentication Service \(AT\)”](#) on page 36.

[Figure 1-3](#) illustrates typical configuration of VCS clusters in secure mode. You can use one of the cluster nodes as AT root broker or you can use a stable system outside the cluster as AT root broker.

**Figure 1-3** Typical configuration of VCS clusters in secure mode



## Typical configuration of VOM-managed SF Oracle RAC clusters


Veritas Operations Manager (VOM) provides a centralized management console for Veritas Storage Foundation and High Availability products.

See “[Veritas Operations Manager](#)” on page 36.

[Figure 1-4](#) illustrates a typical setup of SF Oracle RAC clusters that are centrally managed using Veritas Operations Manager. You can install Symantec Product Authentication Service root broker on the same system as that of VOM Central Server or on a different system.



**Figure 1-4** Typical configuration of VOM-managed clusters



## Typical configuration of SF Oracle RAC campus clusters for disaster recovery

A campus cluster configuration provides local high availability and disaster recovery capability in a single SF Oracle RAC cluster. This configuration uses data mirroring to duplicate data at different sites. No host or array replication is involved in the process. SF Oracle RAC supports campus clusters that employ shared disk groups mirrored with Cluster Volume Manager (CVM).


The SF Oracle RAC campus cluster addresses the following basic challenges in campus cluster configurations:

|                    |                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Latency challenges | An SF Oracle RAC campus cluster handles latency challenges in keeping mirrors synchronized while ensuring the efficient recovery in case of site failures for both data and VxVM metadata. |
| Read performance   | The read performance is enhanced as data is read from local mirrors.                                                                                                                       |
| Site awareness     | SF Oracle RAC makes sure that all the mirrors on a site are detached proactively even when a part of the site goes down.                                                                   |

**Note:** The DiskGroupSnap agent is not supported for SF Oracle RAC.

Figure 1-5 illustrates a basic campus cluster setup.

Figure 1-5 Basic campus cluster setup




For more information, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

## Typical configuration of SF Oracle RAC global clusters for disaster recovery

SF Oracle RAC leverages the global clustering feature of VCS to enable high availability and disaster recovery (HA/DR) for businesses that span wide geographical areas. Global clusters provide protection against outages caused by large-scale disasters such as major floods, hurricanes, and earthquakes. An entire cluster can be affected by such disasters. This type of clustering involves migrating applications between clusters over a considerable distance.

You can set up HA/DR using hardware-based or software-based replication technologies.

**Figure 1-6** Global clusters



To understand how global clusters work, review the example of an Oracle RAC database configured using global clustering. Oracle RAC is installed and configured in cluster A and cluster B. Oracle database is located on shared disks within each cluster and is replicated across clusters to ensure data concurrency. The VCS service groups for Oracle are online on a node in cluster A and are configured to fail over on cluster A and cluster B.

VCS continuously monitors and communicates events between clusters. If cluster A fails, the Oracle database is started on the remote cluster B.

---

**Note:** You must have an SF Oracle RAC HA/DR license to configure global clusters. If you use VVR for replication, you must also have a VVR license. You may configure a basic cluster initially and add the HA/DR and VVR licenses at a later time or you may add the licenses during the SF Oracle RAC installation.

---

For information on supported replication technologies and more information, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide*.

# System requirements

This chapter includes the following topics:

- [Important preinstallation information](#)
- [Hardware requirements](#)
- [I/O fencing requirements](#)
- [Supported Solaris operating systems](#)
- [Required operating system patches](#)
- [Supported database software](#)
- [Veritas File System requirements](#)
- [Supported replication technologies for global clusters](#)

## Important preinstallation information

Before you install SF Oracle RAC, make sure you have reviewed the following information:

- Hardware compatibility list for information about supported hardware:  
<http://entsupport.symantec.com/docs/330441>
- Disk storage array support information:  
<http://entsupport.symantec.com/docs/283282>
- Latest information on support for Oracle database versions:  
<http://www.symantec.com/docs/TECH44807>
- General information regarding the release, installation notes, known issues, and fixed issues:  
See *Veritas Storage Foundation for Oracle RAC Release Notes*.

- Oracle documentation for additional requirements pertaining to your version of Oracle.

# Hardware requirements

Depending on the type of setup planned, make sure you meet the necessary hardware requirements.

For basic clusters                      See [Table 2-1](#) on page 46.

For campus clusters                  See [Table 2-2](#) on page 47.

**Table 2-1**                      Hardware requirements for basic clusters

| Item                  | Description                                                                                                                                                                                                                                                                                                                                            |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC systems | Two to sixteen systems with two or more CPUs at 1 GHz or higher for Solaris SPARC systems.<br><br>Two to sixteen systems with two or more CPUs at 2 GHz or higher for Solaris x64 systems.                                                                                                                                                             |
| DVD drive             | A DVD drive on one of the nodes in the cluster.                                                                                                                                                                                                                                                                                                        |
| Disks                 | SF Oracle RAC requires that all storage disks support SCSI-3 Persistent Reservations (PR).<br><br><b>Note:</b> The coordinator disk does not store data, so configure the disk as the smallest possible LUN on a disk array to avoid wasting space.                                                                                                    |
| Disk space            | You can evaluate your systems for available disk space by running the product installation program. Navigate to the product directory on the product disc and run the following command:<br><br><code># ./installsfrc -precheck node_name</code><br><br>For details on the additional space that is required for Oracle, see the Oracle documentation. |
| RAM                   | Each SF Oracle RAC system requires at least 2 GB.<br><br>Symantec recommends additional amount of at least twice the Oracle SGA size.                                                                                                                                                                                                                  |
| Swap space            | See the Oracle Metalink document: 169706.1                                                                                                                                                                                                                                                                                                             |

**Table 2-1** Hardware requirements for basic clusters (*continued*)

| Item                                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Network links                           | <p>Two or more private links and one public link.</p> <p>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. These links must be of the same type; you cannot mix 100BaseT and gigabit.</p> <p>Symantec recommends gigabit Ethernet using enterprise-class switches for the private links.</p> <p>You can also configure aggregated interfaces.</p> |
| Fiber Channel or SCSI host bus adapters | At least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.                                                                                                                                                                                                                                                                                                                                                                  |

**Table 2-2** lists the hardware requirements for campus clusters in addition to the basic cluster requirements.

**Table 2-2** Hardware requirements for campus clusters

| Item        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage     | <ul style="list-style-type: none"> <li>■ The storage switch (to which each host on a site connects) must have access to storage arrays at all the sites.</li> <li>■ Volumes must be mirrored with storage allocated from at least two sites.</li> <li>■ DWDM links are recommended between sites for storage links. DWDM works at the physical layer and requires multiplexer and de-multiplexer devices.</li> <li>■ The storage and networks must have redundant-loop access between each node and each storage array to prevent the links from becoming a single point of failure.</li> </ul> |
| Network     | <ul style="list-style-type: none"> <li>■ Oracle requires that all nodes use the IP addresses from the same subnet.</li> <li>■ Symantec recommends two Network Interface Cards (NIC) per host for LLT heartbeats. Oracle Clusterware requires one private and one virtual IP for each host.</li> <li>■ Symantec recommends a common cross-site physical infrastructure for storage and LLT private networks.</li> </ul>                                                                                                                                                                          |
| I/O fencing | I/O fencing requires placement of a third coordinator disk at a third site. The DWDM can be extended to the third site or the iSCSI LUN at the third site can be used as the third coordination point. Alternatively Coordination Point Server can be deployed at the third remote site as an arbitration point.                                                                                                                                                                                                                                                                                |

## I/O fencing requirements

Depending on whether you plan to configure disk-based fencing or server-based fencing, make sure that you meet the requirements for coordination points:

- Coordinator disks  
See [“Coordinator disk requirements for I/O fencing”](#) on page 48.
- CP servers  
See [“CP server requirements”](#) on page 48.

### Coordinator disk requirements for I/O fencing

Make sure that the I/O fencing coordinator disks meet the following requirements:

- For disk-based I/O fencing, you must have three coordinator disks.
- The coordinator disks can be raw devices, DMP devices, or iSCSI devices.
- Each of the coordinator disks must use a physically separate disk or LUN.  
Symantec recommends using the smallest possible LUNs for coordinator disks.
- Each of the coordinator disks should exist on a different disk array, if possible.
- The coordinator disks must support SCSI-3 persistent reservations.
- Symantec recommends using hardware-based mirroring for coordinator disks.
- Coordinator disks must not be used to store data or must not be included in disk groups that store user data.
- Coordinator disks cannot be the special devices that array vendors use. For example, you cannot use EMC gatekeeper devices as coordinator disks.

### CP server requirements

SF Oracle RAC 5.1SP1 clusters (application clusters) support CP servers which are hosted on the following VCS and SFHA versions:

- VCS 5.1 or 5.1SP1 single-node cluster  
CP server requires LLT and GAB to be configured on the single-node VCS cluster that hosts CP server. This requirement also applies to any single-node application cluster that uses server-based fencing.
- SFHA 5.1 or 5.1SP1 cluster



---

**Warning:** Before you upgrade CP server nodes to use VCS or SFHA 5.1SP1, you must upgrade all the application clusters that use this CP server to version 5.1SP1. Application clusters at version 5.1 cannot communicate with CP server that runs VCS or SFHA 5.1 SP1.

---

Make sure that you meet the basic hardware requirements for the VCS/SFHA cluster to host the CP server.

See the *Veritas Cluster Server Installation Guide* or the *Veritas Storage Foundation High Availability Installation Guide*.

---

**Note:** While Symantec recommends at least three coordination points for fencing, a single CP server as coordination point is a supported server-based fencing configuration. Such single CP server fencing configuration requires that the coordination point be a highly available CP server that is hosted on an SFHA cluster.

---

Make sure you meet the following additional CP server requirements which are covered in this section before you install and configure CP server:

- Hardware requirements
- Operating system requirements
- Networking requirements (and recommendations)
- Security requirements

[Table 2-3](#) lists additional requirements for hosting the CP server.

**Table 2-3** CP server hardware requirements

| Hardware required | Description                                                                                                                                                                                                                                                                                          |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disk space        | To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: <ul style="list-style-type: none"><li>■ 550 MB in the /opt directory (additionally, the language pack requires another 15 MB)</li><li>■ 300 MB in /usr</li><li>■ 20 MB in /var</li></ul> |
| Storage           | When CP server is hosted on an SFHA cluster, there must be shared storage between the CP servers.                                                                                                                                                                                                    |
| RAM               | Each CP server requires at least 512 MB.                                                                                                                                                                                                                                                             |

**Table 2-3** CP server hardware requirements (*continued*)

| Hardware required                      | Description                                                                                                            |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| CP server to client node physical link | A secure TCP/IP connection is required to connect the CP servers to the SF Oracle RAC clusters (application clusters). |

[Table 2-4](#) displays the CP server supported operating systems and versions. An application cluster can use a CP server that runs any of the following supported operating systems.

**Table 2-4** CP server supported operating systems and versions

| CP server                                                           | Operating system and version                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CP server hosted on a VCS single-node cluster or on an SFHA cluster | <p>CP server supports any of the following operating systems:</p> <ul style="list-style-type: none"> <li>■ AIX 5.3 and 6.1</li> <li>■ HP-UX 11i v3</li> <li>■ Linux: <ul style="list-style-type: none"> <li>■ RHEL 5</li> <li>■ SLES 10</li> <li>■ SLES 11</li> </ul> </li> <li>■ Solaris 9 and 10</li> </ul> <p>Review other details such as supported operating system levels and architecture for the supported operating systems.</p> <p>See the <i>Veritas Cluster Server Installation Guide</i> or the <i>Veritas Storage Foundation High Availability Installation Guide</i>.</p> |

Following are the CP server networking requirements and recommendations:

- Symantec recommends that network access from the application clusters to the CP servers should be made highly-available and redundant. The network connections require either a secure LAN or VPN.
- The CP server uses the TCP/IP protocol to connect to and communicate with the application clusters by these network paths. The CP server listens for messages from the application clusters using TCP port 14250. This is the default port that can be changed during a CP server configuration.
- The CP server supports either Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with the application clusters. If the CP server is configured to use an IPv6 virtual IP address, then the application clusters should also be on the IPv6 network where the CP server is being hosted.

- When placing the CP servers within a specific network configuration, you must take into consideration the number of hops from the different application cluster nodes to the CP servers. As a best practice, Symantec recommends that the number of hops from the different application cluster nodes to the CP servers should be equal. This ensures that if an event occurs that results in an I/O fencing scenario, there is no bias in the race due to the number of hops between the nodes.

For secure communications between the SF Oracle RAC cluster and CP server, consider the following requirements and suggestions:

- In a secure communication environment, all CP servers that are used by the application cluster must be configured with security enabled. A configuration where the application cluster uses some CP servers running with security enabled and other CP servers running with security disabled is not supported.
- The CP server and application clusters should also use the same root broker. If the same root broker is not being used, then trust can be established between the cluster nodes and CP server for the secure communication. Trust can be established by the installer when configuring fencing.
- For non-secure communication between CP server and application clusters, there is no need to configure Symantec Product Authentication Service. In non-secure mode, authorization is still provided by CP server for the application cluster users. The authorization that is performed only ensures that authorized users can perform appropriate actions as per their user privileges on the CP server.

For information about establishing secure communications between the application cluster and CP server, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

## Supported Solaris operating systems

This release of the Veritas products is supported on the following Solaris operating systems:

- Solaris 9 (32-bit and 64-bit, SPARC) with Update 7, 8, and 9  
Symantec VirtualStore is only supported on Solaris 9 (SPARC Platform 64-bit).

---

**Note:** In the next major release, Veritas products will not support Solaris 9.

---

- Solaris 10 (64-bit, SPARC or x86\_64) with Update 6, 7, 8, and 9  
Solaris 10 (SPARC and x86\_64) with Update 9 requires VRTSvxvm patch 142629-08 (SPARC) or 142630-08 (x86\_64)

Symantec VirtualStore is only supported on Solaris 10 (SPARC or X86 Platform 64-bit).

For the most up-to-date list of operating system patches, refer to the Release Notes for your product.

For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/334829>

# Required operating system patches

Some required operating system patches may already be present on your system. Use the `showrev -p` command to display the patches on your system.

If the Solaris patches are not already installed on your system, install them before installing SF Oracle RAC.

Table 2-5 lists the required Solaris patches

Table 2-5 Required Solaris patches

| Solaris version | Required patches                                                                                                                                                                                                                   |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Solaris 10      | <div><div>■ Solaris SPARC</div><div>No additional patches are required on Solaris 10 Update 6 and later.</div><div>■ Solaris x64</div><div>118344-14, 118855-36, 119043-11, 119131-33, 120012-14, 125732-05, 127128-11</div></div> |

**Table 2-5** Required Solaris patches (*continued*)

| Solaris version | Required patches                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Solaris 9       | <p>122300-29 - required for Live Upgrade</p> <p>For Solaris 9 Update 7 (to bring the OS kernel to FS patch level 122300):</p> <ul style="list-style-type: none"><li>■ 117171-17</li><li>■ 113073-14</li></ul> <p>For Solaris 9 Update 7, Update 8, and Update 9:</p> <ul style="list-style-type: none"><li>■ 112908-33 (required for 122300-10)</li><li>■ 118558-39 (required for 122300-10)</li><li>■ 122300-10</li></ul> <p>Additionally, install the SAN patch <code>SAN_s9_4.4.14_install_it</code> to satisfy other package requirements and inter-dependencies.</p> <p><b>Caution:</b> Higher patch versions, such as 122300-35, may cause Oracle RAC 10g versions greater than 10.2.0.1 and Oracle RAC 11g Release 1 installations to freeze.</p> |

## Supported database software

For the latest information on supported Oracle database versions, see the following Technical Support TechNote:

<http://www.symantec.com/docs/TECH44807>

---

**Note:** SF Oracle RAC supports only 64-bit Oracle.

---

The following database versions are supported:

On Solaris SPARC:

- Oracle RAC 10g Release 2
- Oracle RAC 11g Release 1
- Oracle RAC 11g Release 2

On Solaris x64:

- Oracle RAC 10g Release 2
- Oracle RAC 11g Release 2

Additionally, see the Oracle documentation for patches that may be required by Oracle for each release.

## Veritas File System requirements

Veritas File System requires that the values of the Solaris variables `lwp_default_stksize` and `svc_default_stksize` are at least 0x6000. When you install the Veritas File System package, `VRTSvxfs`, the `VRTSvxfs` packaging scripts check the values of these variables in the kernel. If the values are less than the required values, `VRTSvxfs` increases the values and modifies the `/etc/system` file with the required values. If the `VRTSvxfs` scripts increase the values, the installation proceeds as usual except that you must reboot and restart the installation program. A message displays if a reboot is required.

To avoid an unexpected need for a reboot, verify the values of the variables before installing Veritas File System. Use the following commands to check the values of the variables:

```
# echo "lwp_default_stksize/X" | mdb -k
lwp_default_stksize:
lwp_default_stksize:          6000

# echo "svc_default_stksize/X" | mdb -k
svc_default_stksize:
svc_default_stksize:          6000
```

If the values shown are less than 6000, you can expect a reboot after installation.

---

**Note:** The default value of the `svc_default_stksize` variable is 0 (zero), which indicates that the value is set to the value of the `lwp_default_stksize` variable. In this case, no reboot is required, unless the value of the `lwp_default_stksize` variable is too small.

---

To avoid a reboot after installation, you can modify the `/etc/system` file with the appropriate values. Reboot the system prior to installing the packages. Add the following lines to the `/etc/system` file:

```
set lwp_default_stksize=0x6000
set rpcmod:svc_default_stksize=0x6000
```

# Supported replication technologies for global clusters

SF Oracle RAC supports the following hardware-based replication and software-based replication technologies for global cluster configurations:

## Hardware-based replication

- EMC SRDF
- Hitachi TrueCopy
- IBM Metro Mirror
- IBM SAN Volume Controller (SVC)
- EMC MirrorView

## Software-based replication

- Veritas Volume Replicator
- Oracle Data Guard





# Planning to install SF Oracle RAC

This chapter includes the following topics:

- [Planning your network configuration](#)
- [Planning the storage](#)
- [Planning volume layout](#)
- [Planning file system design](#)
- [About planning to configure I/O fencing](#)
- [Planning for cluster management](#)
- [Planning for disaster recovery](#)
- [Planning for Oracle VM Server for SPARC environments](#)

## Planning your network configuration

The following practices are recommended for a resilient network setup:

- Configure the private cluster interconnect over multiple dedicated gigabit Ethernet links. All single point of failures such as network interface cards (NIC), switches, and interconnects should be eliminated.
- The NICs used for the private cluster interconnect should have the same characteristics regarding speed, MTU, and full duplex on all nodes. Do not allow the NICs and switch ports to auto-negotiate speed.
- Configure non-routable IP addresses for private cluster interconnects.

- The default value for LLT peer inactivity timeout is 16 seconds. The value should be set based on service availability requirements and the propagation delay between the cluster nodes in case of campus cluster setup. The LLT peer inactivity timeout value indicates the interval after which SF Oracle RAC on one node declares the other node in the cluster dead, if there is no network communication (heartbeat) from that node.

The default value for the CSS miss-count in case of SF Oracle RAC is 600 seconds. The value of this parameter is much higher than the LLT peer inactivity timeout so that the two clusterwares, VCS and Oracle Clusterware, do not interfere with each other's decisions on which nodes should remain in the cluster in the event of network split-brain. Veritas I/O fencing is allowed to decide on the surviving nodes first, followed by Oracle Clusterware. The CSS miss-count value indicates the amount of time Oracle Clusterware waits before evicting another node from the cluster, when it fails to respond across the interconnect.

For more information, see the Oracle Metalink document: 782148.1

## Planning the public network configuration for Oracle RAC

Identify separate public virtual IP addresses for each node in the cluster. Oracle requires one public virtual IP address for the Oracle listener process on each node. Public virtual IP addresses are used by client applications to connect to the Oracle database and help mitigate TCP/IP timeout delays. Oracle Clusterware manages the virtual IP addresses.

## Planning the private network configuration for Oracle RAC

Oracle RAC requires a minimum of one private IP address for Oracle Clusterware heartbeat.

Depending on the version of Oracle RAC you want to install, use one of the following options for setting up the private network configuration for Oracle database cache fusion:

Oracle RAC 10g    Use either Oracle UDP IPC or VCSIPC/LMX/LLT for the database cache fusion traffic.

By default, the database cache fusion traffic is configured to use VCSIPC/LMX/LLT.

Oracle RAC 11g    You must use UDP IPC for the database cache fusion traffic.

The Oracle UDP IPC protocol requires an IP address. Symantec recommends that an additional private IP address from a different subnet be used for load balancing the Oracle cache fusion traffic.

---

**Note:** Oracle UDP IPC is supported for both Oracle RAC 10g and Oracle RAC 11g with SF Oracle RAC.

---

---

**Note:** The private IP addresses of all nodes that are on the same physical network must be in the same IP subnet.

---

The following practices provide a resilient private network setup:

- Configure Oracle Clusterware interconnects over LLT links to prevent data corruption.

In an SF Oracle RAC cluster, the Oracle Clusterware heartbeat link **MUST** be configured as an LLT link. If Oracle Clusterware and LLT use different links for their communication, then the membership change between VCS and Oracle Clusterware is not coordinated correctly. For example, if only the Oracle Clusterware links are down, Oracle Clusterware kills one set of nodes after the expiry of the `css-miscount` interval and initiates the Oracle Clusterware and database recovery, even before CVM and CFS detect the node failures. This uncoordinated recovery may cause data corruption.
- Oracle Clusterware interconnects need to be protected against NIC failures and link failures. The PrivNIC or MultiPrivNIC agent can be used to protect against NIC failures and link failures, if multiple links are available. Even if link aggregation solutions in the form of bonded NICs are implemented, the PrivNIC or MultiPrivNIC agent can be used to provide additional protection against the failure of the aggregated link by failing over to available alternate links. These alternate links can be simple NIC interfaces or bonded NICs.

An alternative option is to configure the Oracle Clusterware interconnects over bonded NIC interfaces.

See [“High availability solutions for Oracle RAC private network”](#) on page 60.
- Configure Oracle Cache Fusion traffic to take place through the private network. Symantec also recommends that all UDP cache-fusion links be LLT links.

The PrivNIC and MultiPrivNIC agents provide a reliable alternative when operating system limitations prevent you from using NIC bonding to provide high availability and increased bandwidth using multiple network interfaces. In the event of a NIC failure or link failure, the agent fails over the private IP address from the failed link to the connected or available LLT link. To use multiple links for database cache fusion for increased bandwidth, configure the `cluster_interconnects` initialization parameter with multiple IP addresses for each database instance and configure these IP addresses under MultiPrivNIC for high availability.

Oracle database clients use the public network for database services. Whenever there is a node failure or network failure, the client fails over the connection,

for both existing and new connections, to the surviving node in the cluster with which it is able to connect. Client failover occurs as a result of Oracle Fast Application Notification, VIP failover and client connection TCP timeout. It is strongly recommended not to send Oracle Cache Fusion traffic through the public network.

- Use NIC bonding to provide redundancy for public networks so that Oracle can fail over virtual IP addresses if there is a public link failure.

### High availability solutions for Oracle RAC private network

Table 3-1 lists the high availability solutions that you may adopt for your private network.

**Table 3-1** High availability solutions for Oracle RAC private network

| Options                                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Using IPMP for Oracle Clusterware                             | If Oracle Clusterware interconnects are configured over IPMP, all the NICs that are configured under LLT must be configured under the IPMP group. In such a configuration, it is recommended not to manage these links using the PrivNIC/MultiPrivNIC agents.                                                                                                                                                                                                                                                                                                                      |
| Using link aggregation/<br>NIC bonding for Oracle Clusterware | <p>Use a native NIC bonding solution to provide redundancy, in case of NIC failures.</p> <p>Make sure that a link configured under a aggregated link or NIC bond is not configured as a separate LLT link.</p> <p>When LLT is configured over a bonded interface, do one of the following steps to prevent GAB from reporting jeopardy membership:</p> <ul style="list-style-type: none"><li>■ Configure an additional NIC under LLT in addition to the bonded NIC.</li><li>■ Add the following line in the <code>/etc/llttab</code> file:</li></ul> <pre>set-dbg-minlinks 2</pre> |

**Table 3-1** High availability solutions for Oracle RAC private network (*continued*)

| Options                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Using PrivNIC/MultiPrivNIC agents | <p>Use the PrivNIC agent when operating system limitations prevent you from using NIC bonding to provide high availability using multiple network interfaces.</p> <p>Use the MultiPrivNIC agent when operating system limitations prevent you from using NIC bonding to provide high availability and increased bandwidth using multiple network interfaces.</p> <p>For more deployment scenarios that illustrate the use of PrivNIC/MultiPrivNIC deployments, see the appendix "SF Oracle RAC deployment scenarios" in this document.</p> |

## Planning the storage

SF Oracle RAC provides the following options for shared storage:

- CVM  
CVM provides native naming as well as enclosure-based naming (EBN).  
Use enclosure-based naming for easy administration of storage.  
Enclosure-based naming guarantees that the same name is given to a shared LUN on all the nodes, irrespective of the operating system name for the LUN.

- CFS

- Oracle ASM over CVM

The following recommendations ensure better performance and availability of storage.

- Use multiple storage arrays, if possible, to ensure protection against array failures. The minimum recommended configuration is to have two HBAs for each host and two switches.
- Design the storage layout keeping in mind performance and high availability requirements. Use technologies such as striping and mirroring.
- Use appropriate stripe width and depth to optimize I/O performance.
- Use SCSI-3 persistent reservations (PR) compliant storage.
- Provide multiple access paths to disks with HBA/switch combinations to allow DMP to provide high availability against storage link failures and to provide load balancing.

## Planning the storage for SF Oracle RAC

Table 3-2 lists the type of storage required for SF Oracle RAC.

**Table 3-2** Type of storage required for SF Oracle RAC

| SF Oracle RAC files                     | Type of storage |
|-----------------------------------------|-----------------|
| SF Oracle RAC binaries                  | Local           |
| SF Oracle RAC fencing coordinator disks | Shared          |

## Planning the storage for Oracle RAC

Review the storage options and guidelines for Oracle RAC:

- Storage options for OCR and voting disk  
See “[Planning the storage for OCR and voting disk](#)” on page 62.
- Storage options for the Oracle RAC installation directories (ORACLE\_BASE, CRS\_HOME or GRID\_HOME (depending on Oracle RAC version), and ORACLE\_HOME)  
See “[Planning the storage for Oracle RAC binaries and data files](#)” on page 64.


### Planning the storage for OCR and voting disk

Depending on the Oracle RAC version, use one of the following options to place the OCR and voting disk information:

|                                                    |                                                                        |
|----------------------------------------------------|------------------------------------------------------------------------|
| Oracle RAC 10g Release 2/ Oracle RAC 11g Release 1 | Clustered File System<br>CVM raw volumes                               |
| Oracle RAC 11g Release 2                           | Clustered File System<br>ASM disk groups created using CVM raw volumes |

Figure 3-1 illustrates the options for storing OCR and voting disk information.

**Figure 3-1** OCR and voting disk storage options



- If you want to place OCR and voting disk on a clustered file system (option 1), you need to have two separate files for OCR and voting information respectively on CFS mounted on a CVM mirrored volume.
- If you want to place OCR and voting disk on CVM raw volumes or on ASM disk groups that use CVM raw volumes (option 2), you need to use two CVM mirrored volumes for configuring OCR and voting disk on these volumes.

For both option 1 and option 2:

- The installer needs at least two LUNs of 640 MB each for creating the OCR and voting disk storage.  
 Additionally, refer to the Oracle documentation for Oracle's recommendation on the required disk space for OCR and voting disk.
- The option **External Redundancy** must be selected at the time of installing Oracle Clusterware.

---

**Note:** Update the default setting of the CVM disk group fail policy to `leave` for OCR and voting disk. Retain the default setting (`global`) for the disk detach policy.

---

You can use the installer to create the storage for OCR and voting disk. This step is discussed later in the chapter "Installing Oracle RAC".

**Note:** For setting up replicated clusters, OCR and voting disk must be on non-replicated shared storage.

### Planning the storage for Oracle RAC binaries and data files

The Oracle RAC binaries can be stored on local storage or on shared storage, based on your high availability requirements.

**Note:** Symantec recommends that you install the Oracle Clusterware and Oracle database binaries local to each node in the cluster.

Consider the following points while planning the installation:

- Local installations provide improved protection against a single point of failure and also allows for applying Oracle RAC patches in a rolling fashion.
- CFS installations provide a single Oracle installation to manage, regardless of the number of nodes. This scenario offers a reduction in storage requirements and easy addition of nodes.

Table 3-3 lists the type of storage for Oracle RAC binaries and data files.

**Table 3-3** Type of storage for Oracle RAC binaries and data files

| Oracle RAC files                    | Type of storage                                                                                                                                                                                                                                  |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oracle base                         | Local<br><br>The ORACLE_BASE directory specified during the Oracle Clusterware/Grid Infrastructure installation is used to store the Oracle Inventory files. Oracle requires that the Oracle inventory files be placed on the local file system. |
| Oracle Clusterware                  | Local or shared                                                                                                                                                                                                                                  |
| Oracle Grid Infrastructure binaries | Local<br><br>Placing the Oracle Grid Infrastructure binaries on local disks enables rolling upgrade of the cluster.                                                                                                                              |
| Oracle database binaries            | Local or shared<br><br>Placing the Oracle database binaries on local disks enables rolling upgrade of the cluster.                                                                                                                               |



**Table 3-3**
Type of storage for Oracle RAC binaries and data files (*continued*)

| Oracle RAC files                                     | Type of storage                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Database datafiles                                   | <p>Shared</p> <p>Store the Oracle database files on CFS rather than on raw device or CVM raw device for easier management. Create separate clustered file systems for each Oracle database. Keeping the Oracle database datafiles on separate mount points enables you to unmount the database for maintenance purposes without affecting other databases.</p> <p><b>Note:</b> Set the CVM disk group fail policy to <code>leave</code> for disk groups containing the data files.</p> <p>If you plan to store the Oracle database on ASM, configure the ASM disk groups over CVM volumes to take advantage of dynamic multi-pathing.</p> |
| SF Oracle RAC database storage management repository | Shared                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Database recovery data (archive, flash recovery)     | <p>Shared</p> <p>Place archived logs on CFS rather than on local file systems.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Planning for Oracle ASM over CVM**

Review the following information on storage support provided by Oracle ASM:

|                  |                                                                                                                                                                                                |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Supported by ASM | ASM provides storage for data files, control files, online and archive log files, and backup files. Starting with Oracle RAC 11g Release 2, ASM also supports storage for OCR and voting disk. |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Not supported by ASM

Oracle RAC 11g Release 1 and earlier versions:

ASM does not support Oracle binaries, trace files, alert logs, export files, tar files, core files, Oracle Cluster Registry devices (OCR), and voting disk, and application binaries on ASM.

Oracle RAC 11g Release 2:

ASM does not support Oracle binaries, trace files, alert logs, export files, tar files, core files, and application binaries on ASM.

The following practices offer high availability and better performance:

- Use CVM mirrored volumes with dynamic multi-pathing for creating ASM disk groups. Select external redundancy while creating ASM disk groups.
- The CVM raw volumes used for ASM must be used exclusively for ASM. Do not use these volumes for any other purpose, such as creation of file systems. Creating file systems on CVM raw volumes used with ASM may cause data corruption.
- Do not enable ODM when databases are created on ASM. ODM is a disk management interface for data files that reside on the Veritas File System.
- Use a minimum of two Oracle ASM disk groups. Store the data files, one set of redo logs, and one set of control files on one disk group. Store the Flash Recovery Area, archive logs, and a second set of redo logs and control files on the second disk group.

For more information, see Oracle's ASM best practices document.

- Do not configure DMP meta nodes as ASM disks for creating ASM disk groups. Access to DMP meta nodes must be configured to take place through CVM.
- Do not combine DMP with other multipathing software in the cluster.
- Do not use coordinator disks, which are configured for I/O fencing, as ASM disks.
- Volumes presented to a particular ASM disk group should be of the same geometry, speed, and type.

## Planning volume layout

The following recommendations ensure optimal layout of VxVM/CVM volumes:

- Mirror the volumes across two or more storage arrays, if using VxVM mirrors.

- Separate the Oracle recovery structures from the database files to ensure high availability when you design placement policies.
- Separate redo logs and place them on the fastest storage (for example, RAID 1+ 0) for better performance.
- Use "third-mirror break-off" snapshots for cloning the Oracle log volumes. Do not create Oracle log volumes on a Space-Optimized (SO) snapshot.
- Create as many Cache Objects (CO) as possible when you use Space-Optimized (SO) snapshots for Oracle data volumes.
- Distribute the I/O load uniformly on all Cache Objects when you create multiple Cache Objects.
- If using VxVM mirroring, keep the Fast Mirror Resync region size equal to the database block size to reduce the copy-on-write (COW) overhead. Reducing the region size increases the amount of Cache Object allocations leading to performance overheads.
- Implement zoning on SAN switch to control access to shared storage. Be aware that physical disks may be shared by multiple servers or applications and must therefore be protected from accidental access.
- Choose DMP I/O policy based on the storage network topology and the application I/O pattern.
- Exploit thin provisioning for better return on investment.

## Planning file system design

The following recommendations ensure an optimal file system design for databases:

- If using VxVM mirroring, use ODM with CFS for better performance. ODM with SmartSync enables faster recovery of mirrored volumes using Oracle resilvering.
- Create separate file systems for Oracle binaries, data, redo logs, and archive logs. This ensures that recovery data is available if you encounter problems with database data files storage.
- Always place archived logs on CFS file systems rather than local file systems.

## About planning to configure I/O fencing

After you configure SF Oracle RAC with the installer, you must configure I/O fencing in the cluster for data integrity.

You can configure either disk-based I/O fencing or server-based I/O fencing. If your enterprise setup has multiple clusters that use VCS for clustering, Symantec recommends you to configure server-based I/O fencing.

The coordination points in server-based fencing can include only CP servers or a mix of CP servers and coordinator disks. Symantec also supports server-based fencing with a single coordination point which is a single highly available CP server that is hosted on an SFHA cluster.


---

**Warning:** For server-based fencing configurations that use a single coordination point (CP server), the coordination point becomes a single point of failure. In such configurations, the arbitration facility is not available during a failover of the CP server in the SFHA cluster. So, if a network partition occurs on any application cluster during the CP server failover, the application cluster is brought down.

---

[Figure 3-2](#) illustrates a high-level flowchart to configure I/O fencing for the SF Oracle RAC cluster.

**Figure 3-2** Workflow to configure I/O fencing




After you perform the preparatory tasks, you can use any of the following methods to configure I/O fencing:

|                                      |                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Using the installsfrac program       | <p>See <a href="#">“Setting up disk-based I/O fencing using installsfrac program”</a> on page 169.</p> <p>See <a href="#">“Setting up server-based I/O fencing using installsfrac program”</a> on page 184.</p>                                                                                                                                                                                              |
| Using response files                 | <p>See <a href="#">“Response file variables to configure disk-based I/O fencing”</a> on page 469.</p> <p>See <a href="#">“Response file variables to configure server-based I/O fencing”</a> on page 471.</p> <p>See <a href="#">“Response file variables to configure server-based I/O fencing”</a> on page 471.</p> <p>See <a href="#">“Configuring I/O fencing using response files”</a> on page 439.</p> |
| Manually editing configuration files | <p>See <a href="#">“Setting up disk-based I/O fencing manually”</a> on page 178.</p> <p>See <a href="#">“Setting up server-based I/O fencing manually”</a> on page 195.</p>                                                                                                                                                                                                                                  |

## Typical SF Oracle RAC cluster configuration with disk-based I/O fencing

[Figure 3-3](#) displays a typical VCS configuration with two nodes and shared storage. The configuration uses three coordinator disks for I/O fencing.


**Figure 3-3** Typical SF Oracle RAC cluster configuration with disk-based I/O fencing



## Typical SF Oracle RAC cluster configuration with server-based I/O fencing

Figure 3-4 displays a configuration using a SF Oracle RAC cluster (with two nodes), a single CP server, and two coordinator disks. The nodes within the SF Oracle RAC cluster are connected to and communicate with each other using LLT links.

**Figure 3-4** CP server, SF Oracle RAC cluster, and coordinator disks



## Recommended CP server configurations

Following are the recommended CP server configurations:

- Multiple application clusters use three CP servers as their coordination points.  
 See Figure 3-5 on page 72.
- Multiple application clusters use a single CP server and multiple pairs of coordinator disks (two) as their coordination points.  
 See Figure 3-6 on page 73.
- Multiple application clusters use a single CP server as their coordination point  
 This single coordination point fencing configuration must use a highly available CP server that is configured on an SFHA cluster as its coordination point.  
 See Figure 3-7 on page 73.

---

**Warning:** In a single CP server fencing configuration, arbitration facility is not available during a failover of the CP server in the SFHA cluster. So, if a network partition occurs on any application cluster during the CP server failover, the application cluster is brought down.

---

Although the recommended CP server configurations use three coordination points, you can use more than three (must be an odd number) coordination points for I/O fencing. In a configuration where multiple application clusters share a common set of CP server coordination points, the application cluster as well as the CP server use a Universally Unique Identifier (UUID) to uniquely identify an application cluster.

Figure 3-5 displays a configuration using three CP servers that are connected to multiple application clusters.

**Figure 3-5** Three CP servers connecting to multiple application clusters





Figure 3-6 displays a configuration using a single CP server that is connected to multiple application clusters with each application cluster also using two coordinator disks.




**Figure 3-6** Single CP server with two coordinator disks for each application cluster



[Figure 3-7](#) displays a configuration using a single CP server that is connected to multiple application clusters.

**Figure 3-7** Single CP server connecting to multiple application clusters



See [“Configuration diagrams for setting up server-based I/O fencing”](#) on page 802.

# Planning for cluster management

Table 3-4 lists the various agents supported in SF Oracle RAC installations for effective cluster management.

Table 3-4            List of agents

| Agent                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VCS agent for Oracle | <p>Oracle database management</p> <p>The VCS Oracle agent is recommended for managing Oracle databases. VCS controls the Oracle database in this configuration. The configuration without VCS Oracle agent may be used only in a single database setup.</p>                                                                                                                                                                                                                  |
| VCS agent for CVM    | <p>Volume management</p> <p>An SF Oracle RAC installation automatically configures the CVMCluster resource and the CVMVxconfigd resource. You must configure the CVMVolDg agent for each shared disk group.</p>                                                                                                                                                                                                                                                              |
| VCS agent for CFS    | <p>File system management</p> <p>If the database uses cluster file systems, configure the CFSMount agent for each volume in the disk group.</p>                                                                                                                                                                                                                                                                                                                              |
| CSSD agent           | <p>Oracle Clusterware management</p> <p>The CSSD agent starts, stops, and monitors Oracle Clusterware. It ensures that the OCR, the voting disk and the private IP address resources required by Oracle Clusterware are online before Oracle Clusterware starts.</p> <p><b>Note:</b> It is mandatory to use CSSD agent in SF Oracle RAC installations to ensure adequate handling of inter-dependencies and thereby prevent the premature startup of Oracle Clusterware.</p> |
| PrivNIC agent        | <p>High availability for a private IP address</p> <p>The PrivNIC agent provides a reliable alternative when operating system limitations prevent you from using NIC bonding to provide high availability using multiple network interfaces.</p>                                                                                                                                                                                                                              |
| MultiPrivNIC agent   | <p>High availability for multiple private IP addresses</p> <p>The MultiPrivNIC agent provides a reliable alternative when operating system limitations prevent you from using NIC bonding to provide high availability and increased bandwidth using multiple network interfaces.</p>                                                                                                                                                                                        |

**Table 3-4** List of agents (*continued*)

| Agent             | Description                                                                                                                                                                                                                                                                                                                                    |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CRSResource agent | <p>Oracle database management</p> <p>The CRSResource agent provides an alternative mechanism for monitoring the Oracle database in the absence of the VCS Oracle agent. It is useful in scenarios where the database is not managed by VCS and the applications need to be started using VCS after Oracle Clusterware starts the database.</p> |

## Planning for disaster recovery

SF Oracle RAC provides various disaster recovery configurations, such as campus clusters and global clusters, for multi-site clusters. In multi-site clusters, the nodes can be placed in different parts of a building, in separate buildings, or in separate cities. The distance between the nodes depends on the type of disaster from which protection is needed and on the technology used to replicate data. SF Oracle RAC supports various replication technologies for data replication.

To protect clusters against outages caused by disasters, the cluster components must be geographically separated.

### Planning a campus cluster setup

A campus cluster is also known as a stretch cluster or remote mirror configuration. In a campus cluster, the hosts and storage of a cluster span multiple sites separated by a few miles.

Bear in mind the following best practices mind when you configure a SF Oracle RAC campus cluster:

- Campus cluster sites are typically connected using a redundant high-capacity network that provides access to storage and private network communication between the cluster nodes. A single DWDM link can be used for both storage and private network communication.
- Tag the disks or enclosures that belong to a site with the corresponding VxVM site name. VxVM allocates storage from the correct site when creating or resizing a volume and when changing a volume's layout if the disks in the VxVM disk group that contain the volume are tagged with the site name.
- Tag each host with the corresponding VxVM site name. Make sure the read policy of the volumes is set to `SITEREAD`. This setting ensures that the reads on the volumes are satisfied from the local site's plex.

- Turn on the `allsites` attribute for all volumes that have data required by the application, to make sure they are evenly mirrored. Each site must have at least one mirror of all volumes hosting application data, including the FlashSnap log volume.
- In the case of a two-site campus cluster, place the third coordinator disk on the third site. You may use iSCSI disk on the third site as an alternative to Dark Fiber connected FC-SAN disk.  
Starting with version 5.1, you may use the software-based arbitration solution, Coordination Point Server (CPS), as a third coordination point.
- Turn on the `siteconsistent` attribute for the disk groups and the volumes to enable site-aware plex detaches. Snapshot volumes need not be site-consistent.
- Make sure that a DCO log version 20 or higher is attached to the volumes to enable Fast Resync operations.
- Set the CVM disk detach policy as "global" and the disk group fail policy as "leave" for all disk groups containing data volumes as well as disk groups containing the OCR and VOTE volumes.

## Planning a global cluster setup

Global clusters provide the ability to fail over applications between geographically distributed clusters when a disaster occurs.

Global clustering involves two steps:

1. Replication of data between the sites
2. Migration of the application when disaster occurs

The following aspects need to be considered when you design a disaster recovery solution:

- The amount of data lost in the event of a disaster (Recovery Point Objective)
- The acceptable recovery time after the disaster (Recovery Time Objective)

## Data replication considerations

When you choose a replication solution, one of the important factors that you need to consider is the required level of data throughput. Data throughput is the rate at which the application is expected to write data. The impact of write operations on replication are of more significance than the read operations.

In addition to the business needs discussed earlier, the following factors need to be considered while choosing the replication options:

- Mode of replication
- Network bandwidth
- Network latency between the two sites
- Ability of the remote site to keep up with the data changes at the first site

## Planning for Oracle VM Server for SPARC environments

The following best practices are recommended for setting up Oracle VM Server for SPARC environments for deploying SF Oracle RAC:

- Use a minimum of 4 CPU cores and 8 GB RAM for high Oracle workload environments.
- Domains that have direct I/O access are recommended for database applications. Although database applications can run on virtual devices, the performance may degrade.  
 If performance is not a critical concern, as in test deployments, use virtual disks for database applications.
- Distribute LDoms on separate physical machines for clustered applications as it provides high availability in case of physical hardware failure.



# Licensing SF Oracle RAC

This chapter includes the following topics:

- [About Veritas product licensing](#)
- [Setting or changing the product level for keyless licensing](#)
- [Installing Veritas product license keys](#)

## About Veritas product licensing

You have the option to install Veritas products without a license key. Installation without a license does not eliminate the need to obtain a license. A software license is a legal instrument governing the usage or redistribution of copyright protected software. The administrator and company representatives must ensure that a server or cluster is entitled to the license level for the products installed. Symantec reserves the right to ensure entitlement and compliance through auditing.

If you encounter problems while licensing this product, visit the Symantec licensing support website.

[www.symantec.com/techsupp/](http://www.symantec.com/techsupp/)

The Veritas product installer prompts you to select one of the following licensing methods:

- Install a license key for the product and features that you want to install.  
When you purchase a Symantec product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased.
- Continue to install without a license key.  
The installer prompts for the product modes and options that you want to install, and then sets the required product level.

Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled or continue with keyless licensing by managing the server or cluster with a management server. If you do not comply with the above terms, continuing to use the Veritas product is a violation of your end user license agreement, and results in warning messages. For more information about keyless licensing, see the following URL:  
<http://go.symantec.com/sfhakeyless>

If you upgrade to this release from a prior release of the Veritas software, the product installer does not change the license keys that are already installed. The existing license keys may not activate new features in this release.

If you upgrade with the product installer, or if you install or upgrade with a method other than the product installer, you must do one of the following to license the products:

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.  
See “[Setting or changing the product level for keyless licensing](#)” on page 82.  
See the `vxkeyless (1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the products you have purchased.  
See “[Installing Veritas product license keys](#)” on page 83.  
See the `vxlicinst (1m)` manual page.

You can also use the above options to change the product levels to another level that you are authorized to use. For example, you can add the replication option to the installed product. You must ensure that you have the appropriate license for the product level and options in use.

---

**Note:** In order to change from one product group to another, you may need to perform additional steps.

---

## About SF Oracle RAC licenses

[Table 4-1](#) lists the various SF Oracle RAC license levels in keyless licensing and the corresponding features.



**Table 4-1** SF Oracle RAC license levels (keyless licensing)

| License      | Description                                                       | Features enabled                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------|-------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SFRACENT     | SF Oracle RAC Enterprise Edition                                  | <p>The license enables the following features:</p> <ul style="list-style-type: none"> <li>■ Veritas Volume Manager</li> <li>■ Veritas Storage and Availability Management Tools for Oracle databases</li> <li>■ Veritas Extension for ODM</li> <li>■ Veritas File System</li> <li>■ Veritas Cluster Server</li> <li>■ Veritas Mapping Services</li> </ul>                                           |
| SFRACENT_VVR | SF Oracle RAC Enterprise Edition with VVR                         | <p>The license enables the following features:</p> <ul style="list-style-type: none"> <li>■ Veritas Volume Manager<br/>Veritas Volume Replicator is enabled.</li> <li>■ Veritas Storage and Availability Management Tools for Oracle databases</li> <li>■ Veritas Extension for ODM</li> <li>■ Veritas File System</li> <li>■ Veritas Cluster Server</li> <li>■ Veritas Mapping Services</li> </ul> |
| SFRACENT_GCO | SF Oracle RAC Enterprise Edition with Global Cluster Option (GCO) | <p>The license enables the following features:</p> <ul style="list-style-type: none"> <li>■ Veritas Volume Manager</li> <li>■ Veritas Storage and Availability Management Tools for Oracle databases</li> <li>■ Veritas Extension for ODM</li> <li>■ Veritas File System</li> <li>■ Veritas Cluster Server<br/>Global Cluster Option is enabled.</li> <li>■ Veritas Mapping Services</li> </ul>     |

Table 4-1 SF Oracle RAC license levels (keyless licensing) (continued)

| License          | Description                                       | Features enabled                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SFRACENT_VVR_GCO | SF Oracle RAC Enterprise Edition with VVR and GCO | <div>The license enables the following features:</div> <div><div>■ Veritas Volume Manager</div><div>Veritas Volume Replicator is enabled.</div><div>■ Veritas Storage and Availability Management Tools for Oracle databases</div><div>■ Veritas Extension for ODM</div><div>■ Veritas File System</div><div>■ Veritas Cluster Server</div><div>Global Cluster Option is enabled.</div><div>■ Veritas Mapping Services</div></div> |

# Setting or changing the product level for keyless licensing

The keyless licensing method uses product levels to determine the Veritas products and functionality that are licensed. In order to use keyless licensing, you must set up a Management Server to manage your systems.

For more information and to download the management server, see the following URL:

<http://go.symantec.com/vom>

When you set the product license level for the first time, you enable keyless licensing for that system. If you install with the product installer and select the keyless option, you are prompted to select the product and feature level that you want to license.

After you install, you can change product license levels at any time to reflect the products and functionality that you want to license. When you set a product level, you agree that you have the license for that functionality.

### To set or change the product level

- 1 View the current setting for the product level.

```
# vxkeyless -v display
```

- 2 View the possible settings for the product level.

```
# vxkeyless displayall
```

- 3 Set the desired product level.

```
# vxkeyless -q set prod_levels
```

where *prod\_levels* is a comma-separated list of keywords, as shown in step 2

If you want to remove keyless licensing and enter a key, you must clear the keyless licenses. Use the NONE keyword to clear all keys from the system.

---

**Warning:** Clearing the keys disables the Veritas products until you install a new key or set a new product level.

---

### To clear the product license level

- 1 View the current setting for the product license level.

```
# vxkeyless [-v] display
```

- 2 If there are keyless licenses installed, remove all keyless licenses:

```
# vxkeyless [-q] set NONE
```

For more details on using the `vxkeyless` utility, see the `vxkeyless(1m)` manual page.

## Installing Veritas product license keys

The VRTSvlic package enables product licensing. After the VRTSvlic is installed, the following commands and their manual pages are available on the system:

|                        |                                                                    |
|------------------------|--------------------------------------------------------------------|
| <code>vxlicinst</code> | Installs a license key for a Symantec product                      |
| <code>vxlicrep</code>  | Displays currently installed licenses                              |
| <code>vxlictest</code> | Retrieves features and their descriptions encoded in a license key |

Even though other products are included on the enclosed software discs, you can only use the Symantec software products for which you have purchased a license

**To install a new license**

- ◆ Run the following commands. In a cluster environment, run the commands on each node in the cluster:

```
# cd /opt/VRTS/bin
```

```
# ./vxlicinst -k xxxx-xxxx-xxxx-xxxx-xxxx-xxx
```

## Installation and configuration of SF Oracle RAC

- [Chapter 5. Preparing to install SF Oracle RAC](#)
- [Chapter 6. Installing SF Oracle RAC](#)
- [Chapter 7. Preparing to configure SF Oracle RAC](#)
- [Chapter 8. Configuring SF Oracle RAC](#)
- [Chapter 9. Configuring SF Oracle RAC clusters for data integrity](#)
- [Chapter 10. Performing post-installation and configuration tasks](#)



# Preparing to install SF Oracle RAC

This chapter includes the following topics:

- [Setting the umask before installation](#)
- [Synchronizing time settings on cluster nodes](#)
- [Mounting the product disc](#)
- [Setting up shared storage](#)
- [Setting the environment variables](#)
- [Optimizing LLT media speed settings on private NICs](#)
- [Verifying the systems before installation](#)

## Setting the umask before installation

Set the umask to provide appropriate permissions for SF Oracle RAC binaries and files. This setting is valid only for the duration of the current session.

```
# umask 0022
```

## Synchronizing time settings on cluster nodes

Make sure that the time settings on all cluster nodes are synchronized.

---

**Note:** For Oracle RAC 11g Release 2, it is mandatory to configure NTP for synchronizing time on all nodes in the cluster.

---

## Mounting the product disc

You must have superuser (root) privileges to load the SF Oracle RAC software.

You can unmount the product disc after completing the SF Oracle RAC installation.

### To mount the product disc

- 1 Log in as the superuser to a cluster node or a remote node in the same subnet as the cluster nodes.
- 2 Insert the product disc with the SF Oracle RAC software into a drive that is connected to the system.
- 3 If Solaris volume management software is running on your system, the software disc is automatically mounted.
- 4 If Solaris volume management software is not available to mount the disc, you must mount it manually. After inserting the software disc, enter:

```
# mount -F hsfs -o ro /dev/dsk/cXtXdXs2 /dvd_mount
```

Where c0t6d0s2 is the default address for the disc drive.

- 5 Navigate to the location of the packages:

```
# cd /dvd_mount/storage_foundation_for_oracle_rac
```

## Setting up shared storage

You need to set up shared storage to meet the following requirements:

- The LUNs from the shared storage must be visible to all the nodes in the cluster as seen by the following command:

```
# format
```

- The shared storage must support SCSI-3 persistent reservations (PR). Run the `vxfcntlsthaw(1M)` utility to ensure that the shared storage is SCSI-3 compliant. The `vxfcntlsthaw` utility is located in the directory

```
cluster_server/tools/vxfcntlsthaw.
```

```
# ./vxfcntlsthaw
```

For more information on setting up shared storage, see the *Veritas Cluster Server Installation Guide*.



## Setting the environment variables

Set the MANPATH variable in the .profile file (or other appropriate shell setup file for your system) to enable viewing of manual pages.

Based on the shell you use, type one of the following:

```
For sh, ksh, or bash    # MANPATH=/usr/share/man:/opt/VRTS/man
                        # export MANPATH
```

Set the PATH environment variable in the .profile file (or other appropriate shell setup file for your system) on each system to include installation and other commands.

Based on the shell you use, type one of the following:

```
For sh, ksh, or bash    # PATH=/usr/sbin:/sbin:/usr/bin:\
                        /usr/lib/vxvm/bin:/opt/VRTSvcs/bin:\
                        /opt/VRTS/bin:/opt/VRTSvcs/rac/bin:\
                        $PATH; export PATH
```

## Optimizing LLT media speed settings on private NICs

For optimal LLT communication among the cluster nodes, the interface cards on each node must use the same media speed settings. Also, the settings for the switches or the hubs that are used for the LLT interconnections must match that of the interface cards. Incorrect settings can cause poor network performance or even network failure.

If you use different media speed for the private NICs, Symantec recommends that you configure the NICs with lesser speed as low-priority links to enhance LLT performance.

## Guidelines for setting the media speed of the LLT interconnects

Review the following guidelines for setting the media speed of the LLT interconnects:

- Symantec recommends that you manually set the same media speed setting on each Ethernet card on each node.  
If you use different media speed for the private NICs, Symantec recommends that you configure the NICs with lesser speed as low-priority links to enhance LLT performance.

- If you have hubs or switches for LLT interconnects, then set the hub or switch port to the same setting as used on the cards on each node.
- If you use directly connected Ethernet links (using crossover cables), Symantec recommends that you set the media speed to the highest value common to both cards, typically 1000\_Full\_Duplex.

Details for setting the media speeds for specific devices are outside of the scope of this manual. Consult the device's documentation for more information.

## Verifying the systems before installation

Use any of the following options to verify your systems before installation:

- Option 1: Run the Veritas Operations Services (VOS) utility.  
For information on downloading and running VOS:  
<http://go.symantec.com/vos>
- Option 2: Run the installsfrac program with the "-precheck" option as follows:  
Navigate to the directory that contains the installsfrac program. The program is located in the product directory.  
Start the preinstallation check:

```
# ./installsfrac -precheck galaxy nebula
```

The program proceeds in a non-interactive mode, examining the systems for licenses, packages, disk space, and system-to-system communications. The program displays the results of the check and saves them in a log file. The location of the log file is displayed at the end of the precheck process.

- Option 3: Run the Veritas Web-based installation program as follows:

- Start the Web-based installer:

```
# ./webinstaller start
```

Paste the link in the address bar of the Web browser to access the installer.

- On the Select a task and a product page, select **Perform a Pre-installation check** from the **Task** drop-down list.  
Select the product from the **Product** drop-down list, and click **Next**.  
The installer performs the precheck and displays the results.

# Installing SF Oracle RAC

This chapter includes the following topics:

- [About installing SF Oracle RAC](#)
- [About installation and configuration methods](#)
- [Installing SF Oracle RAC using the Veritas script-based installation program](#)
- [Installing SF Oracle RAC using the Veritas Web-based installation program](#)
- [Installing SF Oracle RAC using Solaris JumpStart](#)
- [Installing SF Oracle RAC on an alternate root](#)

## About installing SF Oracle RAC

You can install SF Oracle RAC on clusters of up to 16 nodes.

The installation requires passwordless communication between systems for installing SF Oracle RAC. The installer prompts you to configure passwordless communication at the time of installation, if it is not already set up.

The product installer has improved ability to recover from failed installations. A recovery file is created if an installation fails due to a failed network connection. This file enables the install program to resume from the point where the installation failed.

By default, the installer generates a response file that you can reuse or customize for silent installations on other systems. If you need to suspend the installation, make sure that you select the option **Exit the installer** to create the response file before exiting the installation.

---

**Note:** If you enter **q** to quit the installation, the installer does not create a response file.

---

# About installation and configuration methods

You can use one of the following methods to install and configure SF Oracle RAC.

Table 6-1 Installation and configuration methods

| Method                                                                                                                                                                                                                                                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Interactive installation and configuration using the script-based installer</p> <p><b>Note:</b> If you obtained SF Oracle RAC from an electronic download site, you must use the <code>installsfrac programscript</code> instead of the <code>installer</code> script.</p> | <p>You can use one of the following script-based installers:</p> <ul style="list-style-type: none"><li>■ Common product installer script:<br/><code>installer</code><br/>The common product installer script provides a menu that simplifies the selection of installation and configuration options.<br/>Use this method to install other products, such as the Symantec Product Authentication Service (AT), along with SF Oracle RAC.</li><li>■ Product-specific installation script:<br/><code>installsfrac program</code></li><li>■ The product-specific installation script provides command-line interface options. Installing and configuring with the <code>installsfrac programscript</code> is identical to specifying SF Oracle RAC from the <code>installer</code> script.<br/>Use this method to install or configure only SF Oracle RAC.</li></ul> |
| <p>Silent installation using the response file</p>                                                                                                                                                                                                                            | <p>The response file automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. You can use the script-based installers with the response file to install silently on one or more systems.</p> <p>See <a href="#">“About response files”</a> on page 425.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 6-1 Installation and configuration methods (continued)

| Method              | Description                                                                                                                                                                                                                                                                                                  |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Web-based installer | <p>The Web-based installer provides an interface to manage the installation and configuration from a remote site using a standard Web browser.</p> <p><code>./webinstaller</code></p> <p>See “<a href="#">Installing SF Oracle RAC using the Veritas Web-based installation program</a>” on page 96.</p>     |
| JumpStart           | <p>You can use the Veritas product installer of the product-specific installation script to generate a JumpStart script file. Use the generated script to install Veritas packages from your JumpStart server.</p> <p>See “<a href="#">Installing SF Oracle RAC using Solaris JumpStart</a>” on page 99.</p> |

# Installing SF Oracle RAC using the Veritas script-based installation program

Before you start the installation, make sure that you have the installation and configuration worksheet with the values on hand for the installation.

During the installation, the installer performs the following tasks:

- Verifies system readiness for installation by checking system communication, network speed, installed packages, operating system patches, and available volume space.

---

**Note:** If the installer reports that any of the patches are not available, install the patches on the system before proceeding with the SF Oracle RAC installation.

---

- Installs the SF Oracle RAC 5.1 SP1 packages.

If you encounter issues during the installation, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*, Chapter "Performance and troubleshooting" for information on resolving the issue.

The following sample procedure installs SF Oracle RAC on two systems—galaxy and nebula.

**To install SF Oracle RAC**

- 1** Log in as the superuser on one of the systems.
- 2** Start the installation program:

|                         |                                                                                                                    |
|-------------------------|--------------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC installer | Navigate to the directory that contains the installation program. The program is located in the product directory. |
|-------------------------|--------------------------------------------------------------------------------------------------------------------|

Run the program:

```
# ./installsfprac galaxy nebula
```

|                          |                                                                   |
|--------------------------|-------------------------------------------------------------------|
| Common product installer | Navigate to the directory that contains the installation program. |
|                          | Run the program:                                                  |

```
# ./installer galaxy nebula
```

From the opening Selection Menu, choose: **"I"** for "Install a Product."

From the displayed list of products to install, choose **Veritas Storage Foundation for Oracle RAC**.

The installer displays the copyright message and specifies the directory where the running logs are created.

- 3** Set up the systems so that commands between systems execute without prompting for passwords or confirmations.
- 4** If you had quit the installer in the process of an active installation, the installer discovers that installer process and provides the option of resuming the installation or starting a new installation. Provide a suitable response.

```
The installer has discovered an existing installer process.
The process exited while performing configure of SF Oracle RAC
on galaxy.
Do you want to resume this process? [y,n,q,?] (y) n
```

- 5** Enter **y** to agree to the End User License Agreement (EULA).
- 6** Select the type of installation—Minimal, Recommended, All. Symantec recommends you to choose option **3 Install all Veritas Storage Foundation for Oracle RAC packages**.

```
1) Install minimal required Veritas Storage Foundation for
Oracle RAC packages
```

```

2)  Install recommended Veritas Storage Foundation for
    Oracle RAC packages
3)  Install all Veritas Storage Foundation for
    Oracle RAC packages
4)  Display packages to be installed for each option
Select the packages to be installed on all systems? [1-4,q,?] (2) 3

```

The installer verifies the systems for compatibility and displays the list of packages and patches that will be installed.

The installer installs the SF Oracle RAC packages and patches.

## 7 Select the appropriate license option.

```

1) Enter a valid license key
2) Enable keyless licensing
How would you like to license the systems? [1-2,q]

```

- Enter **1** if you have a valid license key. When prompted, enter the license key.

```

Enter a SF Oracle RAC license key:
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-X

```

If you plan to enable additional capabilities, enter the corresponding license keys when you are prompted for additional licenses.

```

Do you wish to enter additional licenses? [y,n,q,b] (n)

```

- Enter **2** to enable keyless licensing.

---

**Note:** The keyless license option enables you to install SF Oracle RAC without entering a key. However, you must still acquire a valid license to install and use SF Oracle RAC. Keyless licensing requires that you manage the systems with a Management Server.

---

Enter **y** if you want to configure Veritas Volume Replicator (VVR) and Global Cluster Option (GCO) during the installation. VVR is configured with default values while GCO is configured with the settings you specify. You can reconfigure VVR and GCO manually at any time.

```

Would you like to enable the Veritas Volume Replicator? [y,n,q] y

```

```

Would you like to enable the Global Cluster Option? [y,n,q] y

```

The installer registers the license.

- 8 Verify that the installation process completed successfully. Review the output at the end of the installation and note the location of the summary and log files for future reference.
- 9 Enter **y** to configure SF Oracle RAC:

---

**Note:** If you plan to configure clusters in secure mode or set up server-based fencing, enter **n** at the prompt. You need to complete the preparatory tasks before starting the SF Oracle RAC configuration. For instructions, see the chapter "Preparing to configure SF Oracle RAC" in this document.

---

```
Would you like to configure SF Oracle RAC on galaxy nebula [y,n,q] (n) y
```

Proceed to configure the SF Oracle RAC components.

- 10 Enter **y** if you want to send the installation information to Symantec.

```
Would you like to send the information about this installation  
to Symantec to help improve installation in the future? [y,n,q,?] (y) y
```

## Installing SF Oracle RAC using the Veritas Web-based installation program

Before you start the installation, make sure that you have the installation and configuration worksheet with the values on hand for the installation.

The installation of SF Oracle RAC using the Web-based installation program involves the following tasks:

- Review the system configuration requirements for a Web-based installation. See [“Before using the Veritas Web-based installer”](#) on page 97.
- Start the Veritas Web installer. See [“Starting the Veritas Web-based installer”](#) on page 97.
- Install SF Oracle RAC. See [“Installing products with the Veritas Web-based installer”](#) on page 98.

During the installation, the installer performs the following tasks:

- Verifies system readiness for installation by checking system communication, network speed, installed packages, operating system patches, and available volume space.



---

**Note:** If the installer reports that any of the patches are not available, install the patches on the system before proceeding with the SF Oracle RAC installation.

---

- Installs the SF Oracle RAC 5.1 SP1 packages.

If you encounter issues during the installation, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*, Chapter "Performance and troubleshooting" for information on resolving the issue.

## Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

**Table 6-2** Web-based installer requirements

| System                | Function                                                                                                        | Requirements                                                                                                                                             |
|-----------------------|-----------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Target system         | The systems where you plan to install the Veritas products.                                                     | Must be a supported platform for SF Oracle RAC 5.1 SP1.                                                                                                  |
| Installation server   | The server where you start the installation. The installation media is accessible from the installation server. | Must use the same operating system as the target systems and must be at one of the supported operating system update levels.                             |
| Administrative system | The system where you run the Web browser to perform the installation.                                           | Must have a Web browser.<br>Supported browsers: <ul style="list-style-type: none"> <li>■ Internet Explorer 6, 7, and 8</li> <li>■ Firefox 3.x</li> </ul> |

## Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

### To start the Web-based installer

- 1 Start the Veritas XPortal Server process `xprtlwid`, on the installation server:

```
# ./webinstaller start
```

The webinstaller script displays a URL. Note this URL.

---

**Note:** If you do not see the URL, run the command again.

---

- 2 On the administrative server, start the Web browser.
- 3 Navigate to the URL that the script displayed.
- 4 The browser may display the following message:  
  

```
Secure Connection Failed
```

  
Obtain a security exception for your browser.
- 5 When prompted, enter `root` and root's password of the installation server.

## Installing products with the Veritas Web-based installer

This section describes installing SF Oracle RAC with the Veritas Web-based installer.

### To install SF Oracle RAC

- 1 Perform preliminary steps.
- 2 Start the Web-based installer.  
See [“Starting the Veritas Web-based installer”](#) on page 97.
- 3 Select **Veritas Storage Foundation for Oracle RAC** from the Product drop-down list, and click **Next**.
- 4 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 5 Choose minimal, recommended, or all packages. Click **Next**.
- 6 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.
- 7 After the validation completes successfully, click **Next** to install SF Oracle RAC on the selected system.
- 8 After the installation completes, you must choose your licensing method.

On the license page, select one of the following tabs:

■ Keyless licensing

---

**Note:** The keyless license option enables you to install without entering a key. However, in order to ensure compliance you must manage the systems with a management server.

For more information, go to the following website:

<http://go.symantec.com/sfhakeyless>

---

Complete the following information:

Choose whether you want to enable Global Cluster option.

Click Register.

■ Enter license key

If you have a valid license key, select this tab. Enter the license key for each system. Click **Register**.

- 9 The installer prompts you to configure the cluster. Select Yes to continue with configuring the product.

If you select No, you can exit the installer. You must configure the product before you can use SF Oracle RAC.

After the installation completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

- 10 Select the checkbox to specify whether you want to send your installation information to Symantec.

Would you like to send the information about this installation to Symantec to help improve installation in the future?

Click **Finish**. The installer prompts you for another task.

## Installing SF Oracle RAC using Solaris JumpStart

This section provides instructions for installing SF Oracle RAC using Solaris JumpStart. The instructions assume a working knowledge of Solaris JumpStart. See the operating system documentation for detailed information on using Solaris JumpStart.

---

**Note:** Only new installations of SF Oracle RAC are supported using Solaris JumpStart.

---

Before you perform the instructions in this section, complete the preparatory tasks for installing SF Oracle RAC.

## Task overview for SF Oracle RAC installation using JumpStart

This section provides a summary of the tasks for installing SF Oracle RAC using Solaris JumpStart.

1. Set up a central Solaris JumpStart server on the network.  
For instructions, see the Solaris JumpStart documentation.
2. Add the systems, on which you want to install SF Oracle RAC, as clients to the JumpStart server.  
For instructions, see the Solaris JumpStart documentation.
3. Prepare the installation resources.  
See [“Preparing the JumpStart installation resources”](#) on page 100.
4. Install and configure SF Oracle RAC.  
See [“Installing and configuring SF Oracle RAC using JumpStart”](#) on page 108.

## Preparing the JumpStart installation resources

This section contains instructions for creating the installation resources.

[Table 6-3](#) lists the installation resources you must prepare before you install SF Oracle RAC using Solaris JumpStart.

**Table 6-3** Installation resources

| Files          | Description                                                                                                                                                                                                        |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Finish scripts | You need to generate the following finish scripts: <ul style="list-style-type: none"><li>■ <code>jumpstart_sfrac51sp1.fin</code> (Required)</li><li>■ <code>encap_bootdisk_vm51sp1.fin</code> (Optional)</li></ul> |
| Response files | You need to create empty response files for the following packages:<br><br><code>VRTSaslapm, VRTSvxvm</code>                                                                                                       |
| admin file     | You need to create an admin file if you plan to perform a non-interactive installation.                                                                                                                            |
| rules file     | You need to modify the rules file as appropriate for your systems.                                                                                                                                                 |

Table 6-4 lists the sample directories used in the procedure.

**Table 6-4** Sample directories used in the procedure

| Files                                        | Sample directories                                                                                                                                                                 |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC product disc content           | /export/config                                                                                                                                                                     |
| Installation and finish scripts              | /export/config                                                                                                                                                                     |
| Solaris 9 packages and patches               | <p>Solaris 9 packages:</p> <p>/export/pkgs</p> <p>Solaris 9 patches:</p> <p>/export/patches</p> <p><b>Note:</b> No additional packages or patches are required for Solaris 10.</p> |
| Response files for installation              | /export/config/pkgs                                                                                                                                                                |
| Admin file for non-interactive installations | /export/config/pkgs                                                                                                                                                                |

**Note:** The directories must be mounted as NFS-accessible directories to the JumpStart server.

### To prepare the installation resources

- 1 Copy the packages from the product disc to the Solaris JumpStart server under a shared directory. The packages are in .pkg format and the patches are in the .tar.gz format. Uncompress the patches before you proceed otherwise the installation fails.
  - First, create directories for installation.
 

```
# mkdir /export/config
```
  - Insert the product disc into a drive that is connected to the system. The Solaris volume management software automatically mounts the disc as /dvdrom/dvd1. Type the command:
 

```
# cd /dvd_mount/
```
  - Copy the package files and patches to the server.

```
# cp -r * /export/config
```

- 2 Create response files for the SF Oracle RAC packages.  
See [“Creating JumpStart response files”](#) on page 102.
- 3 For non-interactive installations, create the file `admin` in the current directory on your JumpStart server (`/export/config/pkgs`), and modify the file as follows:

```
mail=
instance=overwrite
partial=nocheck
runlevel=quit
idepend=quit
rdepend=nocheck
space=quit
setuid=nocheck
conflict=nocheck
action=nocheck
basedir=default
```

---

**Note:** Specify the `-a adminfile` option with the `pkgadd` command in the finish script you generate in the next step.

---

- 4 Generate the installation and finish scripts.  
See [“Generating the JumpStart installation and finish scripts”](#) on page 103.
- 5 Modify the rules file as required.

For example:

```
any - - profile_sfrac jumpstart_sfrac51spl.fin
```

If you generated the root disk encapsulation finish file:

```
any - - profile_sfrac encap_bootdisk_vm51spl.fin
```

For detailed instructions, see the Solaris JumpStart documentation.

## Creating JumpStart response files

Response files contain the installation profile for packages. Some packages need empty response files, while some packages require specific settings.

---

**Note:** Make sure that you edit the `finish` script to use the `-r` option with the `pkgadd` command to install the packages using the corresponding response files:

```
# pkgadd -r responsefile_name package_name
```

For example, to install the `VRTSvxvm` package using its response file:

```
# pkgadd -r VRTSvxvm.responsefile -d VRTSvxvm.pkg
```

---

The sample procedure places the response files in the directory `/export/config/pkgs`.

### To create response files

- 1 Change to the directory `/export/config/pkgs`.

```
# cd /export/config/pkgs
```

- 2 For Solaris SPARC systems: Create a response file for each of the following packages: `VRTSaslapm`, `VRTSvxvm`

```
# pkgask -r package_name.response -d package_name.pkg
```

For example:

```
# pkgask -r VRTSvxvm.responsefile -d VRTSvxvm.pkg
```

## Generating the JumpStart installation and finish scripts

Run the SF Oracle RAC installer to generate the installation and finish scripts.

The installer generates the following scripts:

|                                         |                                                         |
|-----------------------------------------|---------------------------------------------------------|
| <code>jumpstart_sfrac51sp1.fin</code>   | Finish script for installing SF Oracle RAC              |
| <code>encap_bootdisk_vm51sp1.fin</code> | Encapsulation finish script for root disk encapsulation |
| <code>installsfrac</code>               | Installer script for configuring SF Oracle RAC          |
| <code>uninstallsfrac</code>             | Installer script for uninstalling SF Oracle RAC         |

### To generate the JumpStart installation and finish scripts

- 1 Run the SF Oracle RAC installer to generate the installation and finish scripts:

```
# cd /dvd_mount/storage_foundation_for_oracle_rac
# ./installsfrac -jumpstart dir_path
```

Where *dir\_path* is the full path to the directory where the scripts are placed.

For example:

```
# ./installsfrac -jumpstart /export/config/
```

- 2 If you want to encapsulate the root disk automatically, generate a separate finish script for root disk encapsulation.

Enter **y** to generate a sample finish script for root disk encapsulation.

```
Would you like to generate the finish script to encapsulate
the boot disk? [y,n,q,?] (y)
```

Enter the disk group name, private region length, and the disk media name of the root disk to be encapsulated.

```
Specify the disk group name of the root disk to be encapsulated: rootdg
Specify the private region length of the root disk
to be encapsulated: (65536)
Specify the disk media name of the root disk to
be encapsulated: (rootdg_01)
```

- 3 View the list of generated scripts.

```
# ls /export/config
```

The following scripts will be listed:

```
encap_bootdisk_vm51sp1.fin installsfrac
jumpstart_sfrac51sp1.fin uninstallsfrac
```

The root disk encapsulation script will be listed only if you chose to encapsulate the root disk automatically.

- 4 Modify the finish files, as required.

You will need to update the following information in the finish file:

- Installation order for packages and patches  
The finish script must contain the correct order of the SF Oracle RAC packages and patches and the operating system packages and patches.



See [“Installation order of packages and patches for JumpStart”](#) on page 106. Use the list of packages and patches that is generated to replace the package list in the finish scripts.

- **Package and patch information for Solaris 9 installations**  
For Solaris 9 installations, provide the necessary package and patch information.  
See [“Installation order of packages and patches for JumpStart”](#) on page 106.

- **BUILDSRC value**  
The path indicated in the BUILDSRC variable must contain the product disc content. It must also contain the product installation and uninstallation scripts (generated in step 1 and step 2 of this procedure) if you want to configure or uninstall SF Oracle RAC from the `/opt/VRTS/install` directory. Otherwise, you need to configure and uninstall SF Oracle RAC from the product disc.  
The value must be in the following format:  
*hostname\_or\_ip:/path\_to\_pkgs\_patches\_scripts*  
For example:

```
192.168.12.1:/export/config
```

- **ENCAPSRC value**  
The path indicated in the ENCAPSRC variable must contain the root encapsulation finish script.  
The value must be in the following format:  
*hostname\_or\_ip:/path\_to\_encap\_script*  
For example:

```
192.168.12.1:/export/config
```

- **License information for root disk encapsulation**  
If you want the root disk to be encapsulated, you must provide the Veritas Volume Manager license information in the root disk encapsulation finish file.
- **Language pack information**  
If you want to install SF Oracle RAC in a language other than English, add the language pack information to the basic finish file.  
See [“Adding language pack information to the JumpStart finish file”](#) on page 107.

For a basic sample finish file:

See [“Sample JumpStart finish file \(basic installation\)”](#) on page 108.

For a sample root disk encapsulation finish file:  
See [“Sample JumpStart finish file \(for root encapsulation\)”](#) on page 113.

## Installation order of packages and patches for JumpStart

The correct installation order of packages and patches for JumpStart is as follows:

SF Oracle RAC packages and patches

The correct order of SF Oracle RAC packages and patches can be viewed by running the `installsfprac` program with one of the following options:

■

-minpkgs

Install SF Oracle RAC with basic functionality.

■

-recpkgs

Installs the full feature set without optional packages.

■

-allpkgs

Installs all available packages.  
Symantec recommends installation of all the packages.

For example, to view the installation order for installing SF Oracle RAC 5.1 SP1:

```
# cd /dvd_mount/storage_foundation_for_oracle_rac
# ./installsfprac -allpkgs
```

Operating system packages and patches

See [Table 6-5](#).

[Table 6-5](#) lists the correct order for the Solaris 9 operating system packages and patches.

**Table 6-5** Solaris 9 operating system packages and patches

| Operating system version | Installation order of packages and patches                                                                                                                                                                                                |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Solaris 9 packages       | SUNWsan, SUNWcfpl, SUNWcfplx, SUNWcfclr, SUNWcfcl, SUNWcfclx, SUNWfchbr, SUNWfchba, SUNWfchbx, SUNWfcsn, SUNWfcsnx, SUNWmdiu, SUNWqlc, SUNWqlcx, SUNWjfca, SUNWjfcax, SUNWjfcay, SUNWjfcayx, SUNWemlxs, SUNWemlxsx, SUNWemlxu, SUNWemlxux |

**Table 6-5** Solaris 9 operating system packages and patches (*continued*)

| Operating system version   | Installation order of packages and patches                                                                                                                                                                   |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Solaris 9 Update 7 patches | 117171-17, 113073-14, 112908-33, 118558-39, 122300-10, 111847-08, 113046-01 113049-01, 113039-20, 113040-25, 113041-14, 113042-19 113043-15, 113044-07, 114476-09, 114477-04, 114478-08 114878-10, 119914-13 |
| Solaris 9 Update 9 patches | 112908-33, 118558-39, 122300-10, 111847-08, 113046-01 113049-01, 113039-20, 113040-25, 113041-14, 113042-19 113043-15, 113044-07, 114476-09, 114477-04, 114478-08 114878-10, 119914-13                       |

**Note:** No additional packages or patches are required for Solaris 10 Update 6 and later update versions.

## Adding language pack information to the JumpStart finish file

Perform the steps in the following procedure to install SF Oracle RAC in a language other than English. The packages need not be manually ordered as there are no package inter-dependencies.

### To add language pack information to the finish file

- 1 Copy the packages from the language pack installation disc.

```
# cp -r pkgs/* /export/config/pkgs
```

- 2 Add the following lines in the finish script.

```
for PKG in VRTSjacse VRTSjadbe VRTSmulic \
VRTSatJA VRTSjacs VRTSjafs \
VRTSatZH VRTSjacsu VRTSjaodm VRTSzhvm \
VRTSjacav VRTSjadba VRTSjavm
do
<...language pack instructions>
done
```

For sample finish file:

See [“Sample JumpStart finish file \(basic installation\)”](#) on page 108.

## Installing and configuring SF Oracle RAC using JumpStart

Perform the steps in the following procedure to install SF Oracle RAC using Solaris JumpStart.

### To install and configure SF Oracle RAC using JumpStart on Solaris SPARC systems

- 1 On each client node, run the following command to install the SF Oracle RAC packages:

```
ok> boot net - install
```

For Solaris x64 systems, press F12 and select the network boot mode.

The system is restarted after the packages are installed. If you choose to encapsulate the root disk on your systems, the systems start with an encapsulated root disk.

- 2 Configure SF Oracle RAC.

---

**Note:** Before you start the configuration, complete the preparatory tasks. For instructions, see the chapter *Preparing to configure SF Oracle RAC* in this document.

---

```
# /opt/VRTS/install/installsfrc -configure
```

For instructions on configuring SF Oracle RAC, see the following chapters in this document:

*Configuring SF Oracle RAC*

*Configuring SF Oracle RAC clusters for data integrity*

- 3 Complete the post-installation tasks.

For instructions, see the chapter *Performing post-installation and configuration tasks* in this document.

## Sample JumpStart finish file (basic installation)

The following extract is a sample finish file generated using the `installsfrc` program for SF Oracle RAC installations on Solaris 9.

The text in bold indicates modifications required for installing SF Oracle RAC.

```
#!/bin/sh
```

```
# $Copyright: Copyright (c) 2009 Symantec Corporation.  
# All rights reserved.
```

```
#
# THIS SOFTWARE CONTAINS CONFIDENTIAL INFORMATION AND TRADE SECRETS OF
# SYMANTEC CORPORATION. USE, DISCLOSURE OR REPRODUCTION IS PROHIBITED
# WITHOUT THE PRIOR EXPRESS WRITTEN PERMISSION OF SYMANTEC CORPORATION.
#
# The Licensed Software and Documentation are deemed to be commercial
# computer software as defined in FAR 12.212 and subject to restricted
# rights as defined in FAR Section 52.227-19 "Commercial Computer
# Software - Restricted Rights" and DFARS 227.7202, "Rights in
# Commercial Computer Software or Commercial Computer Software
# Documentation", as applicable, and any successor regulations. Any use,
# modification, reproduction release, performance, display or disclosure
# of the Licensed Software and Documentation by the U.S. Government
# shall be solely in accordance with the terms of this Agreement.  $0

echo "==== Executing finish script: $me ===="

PATH=$PATH:/sbin:/usr/sbin
export PATH

#
# Notice:
# * Please modify the BUILDSRC and ENCAPSRC below according to your
# * real environment
# * The location specified with BUILDSRC and ENCAPSRC should be NFS
# * accessible to the Jumpstart Server
# * It's required to set ENCAPSRC only if you are using jumpstart for
# * automatic boot disk encapsulation
# * Please copy the whole directories of pkgs and scripts from
# * installation media to the BUILDSRC
# * Please create the admin and response file for pkgadd according
# * to 'jumpstart_readme.txt' in the DVD
# * Please put generated install<product> scripts into the BUILDSRC
#

# The ENCAPSRC variable is required only if you want to encapsulate
# the root disk.
# The PKGSAN and PATCHSAN variables indicate the full path to the
# SAN packages and patches required for installing SF Oracle RAC.
BUILDSRC="192.168.12.1:/export/config"
ENCAPSRC="192.168.12.1:/export/config"
PKGSAN="192.168.12.1:/export/pkgs"
PATCHSAN="192.168.12.1:/export/patches"
```

```
# Scripts below are no need to change

ROOT=/a

# The MNT and MNT1 variables indicate the local mount points for the
# SAN packages and patches.
MNT=/a/pkg
MNT1=/a/patch
BUILDDIR="${ROOT}/build"
PKGDIR="${BUILDDIR}/pkgs"
# PATCHDIR="${BUILDDIR}/patches"
SCRIPTDIR="${BUILDDIR}/scripts"
CPIDIR="${ROOT}/opt/VRTS/install"
CPIMODDIR="${ROOT}/opt/VRTSperl/lib/site_perl/UXRT51"
ENCAPDIR="${BUILDDIR}/encap_script"

mkdir -p ${BUILDDIR}
mount -F nfs -o vers=3 ${BUILDSRC} ${BUILDDIR}

# The following procedure creates the mount directories and adds
the packages and patches.
mkdir ${MNT}
mkdir ${MNT1}
mount -F nfs -o vers=3 ${PKGDIR} ${MNT}
mount -F nfs -o vers=3 ${PATCHDIR} ${MNT1}
for pkg in SUNWsan SUNWcflr SUNWcflx SUNWcflr SUNWcflr \
SUNWcflx SUNWfchbr SUNWfchba SUNWfchbx SUNWfcsn SUNWfcsn \
SUNWmdiu SUNWqlc SUNWqlcx SUNWjfca SUNWjfcax SUNWjfcax \
SUNWjfcax SUNWemlxs SUNWemlxsx SUNWemlxu SUNWemlxu
do
    pkgadd -n -a ${MNT}/admin -d ${MNT}/${pkg} -R ${ROOT} ${pkg}
done
for patch in `cat ${MNT1}/patch_order`
do
    patchadd -u -d -R ${ROOT} ${MNT1}/${patch}
done

for PKG in VRTSvlic VRTSperl VRTSspt VRTSvxvm VRTSaslapm VRTSob \
VRTSsfmh VRTSvxfs VRTSfssdk VRTSat VRTS1lt VRTSgab VRTSvxfen \
VRTSamf VRTSvcs VRTScps VRTSvcsag VRTSvcsea VRTSdbed VRTSglm \
VRTScavf VRTSgms VRTSodm VRTSdbac;
do
```

```
RESP="${PKGDIR}/${PKG}.response"
echo "Installing package -- $PKG"
if [ -f $RESP ]
then
    pkgadd -n -a ${PKGDIR}/admin -d ${PKGDIR}/${PKG}.pkg -r $RESP \
-R ${ROOT} ${PKG}
else
    pkgadd -v -a ${PKGDIR}/admin -d ${PKGDIR}/${PKG}.pkg -R ${ROOT} \
${PKG}
fi
done

for PATCH in 142629-06 143687-02 142634-05 142633-05 143282-01 143281-01 \
143284-01 143283-01 145455-01 143287-01 143289-01 143288-01 143290-01 \
142631-04 143681-01 143680-01 143274-05 143273-05 143678-01 143677-01 \
143271-05 143270-05 145451-01 145450-01
do
    patchadd -R ${ROOT} -M ${PATCHDIR} ${PATCH}
done

# Required for language package installation
for PKG in VRTSjacse VRTSjadbe VRTSmulic \
VRTSatJA VRTSjacs VRTSjafs \
VRTSatZH VRTSjacsu VRTSjaodm VRTSzhvm \
VRTSjacav VRTSjadba VRTSjavm
do
    echo "Installing package -- $PKG"
    pkgadd -v -a ${PKGDIR}/admin -d ${PKGDIR}/${PKG}.pkg -R ${ROOT} ${PKG}
done

# copy cpi perl modules
mkdir -p ${CPIMODDIR}
for MODULE in EDR CPIC CPIP
do
    if [ -d "${SCRIPTDIR}/${MODULE}" ]
    then
        cp -rp ${SCRIPTDIR}/${MODULE} ${CPIMODDIR}
        echo "Perl modules from ${SCRIPTDIR}/${MODULE} are copied
to ${CPIMODDIR}"
    else
        echo "Required perl modules ${SCRIPTDIR}/${MODULE} not found"
    fi
done
```

```

done

# copy cpi message catalogs
if [ -d "${SCRIPTDIR}/messages" ]
then
    cp -rp "${SCRIPTDIR}/messages" ${CPIMODDIR}
    echo "Messages from ${SCRIPTDIR}/messages are copied to ${CPIMODDIR}"
fi
if [ -d "${SCRIPTDIR}/bin" ]
then
    cp -rp "${SCRIPTDIR}/bin" ${CPIMODDIR}
    echo "Commands from ${SCRIPTDIR}/bin are copied to ${CPIMODDIR}"
fi

# copy cpi installer scripts
mkdir -p ${CPIDIR}
touch "${CPIDIR}/.cpi5"
for FILE in `ls ${BUILDDIR} 2>/dev/null | grep 'install'`
do
    cp -p ${BUILDDIR}/${FILE} ${CPIDIR}
    echo "Copy ${BUILDDIR}/${FILE} to ${CPIDIR}"
done

touch ${ROOT}/noautoshutdown

umount ${BUILDDIR}

mkdir -p ${ENCAPDIR}
mount -F nfs -o vers=3 ${ENCAPSRC} ${ENCAPDIR}

cp ${ENCAPDIR}/encap_bootdisk_vm51.fin ${ROOT}/etc/init.d/vxvm-jumpstart
ln ${ROOT}/etc/init.d/vxvm-jumpstart ${ROOT}/etc/rc2.d/S99vxvm-jumpstart
chmod 755 ${ROOT}/etc/init.d/vxvm-jumpstart

# Unmount the
umount ${MNT}
umount ${MNT1}
rmdir ${MNT}
rmdir ${MNT1}
echo "==== Completed finish script $me ====="

exit 0

```



## Sample JumpStart finish file (for root encapsulation)

The following extract is a sample finish file generated using the `installsffrac` program for SF Oracle RAC installations on Solaris 9 for encapsulating the root disk.

The text in bold indicates the license key required for installing SF Oracle RAC.

---

**Note:** Do not modify the other statements in the script.

---

```
#!/bin/sh

# Copyright: Copyright (c) 2009 Symantec Corporation.
# All rights reserved.
#
# THIS SOFTWARE CONTAINS CONFIDENTIAL INFORMATION AND TRADE SECRETS OF
# SYMANTEC CORPORATION. USE, DISCLOSURE OR REPRODUCTION IS PROHIBITED
# WITHOUT THE PRIOR EXPRESS WRITTEN PERMISSION OF SYMANTEC CORPORATION.
#
# The Licensed Software and Documentation are deemed to be commercial
# computer software as defined in FAR 12.212 and subject to restricted
# rights as defined in FAR Section 52.227-19 "Commercial Computer
# Software - Restricted Rights" and DFARS 227.7202, "Rights in
# Commercial Computer Software or Commercial Computer Software
# Documentation", as applicable, and any successor regulations. Any use,
# modification, reproduction release, performance, display or disclosure
# of the Licensed Software and Documentation by the U.S. Government
# shall be solely in accordance with the terms of this Agreement.

#####
#
# The following init script encapsulates the root disk.
# The script was copied to the /etc/rc2.d directory remotely
# as part of the vxvm jumpstart installation procedure.
#
#####

: ${VOLROOT_DIR:=${__VXVM_ROOT_DIR}}
. ${VOL_SCRIPTS_LIB:-/usr/lib/vxvm/lib}/vxcommon

CMD=`basename $0`

quit()
```

```

{
    code=$1
    if [ -n "$DEBUG" ]; then
        set -x
    fi
    rm -f /etc/init.d/vxvm-jumpstart /etc/rc2.d/S01vxvm-jumpstart
    if [ "$code" -eq 100 ]; then
        shutdown -g0 -y -i6
        code=0
    fi
    exit $code
}

trap 'quit 2' INT HUP QUIT TERM

if [ -n "$DEBUG" ]; then
    set -x
fi

# if system is already encapsulated, then exit init script
df / | grep rootvol > /dev/null
if [ $? -eq 0 ]; then
    echo "INFO: $CMD: system is already encapsulated."
    quit 0
fi

# Do minimal vxvm installation
if [ -d /dev/vx/dmp ]
then
    /sbin/mount -F tmpfs dmpfs /dev/vx/dmp
fi
if [ -d /dev/vx/rdmp ]
then
    /sbin/mount -F tmpfs dmpfs /dev/vx/rdmp
fi

# set the license for vxconfigd to work
mount /opt 2> /dev/null
/opt/VRTS/bin/vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXX

vxconfigd -k -m disable > /dev/null 2>&1
vxdctl init > /dev/null 2>&1
vxdctl enable

```

```
voldmode=`vxctl mode 2>/dev/null`
if [ "X$voldmode" != "Xmode: enabled" ]
then
    echo "ERROR: $CMD: vold could not be enabled."
    quit 1
fi

rm -f $mkdbfile

# Determine root disk of system
set_rootdisk
if [ -z "$rootdisk" ]; then
    echo "ERROR: $CMD: Could not locate root disk : $rootdisk."
    quit 2
fi

# Encapsulate root disk
/usr/lib/vxvm/bin/vxencap -c -g rootdg -f sliced -s 65536 \
rootdg_01=$rootdisk

# Exit if encapsulation of root disk failed
if [ ! -s /etc/vx/reconfig.d/disk.d/$rootdisk/newpart ]
then
    echo "ERROR: $CMD: Encapsulation of root disk failed."
    quit 3
fi


# encapsulation was successful. Shutdown the system to complete
# encapsulation.
quit 100
```

## Installing SF Oracle RAC on an alternate root

Installing SF Oracle RAC on an alternate root enables you to boot from the second disk instead of the default disk. Installing on an alternate root also enables you to upgrade the OS on a Solaris system without affecting the existing configuration or requiring much downtime. Using an alternate root is required when using Live Upgrade to upgrade to SF Oracle RAC 5.1 SP1.

[Figure 6-1](#) illustrates the installation process.

**Figure 6-1** Installing SF Oracle RAC on an alternate root disk



Perform the steps in the following procedure on the active root disk of each node.

**To install SF Oracle RAC 5.1 SP1 on the alternate root disk of your system**

- 1 Verify that the Solaris operating system is installed on the alternate root disk of the system.

For example, `/dev/dsk/cXtXdXs2`, where `cXtXdXs2` is the alternate root disk.

- 2 Mount your alternate root disk.

```
# mkdir /altroot
# mount /dev/dsk/cXtXdXs0 /altroot
```

The mount point must have the same name on all systems.

- 3 Start the installer with the `-rootpath` option.

```
# cd /dvd_mount/storage_foundation_for_oracle_rac
# ./installsfprac -rootpath /altroot galaxy nebula
```

- 4 Stop the applications that are running on the current root disk using native application commands.

- 5 Restart the systems with the alternate root `/altroot`.

- Display the current boot disk device and device aliases.

```
# eeprom
boot-device=vx-rootdg vx-int_disk
use-nvramrc?=true
```

```
nvrsrc=devalias vx-int_disk /pci@1c,600000/scsi@2/disk@0,0:a  
devalias vx-rootdg01 /pci@1c,600000/scsi@2/disk@1,0:a
```

- Set the device from which to boot using the `eeprom` command. This example shows booting from the alternate root disk.

```
# eeprom boot-device=/pci@780/pci@0/pci@9/scsi@0/disk@1
```

- Reboot the system.

```
# shutdown -g0 -i6 -y
```

## 6 Configure SF Oracle RAC.

```
# ./installsfrc -configure galaxy nebula
```

## 7 Install the language packages if you would like to run SF Oracle RAC in a language other than English. Follow the procedure for the appropriate language packages.

See [“Installing language packages”](#) on page 207.

## 8 Complete the post-installation tasks.

For instructions, see the chapter *Performing post-installation and configuration tasks* in this document.



# Preparing to configure SF Oracle RAC

This chapter includes the following topics:

- [Preparing to configure the clusters in secure mode](#)
- [Setting up the CP server](#)

## Preparing to configure the clusters in secure mode

You can set up Symantec Product Authentication Service (AT) for the cluster during or after the SF Oracle RAC configuration.

In a cluster that is online, if you want to enable or disable AT using the `installsf -security` command, see the *Veritas Cluster Server Administrator's Guide* for instructions.

The prerequisites to configure a cluster in secure mode are as follows:

- A system in your enterprise that serves as root broker (RB).

You can either use an external system as root broker, or use one of the cluster nodes as root broker.

- To use an external root broker, identify an existing root broker system in your enterprise or install and configure root broker on a stable system. See [“Installing the root broker for the security infrastructure”](#) on page 123.
- To use one of the cluster nodes as root broker, the installer does not require you to do any preparatory tasks.

When you configure the cluster in secure mode using the script-based installer, choose the automatic mode and choose one of the nodes for the installer to configure as root broker.

Symantec recommends that you configure a single root broker system for your entire enterprise. If you use different root broker systems, then you must establish trust between the root brokers.

For example, if the management server and the cluster use different root brokers, then you must establish trust.

- For external root broker, an authentication broker (AB) account for each node in the cluster is set up on the root broker system.  
See [“Creating authentication broker accounts on root broker system”](#) on page 124.
- The system clocks of the external root broker and authentication brokers must be in sync.

The script-based installer provides the following configuration modes:

|                     |                                                                                                                                                                                                                                                                                                     |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Automatic mode      | The external root broker system must allow rsh or ssh passwordless login to use this mode.                                                                                                                                                                                                          |
| Semi-automatic mode | <p>This mode requires encrypted files (BLOB files) from the AT administrator to configure a cluster in secure mode.</p> <p>The nodes in the cluster must allow rsh or ssh passwordless login.</p> <p>See <a href="#">“Setting up inter-system communication”</a> on page 702.</p>                   |
| Manual mode         | <p>This mode requires root_hash file and the root broker information from the AT administrator to configure a cluster in secure mode.</p> <p>The nodes in the cluster must allow rsh or ssh passwordless login.</p> <p>See <a href="#">“Setting up inter-system communication”</a> on page 702.</p> |

[Figure 7-1](#) depicts the flow of configuring SF Oracle RAC cluster in secure mode.



Figure 7-1 Workflow to configure SF Oracle RAC cluster in secure mode




Table 7-1 lists the preparatory tasks in the order which the AT and VCS administrators must perform. These preparatory tasks apply only when you use an external root broker system for the cluster.

**Table 7-1** Preparatory tasks to configure a cluster in secure mode (with an external root broker)

| Tasks                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Who performs this task |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| <p>Decide one of the following configuration modes to set up a cluster in secure mode:</p> <ul style="list-style-type: none"> <li>■ Automatic mode</li> <li>■ Semi-automatic mode</li> <li>■ Manual mode</li> </ul>                                                                                                                                                                                                                                                                                                                  | VCS administrator      |
| <p>Install the root broker on a stable system in the enterprise.</p> <p>See <a href="#">“Installing the root broker for the security infrastructure”</a> on page 123.</p>                                                                                                                                                                                                                                                                                                                                                            | AT administrator       |
| <p>To use the semi-automatic mode or the manual mode, on the root broker system, create authentication broker accounts for each node in the cluster.</p> <p>See <a href="#">“Creating authentication broker accounts on root broker system”</a> on page 124.</p> <p>The AT administrator requires the following information from the VCS administrator:</p> <ul style="list-style-type: none"> <li>■ Node names that are designated to serve as authentication brokers</li> <li>■ Password for each authentication broker</li> </ul> | AT administrator       |
| <p>To use the semi-automatic mode, create the encrypted files (BLOB files) for each node and provide the files to the VCS administrator.</p> <p>See <a href="#">“Creating encrypted files for the security infrastructure”</a> on page 125.</p> <p>The AT administrator requires the following additional information from the VCS administrator:</p> <ul style="list-style-type: none"> <li>■ Administrator password for each authentication broker</li> </ul> <p>Typically, the password is the same for all nodes.</p>            | AT administrator       |
| <p>To use the manual mode, provide the root_hash file (/opt/VRTSat/bin/root_hash) from the root broker system to the VCS administrator.</p>                                                                                                                                                                                                                                                                                                                                                                                          | AT administrator       |

**Table 7-1** Preparatory tasks to configure a cluster in secure mode (with an external root broker) *(continued)*

| Tasks                                                                                                                                                                                                                                                        | Who performs this task |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| Copy the files that are required to configure a cluster in secure mode to the system from where you plan to install and configure SF Oracle RAC.<br><br>See <a href="#">“Preparing the installation system for the security infrastructure”</a> on page 127. | VCS administrator      |

## Installing the root broker for the security infrastructure

Install the root broker only if you plan to use AT to configure the cluster in secure mode. You can use a system outside the cluster or one of the systems within the cluster as root broker. If you plan to use an external broker, the root broker administrator must install and configure the root broker before you configure the Authentication Service for SF Oracle RAC. Symantec recommends that you install the root broker on a stable system that is outside the cluster.

You can also identify an existing root broker system in the data center to configure the cluster in secure mode. The root broker system can run AIX, HP-UX, Linux, or Solaris operating system.

See Symantec Product Authentication Service documentation for more information.

See [“About Symantec Product Authentication Service \(AT\)”](#) on page 36.

### To install the root broker

- 1 Mount the product disc and start the installer.  
  
`# ./installer`
- 2 From the Task Menu, choose I for "Install a Product."
- 3 From the displayed list of products to install, choose: Symantec Product Authentication Service (AT).
- 4 Enter **y** to agree to the End User License Agreement (EULA).
- 5 Enter 2 to install the recommended packages.
- 6 Enter the name of the system where you want to install the Root Broker.

Enter the `operating_system` system names separated by space [q,?]: **venus**

- 7 Review the output as the installer does the following:
  - Checks to make sure that AT supports the operating system
  - Verifies that you install from the global zone
  - Checks if the packages are already on the system.

The installer lists the packages that the program is about to install on the system. Press Enter to continue.

- 8 Review the output as the installer installs the root broker on the system.
- 9 After the installation, configure the root broker.
- 10 Select a mode to configure the root broker from the three choices that the installer presents:

```
1)Root+AB Mode
2)Root Mode
3)AB Mode
```

```
Enter the mode in which you would like AT to be configured? [1-3,q] 2
```

```
All AT processes that are currently running must be stopped
```

```
Do you want to stop AT processes now? [y,n,q,?] (y)
```

- 11 Press Enter to continue and review the output as the installer starts the Authentication Service.

## Creating authentication broker accounts on root broker system

On the root broker system, the administrator must create an authentication broker (AB) account for each node in the cluster.

### To create authentication broker accounts on root broker system

- 1 Determine the root broker domain name. Enter the following command on the root broker system:

```
venus> # vssat showalltrustedcreds
```

For example, the domain name resembles "Domain Name:  
root@venus.symantecexample.com" in the output.

- 2 For each node in the cluster, verify whether an account exists on the root broker system.

For example, to verify that an account exists for node galaxy:

```
venus> # vssat showprpl --pdrtype root \
--domain root@venus.symantecexample.com --prplname galaxy
```

- If the output displays the principal account on root broker for the authentication broker on the node, then delete the existing principal accounts. For example:

```
venus> # vssat deleteprpl --pdrtype root \
--domain root@venus.symantecexample.com \
--prplname galaxy --silent
```

- If the output displays the following error, then the account for the given authentication broker is not created on this root broker:

```
"Failed To Get Attributes For Principal"
```

Proceed to step 3.

- 3 Create a principal account for each authentication broker in the cluster. For example:

```
venus> # vssat addprpl --pdrtype root --domain \
root@venus.symantecexample.com --prplname galaxy \
--password password --prpltype service
```

You must use this password that you create in the input file for the encrypted file.

## Creating encrypted files for the security infrastructure

Create encrypted files (BLOB files) only if you plan to choose the semiautomatic mode that uses an encrypted file to configure the Authentication Service. The administrator must create the encrypted files on the root broker node. The administrator must create encrypted files for each node that is going to be a part of the cluster before you configure the Authentication Service for SF Oracle RAC.

### To create encrypted files

- 1 Make a note of the following root broker information. This information is required for the input file for the encrypted file:

|      |                                                                                                                       |
|------|-----------------------------------------------------------------------------------------------------------------------|
| hash | The value of the root hash string, which consists of 40 characters. Execute the following command to find this value: |
|------|-----------------------------------------------------------------------------------------------------------------------|

```
venus> # vssat showbrokerhash
```

**root\_domain**                      The value for the domain name of the root broker system. Execute the following command to find this value:

```
venus> # vssat showalltrustedcreds
```

- 2    Make a note of the following authentication broker information for each node. This information is required for the input file for the encrypted file:

**identity**                      The value for the authentication broker identity, which you provided to create authentication broker principal on the root broker system.

This is the value for the **--prplname** option of the **addprpl** command.

See [“Creating authentication broker accounts on root broker system”](#) on page 124.

**password**                      The value for the authentication broker password, which you provided to create authentication broker principal on the root broker system.

This is the value for the **--password** option of the **addprpl** command.

See [“Creating authentication broker accounts on root broker system”](#) on page 124.

- 3    For each node in the cluster, create the input file for the encrypted file.

The installer presents the format of the input file for the encrypted file when you proceed to configure the Authentication Service using encrypted file. For example, the input file for authentication broker on galaxy resembles:

```
[setuptrust]
broker=venus.symantecexample.com
hash=758a33dbd6fae751630058ace3dedb54e562fe98
securitylevel=high
```

```
[configab]
identity=galaxy
password=password
root_domain=root@venus.symantecexample.com
root_broker=venus.symantecexample.com:2821
```

```
start_broker=false
enable_pbx=false
```

- 4 Back up these input files that you created for the authentication broker on each node in the cluster.

Note that for security purposes, the command to create the output file for the encrypted file deletes the input file.

- 5 For each node in the cluster, create the output file for the encrypted file from the root broker system using the following command:

```
RootBroker> # vssat createpkg \
--in /path/to/blob/input/file.txt \
--out /path/to/encrypted/blob/file.txt \
--host_ctx AB-hostname
```

For example:

```
venus> # vssat createpkg --in /tmp/galaxy.blob.in \
--out /tmp/galaxy.blob.out --host_ctx galaxy
```

Note that this command creates an encrypted file even if you provide wrong password for "password=" entry. But such an encrypted file with wrong password fails to install on authentication broker node.

- 6 After you complete creating the output files for the encrypted file, you must copy these encrypted BLOB files for each node in the cluster.

## Preparing the installation system for the security infrastructure

The VCS administrator must gather the required information and prepare the installation system to configure a cluster in secure mode.

### To prepare the installation system for the security infrastructure

- ◆ Depending on the configuration mode you decided to use, do one of the following:

Automatic mode    Do the following:

- Gather the root broker system name from the AT administrator.
- During SF Oracle RAC configuration, choose the configuration option 1 when the installsfrac program prompts.

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Semi-automatic mode | <div>Do the following:</div> <ul style="list-style-type: none"><li>■ Copy the encrypted files (BLOB files) to the system from where you plan to install VCS.<br/>Note the path of these files that you copied to the installation system.</li><li>■ During SF Oracle RAC configuration, choose the configuration option 2 when the installsfrac program prompts.</li></ul>                                                                                                                                                                                                                                                                                                                                 |
| Manual mode         | <div>Do the following:</div> <ul style="list-style-type: none"><li>■ Copy the root_hash file that you fetched to the system from where you plan to install VCS.<br/>Note the path of the root hash file that you copied to the installation system.</li><li>■ Gather the root broker information such as name, fully qualified domain name, domain, and port from the AT administrator.</li><li>■ Note the principal name and password information for each authentication broker that you provided to the AT administrator to create the authentication broker accounts.</li><li>■ During SF Oracle RAC configuration, choose the configuration option 3 when the installsfrac program prompts.</li></ul> |

## Setting up the CP server

[Table 7-2](#) lists the tasks to set up the CP server for server-based I/O fencing.

**Table 7-2** Tasks to set up CP server for server-based I/O fencing

| Task                                             | Reference                                                                               |
|--------------------------------------------------|-----------------------------------------------------------------------------------------|
| Plan your CP server setup                        | See <a href="#">“Planning your CP server setup”</a> on page 129.                        |
| Install the CP server                            | See <a href="#">“Installing the CP server using the installer”</a> on page 130.         |
| Configure the CP server cluster in secure mode   | See <a href="#">“Configuring the CP server cluster in secure mode”</a> on page 131.     |
| Set up shared storage for the CP server database | See <a href="#">“Setting up shared storage for the CP server database”</a> on page 132. |



**Table 7-2** Tasks to set up CP server for server-based I/O fencing (*continued*)

| Task                               | Reference                                                                                                                                                                     |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the CP server            | See “ <a href="#">Configuring the CP server using the configuration utility</a> ” on page 133.<br><br>See “ <a href="#">Configuring the CP server manually</a> ” on page 141. |
| Verify the CP server configuration | See “ <a href="#">Verifying the CP server configuration</a> ” on page 143.                                                                                                    |

## Planning your CP server setup

Follow the planning instructions to set up CP server for server-based I/O fencing.

### To plan your CP server setup

- 1 Decide whether you want to host the CP server on a single-node VCS cluster, or on an SFHA cluster.

Symantec recommends hosting the CP server on an SFHA cluster.

- 2 If you host the CP server on an SFHA cluster, review the following information. Make sure you make the decisions and meet these prerequisites when you set up the CP server:

- You must configure fencing in enabled mode during the SFHA configuration.
- You must set up shared storage for the CP server database during your CP server setup.
- Decide whether you want to configure server-based fencing for the SF Oracle RAC cluster (application cluster) with a single CP server as coordination point or with at least three coordination points.  
Symantec recommends using at least three coordination points.

- 3 Decide whether you want to configure the CP server cluster in secure mode using the Symantec Product Authentication Service (AT).

Symantec recommends configuring the CP server cluster in secure mode. Setting up AT secures the communication between the CP server and its clients (SF Oracle RAC clusters). It also secures the HAD communication on the CP server cluster, and makes the authentication broker highly available.

- 4 Set up the hardware and network for your CP server.

See “[CP server requirements](#)” on page 48.

**5** Have the following information handy for CP server configuration:

- **Name for the CP server**  
The CP server name should not contain any special characters.
- **Port number for the CP server**  
Allocate a TCP/IP port for use by the CP server.  
Valid port range is between 49152 and 65535. The default port number is 14250.
- **Virtual IP address, network interface, netmask, and networkhosts for the CP server**

---

**Note:** If a CP server is configured to use an IPv6 network address for the virtual IP, then the operating system of the cluster nodes where SF Oracle RAC is hosted should also support the IPv6 network stack along with the IPv4 network stack. Note that even if SF Oracle RAC supports IPv6, it cannot be configured to use IPv6 network addresses for its agents, because Oracle RAC does not support IPv6.

---

## Installing the CP server using the installer

Perform the following procedure to install and configure VCS or SFHA on CP server systems.

**To install and configure VCS or SFHA on the CP server systems**

- ◆ Depending on whether your CP server uses a single system or multiple systems, perform the following tasks:

CP server setup uses a single system

Install and configure VCS to create a single-node VCS cluster.

Meet the following requirements for CP server:

- During installation, make sure to select all packages for installation. The VRTScps package is installed only if you select to install all packages.
- During configuration, make sure to configure LLT and GAB.
- During configuration, set up the cluster in secure mode if you want secure communication between the CP server and the SF Oracle RAC cluster (application cluster).

See the *Veritas Cluster Server Installation Guide* for instructions on installing and configuring VCS.

Proceed to configure the CP server.

See [“Configuring the CP server using the configuration utility”](#) on page 133.

See [“Configuring the CP server manually”](#) on page 141.

CP server setup uses multiple systems

Install and configure SFHA to create an SFHA cluster. This makes the CP server highly available.

Meet the following requirements for CP server:

- During installation, make sure to select all packages for installation. The VRTScps package is installed only if you select to install all packages.
- During configuration, set up the cluster in secure mode if you want secure communication between the CP server and the SF Oracle RAC cluster (application cluster).

See [“Preparing to configure the clusters in secure mode”](#) on page 119.

- During configuration, configure disk-based fencing (scsi3 mode).

See the *Veritas Storage Foundation and High Availability Installation Guide* for instructions on installing and configuring SFHA.

Proceed to set up shared storage for the CP server database.

## Configuring the CP server cluster in secure mode

You must configure security on the CP server only if you want to secure the communication between the CP server and the SF Oracle RAC cluster (CP client).

This step secures the HAD communication on the CP server cluster, and makes the authentication broker highly available.

---

**Note:** If you already configured Symantec Product Authentication Service (AT) during VCS configuration, you can skip this section.

---

### To configure the CP server cluster in secure mode

- ◆ Run the installer as follows to configure the CP server cluster in secure mode:

```
# installsfrac -security
```

See [“Preparing to configure the clusters in secure mode”](#) on page 119.

## Setting up shared storage for the CP server database

Symantec recommends that you create a mirrored volume for the CP server database and that you use the vxfs file system type.

If you configured SFHA on the CP server cluster, perform the following procedure to set up shared storage for the CP server database.

### To set up shared storage for the CP server database

- 1 Create a disk group containing the disks. You require two disks to create a mirrored volume.

For example:

```
# vxdg init cps_dg disk1 disk2
```

- 2 Import the disk group if it's not already imported.

For example:

```
# vxdg import cps_dg
```

### 3 Create a mirrored volume over the disk group.

For example:

```
# vxassist -g cps_dg make cps_vol volume_size layout=mirror
```

### 4 Create a file system over the volume.

The CP server configuration utility only supports vxfs file system type. If you use an alternate file system, then you must configure CP server manually.

Depending on the operating system that your CP server runs, enter the following command:

```
AIX # mkfs -V vxfs /dev/vx/rdisk/cps_dg/cps_volume
```

```
HP-UX # mkfs -F vxfs /dev/vx/rdisk/cps_dg/cps_volume
```

```
Linux # mkfs -t vxfs /dev/vx/rdisk/cps_dg/cps_volume
```

```
Solaris # mkfs -F vxfs /dev/vx/rdisk/cps_dg/cps_volume
```

## Configuring the CP server using the configuration utility

The CP server configuration utility (`configure_cps.pl`) is part of the VRTScps package.

Perform one of the following procedures:

|                                            |                                                                                            |
|--------------------------------------------|--------------------------------------------------------------------------------------------|
| For CP servers on single-node VCS cluster: | See <a href="#">“To configure the CP server on a single-node VCS cluster”</a> on page 133. |
|--------------------------------------------|--------------------------------------------------------------------------------------------|

|                                    |                                                                                  |
|------------------------------------|----------------------------------------------------------------------------------|
| For CP servers on an SFHA cluster: | See <a href="#">“To configure the CP server on an SFHA cluster”</a> on page 137. |
|------------------------------------|----------------------------------------------------------------------------------|

### To configure the CP server on a single-node VCS cluster

- 1 Verify that the VRTScps package is installed on the node.
- 2 Run the CP server configuration script on the node where you want to configure the CP server:

```
# /opt/VRTScps/bin/configure_cps.pl
```

- 3** Enter **1** at the prompt to configure CP server on a single-node VCS cluster. The configuration utility then runs the following preconfiguration checks:

- Checks to see if a single-node VCS cluster is running with the supported platform.  
The CP server requires VCS to be installed and configured before its configuration.
- Checks to see if the CP server is already configured on the system.  
If the CP server is already configured, then the configuration utility informs the user and requests that the user unconfigure the CP server before trying to configure it.

- 4** Enter the name of the CP server.

```
Enter the name of the CP Server: mycps1.symantecexample.com
```

- 5** Enter a valid virtual IP address on which the CP server process should depend on.

```
Enter a valid Virtual IP address on which  
the CP Server process should depend on:  
10.209.83.85
```

- 6** Enter the CP server port number or press Enter to accept the default value (14250).

```
Enter a port number in range [49152, 65535], or  
press <enter> for default port (14250):
```

- 7** Choose whether the communication between the CP server and the SF Oracle RAC clusters has to be made secure.

If you have not configured the CP server cluster in secure mode, enter **n** at the prompt.

---

**Warning:** If the CP server cluster is not configured in secure mode, and if you enter y, then the script immediately exits. You must configure the CP server cluster in secure mode and rerun the CP server configuration script.

---

Veritas recommends secure communication between the CP server and application clusters. Enabling security requires Symantec Product Authentication Service to be installed and configured on the cluster.

Do you want to enable Security for the communications? (y/n)  
(Default:y) :

- 8** Enter the absolute path of the CP server database or press Enter to accept the default value (/etc/VRTScps/db).

CP Server uses an internal database to store the client information.

Note: As the CP Server is being configured on a single node VCS, the database can reside on local file system.

Enter absolute path of the database (Default:/etc/VRTScps/db):

- 9** Verify and confirm the CP server configuration information.

Following is the CP Server configuration information:

```
-----
(a)CP Server Name: mycps1.symantecexample.com
(b)CP Server Virtual IP: 10.209.83.85
(c)CP Server Port: 14250
(d)CP Server Security : 1
(e)CP Server Database Dir: /etc/VRTScps/db
-----
```

Press b if you want to change the configuration, <enter> to continue :

- 10** The configuration utility proceeds with the configuration process, and creates a vxcps.conf configuration file.

```
Successfully generated the /etc/vxcps.conf configuration file.  
Successfully created directory /etc/VRTScps/db.
```

```
Configuring CP Server Service Group (CPSSG) for this cluster  
-----
```

```
NOTE: Please ensure that the supplied network interface is a  
public NIC
```

- 11** Enter a valid network interface for the virtual IP address for the CP server process.

```
Enter a valid network interface for virtual IP 10.209.83.85  
on mycps1.symantecexample.com: bge0
```

- 12** Enter networkhosts information for the NIC resource.

```
Symantec recommends configuring NetworkHosts attribute to ensure  
NIC resource to be online always.  
Do you want to add NetworkHosts attribute for the NIC resource bge0 on  
system mycps1? [y/n] : y  
Enter a valid IP address to configure NetworkHosts for NIC bge0 on  
system mycps1 : 10.209.83.86  
Do you want to add another Network Host ?[y/n] : n
```

- 13** Enter the netmask for the virtual IP address. For example:

```
Enter the netmask for virtual IP 10.209.83.85 :  
255.255.252.0
```



- 14** After the configuration process has completed, a success message appears. For example:

```
Successfully added the CPSSG service group to  
VCS configuration. Bringing the CPSSG service  
group online. Please wait...
```

```
The Veritas Coordination Point Server has been  
configured on your system.
```

- 15** Run the `hagrp -state` command to ensure that the CPSSG service group has been added.

For example:

```
# hagrp -state CPSSG
```

```
#Group      Attribute      System                               Value  
CPSSG       State         mycps1.symantecexample.com         |ONLINE|
```

It also generates the configuration file for CP server (`/etc/vxcps.conf`).

The configuration utility adds the `vxcperv` process and other resources to the VCS configuration in the CP server service group (CPSSG).

For information about the CPSSG, refer to the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

In addition, the `main.cf` samples contain details about the `vxcperv` resource and its dependencies.

See [“Sample configuration files for CP server”](#) on page 693.

### To configure the CP server on an SFHA cluster

- 1** Verify that the VRTScps package is installed on each node.
- 2** Make sure that you have configured passwordless ssh or rsh on the CP server cluster nodes.
- 3** Run the CP server configuration script on the node where you want to configure the CP server:

```
# /opt/VRTScps/bin/configure_cps.pl [-n]
```

The CP server configuration utility uses ssh by default to communicate between systems. Use the `-n` option for rsh communication.

- 4** Enter **2** at the prompt to configure CP server on an SFHA cluster.

The configuration utility then runs the following preconfiguration checks:

- Checks to see if an SFHA cluster is running with the supported platform. The CP server requires SFHA to be installed and configured before its configuration.
- Checks to see if the CP server is already configured on the system. If the CP server is already configured, then the configuration utility informs the user and requests that the user unconfigure the CP server before trying to configure it.

**5** Enter the name of the CP server.

```
Enter the name of the CP Server: mycps1.symantecexample.com
```

**6** Enter a valid virtual IP address on which the CP server process should depend on.

```
Enter a valid Virtual IP address on which  
the CP Server process should depend on:  
10.209.83.85
```

**7** Enter the CP server port number or press Enter to accept the default value (14250).

```
Enter a port number in range [49152, 65535], or  
press <enter> for default port (14250):
```

**8** Choose whether the communication between the CP server and the SF Oracle RAC clusters has to be made secure.

If you have not configured the CP server cluster in secure mode, enter **n** at the prompt.

---

**Warning:** If the CP server cluster is not configured in secure mode, and if you enter **y**, then the script immediately exits. You must configure the CP server cluster in secure mode and rerun the CP server configuration script.

---

Veritas recommends secure communication between the CP server and application clusters. Enabling security requires Symantec Product Authentication Service to be installed and configured on the cluster.

```
Do you want to enable Security for the communications? (y/n)  
(Default:y) :
```

## 9 Enter the absolute path of the CP server database or press Enter to accept the default value (/etc/VRTScps/db).

CP Server uses an internal database to store the client information.

Note: As the CP Server is being configured on SFHA cluster, the database should reside on shared storage with vxfs file system.

Please refer to documentation for information on setting up of shared storage for CP server database.

Enter absolute path of the database (Default:/etc/VRTScps/db):

## 10 Verify and confirm the CP server configuration information.

Following is the CP Server configuration information:

```
-----
(a) CP Server Name: mycps1.symantecexample.com
(b) CP Server Virtual IP: 10.209.83.85
(c) CP Server Port: 14250
(d) CP Server Security : 1
(e) CP Server Database Dir: /etc/VRTScps/db
-----
```

Press b if you want to change the configuration, <enter> to continue :

## 11 The configuration utility proceeds with the configuration process, and creates a vxcps.conf configuration file on each node.

The following output is for one node:

```
Successfully generated the /etc/vxcps.conf
configuration file.
Successfully created directory /etc/VRTScps/db.
Creating mount point /etc/VRTScps/db on
mycps1.symantecexample.com.
Copying configuration file /etc/vxcps.conf to
mycps1.symantecexample.com
```

Configuring CP Server Service Group (CPSSG) for this cluster

```
-----
```

**12 Confirm whether you use the same NIC name for the virtual IP on all the systems in the cluster.**

```
Is the name of NIC for virtual IP 10.209.83.85 same on all the systems?  
[y/n] : y
```

NOTE: Please ensure that the supplied network interface is a public NIC

**13 Enter a valid network interface for the virtual IP address for the CP server process.**

```
Enter a valid interface for virtual IP 10.209.83.85  
on all the systems : bge0
```

**14 Enter networkhosts information for the NIC resource.**

```
Symantec recommends configuring NetworkHosts attribute to ensure  
NIC resource to be online always.  
Do you want to add NetworkHosts attribute for the NIC resource bge0 on  
system mycps1? [y/n] : y  
Enter a valid IP address to configure NetworkHosts for NIC bge0 on  
system mycps1 : 10.209.83.86  
Do you want to add another Network Host ?[y/n] : n
```

**15 Enter the netmask for the virtual IP address.**

```
Enter the netmask for virtual IP 10.209.83.85 :  
255.255.252.0
```

**16 Enter the name of the disk group for the CP server database.**

```
Enter the name of diskgroup for cps database :  
cps_dg
```

**17 Enter the name of the volume that is created on the above disk group.**

```
Enter the name of volume created on diskgroup cps_dg :  
cps_volume
```

- 18** After the configuration process has completed, a success message appears. For example:

```
Successfully added the CPSSG service group to
VCS configuration. Bringing the CPSSG service
group online. Please wait...
```

```
The Veritas Coordination Point Server has been
configured on your system.
```

- 19** Run the `hagrp -state` command to ensure that the CPSSG service group has been added.

For example:

```
# hagrp -state CPSSG
```

| #Group | Attribute | System                     | Value   |
|--------|-----------|----------------------------|---------|
| CPSSG  | State     | mycps1.symantecexample.com | ONLINE  |
| CPSSG  | State     | mycps2.symantecexample.com | OFFLINE |

It also generates the configuration file for CP server (`/etc/vxcps.conf`).

The configuration utility adds the `vxcperv` process and other resources to the VCS configuration in the CP server service group (CPSSG).

For information about the CPSSG, refer to the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

In addition, the `main.cf` samples contain details about the `vxcperv` resource and its dependencies.

See [“Sample configuration files for CP server”](#) on page 693.

## Configuring the CP server manually

Perform the following steps to manually configure the CP server.

### To manually configure the CP server

- 1 Stop VCS on each node in the CP server cluster using the following command:

```
# hastop -local
```

- 2 Edit the `main.cf` file to add the CPSSG service group on any node. Use the CPSSG service group in the `main.cf` as an example:

See [“Sample configuration files for CP server”](#) on page 693.

Customize the resources under the CPSSG service group as per your configuration.

- 3 Verify the `main.cf` file using the following command:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

If successfully verified, copy this `main.cf` to all other cluster nodes.

- 4 Create the `/etc/vxcps.conf` file using the sample configuration file provided at `/etc/vxcps/vxcps.conf.sample`.

Based on whether you have configured the CP server cluster in secure mode or not, do the following:

- For a CP server cluster which is configured in secure mode, edit the `/etc/vxcps.conf` file to set `security=1`.
- For a CP server cluster which is not configured in secure mode, edit the `/etc/vxcps.conf` file to set `security=0`.

Symantec recommends enabling security for communication between CP server and the application clusters.

- 5 Start VCS on all the cluster nodes.

```
# hastart
```

- 6 Verify that the CP server service group (CPSSG) is online.

```
# hagrps -state CPSSG
```

Output similar to the following appears:

| # | Group Attribute | System                     | Value  |
|---|-----------------|----------------------------|--------|
|   | CPSSG State     | mycps1.symantecexample.com | ONLINE |

## Verifying the CP server configuration

Perform the following steps to verify the CP server configuration.

### To verify the CP server configuration

- 1 Verify that the following configuration files are updated with the information you provided during the CP server configuration process:
  - `/etc/vxcps.conf` (CP server configuration file)
  - `/etc/VRTSvcs/conf/config/main.cf` (VCS configuration file)
  - `/etc/VRTScps/db` (default location for CP server database)
- 2 Run the `cpsadm` command to check if the `vxcpsserv` process is listening on the configured Virtual IP.

```
# cpsadm -s cp_server -a ping_cps
```

where *cp\_server* is the virtual IP address or the virtual hostname of the CP server.





# Configuring SF Oracle RAC

This chapter includes the following topics:

- [About configuring SF Oracle RAC](#)
- [Configuring the SF Oracle RAC components using the script-based installer](#)
- [Configuring SF Oracle RAC using the Web-based installer](#)

## About configuring SF Oracle RAC

You need to configure SF Oracle RAC when:

- You have completed installation of SF Oracle RAC on your systems.
- You want to reconfigure an existing SF Oracle RAC cluster.

---

**Note:** Before you reconfigure a cluster, make sure that you stop any running applications that use VxFS/CFS. Then, unmount the VxFS/CFS mounts.

---

SF Oracle RAC configuration involves the following high-level tasks:

- Starting the product installer (if you quit the installer after installation or want to reconfigure the cluster)
- Configuring the SF Oracle RAC components—VCS, CVM, and CFS
- Configuring the SF Oracle RAC clusters for data integrity

During the configuration process, the installer performs the following tasks:

- Verifies the cluster information.
- Stops SF Oracle RAC processes.
- Starts SF Oracle RAC processes.

- Creates a new directory with a log file that contains any system commands executed, and their output, a response file that can be used with the `-responsefile` option of the installer, and a summary file that contains the output of the install scripts. The location of the files is indicated by the installer.

## Configuring the SF Oracle RAC components using the script-based installer

Make sure that you have performed the necessary pre-configuration tasks if you want to configure the cluster in secure mode.

Start the `installsfrac` or `installer` program if you quit the installer after installation.

At the end of the configuration, the VCS, CVM, and CFS components are configured to provide a cluster-aware environment.

---

**Note:** If you want to reconfigure SF Oracle RAC, before you start the installer you must stop all the resources that are under VCS control using the `hastop` command or the `hagrp -offline` command.

---

If you encounter issues during the configuration, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*, Chapter "Performance and troubleshooting" for information on resolving the issue.

### To configure the SF Oracle RAC components

- 1 Log in as the superuser on any of the nodes in the cluster.
- 2 Start the configuration program.

|                         |                                                             |
|-------------------------|-------------------------------------------------------------|
| SF Oracle RAC installer | Run the program:                                            |
|                         | <pre># cd /opt/VRTS/install</pre>                           |
|                         | <pre># ./installsfrac -configure <i>galaxy nebula</i></pre> |

|                          |                                                                                        |
|--------------------------|----------------------------------------------------------------------------------------|
| Common product installer | Run the program:                                                                       |
|                          | <pre># ./installer -configure <i>galaxy nebula</i></pre>                               |
|                          | Choose <b>Veritas Storage Foundation for Oracle RAC</b><br>to configure SF Oracle RAC. |

The installer displays the copyright message and specifies the directory where the logs are created.

- 3 Enter **1** to select the option **Configure SF Oracle RAC sub-components**.

```
1) Configure SF Oracle RAC sub-components
2) SF Oracle RAC Installation and Configuration Checks
3) Prepare to Install Oracle
4) Install Oracle Clusterware/Grid Infrastructure and Database
5) Post Oracle Installation Tasks
6) Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1)
```

- 4 If you had quit the installer in the process of an active configuration, the installer discovers that installer process and provides the option of resuming the configuration or starting a new configuration. Provide a suitable response.

```
The installer has discovered an existing installer process.
The process exited while performing configure of
SF Oracle RAC on galaxy.
Do you want to resume this process? [y,n,q,?] (y) n
```

- 5 Configure the Veritas Cluster Server component to set up the SF Oracle RAC cluster.

See [“Configuring the SF Oracle RAC cluster”](#) on page 148.

- 6 Add VCS users.  
See [“Adding VCS users”](#) on page 156.
- 7 Configure SMTP email notification.  
See [“Configuring SMTP email notification”](#) on page 156.
- 8 Configure SNMP trap notification.  
See [“Configuring SNMP trap notification”](#) on page 158.
- 9 Configure global clusters, if you chose to enable GCO during the installation.  
See [“Configuring global clusters”](#) on page 160.
- 10 Stop the SF Oracle RAC resources.  
See [“Stopping and starting SF Oracle RAC processes”](#) on page 161.

## Configuring the SF Oracle RAC cluster

You must configure the Veritas Cluster Server component to set up the SF Oracle RAC cluster. You can configure a basic cluster with only the required components, or an advanced cluster with any or all the optional features that meet your configuration requirements.

Refer to the *Veritas Cluster Server Installation Guide* for more information.

### Configuring the cluster name and ID

Enter the cluster information when the installer prompts you.

#### To configure the cluster

- 1 Review the configuration instructions that the installer presents.
- 2 Enter the unique cluster name and cluster ID.

```
Enter the unique cluster name: [q,?] rac_cluster101
```

```
Enter a unique Cluster ID number between 0-65535: [b,q,?] 7
```

### Configuring private heartbeat links

You now configure the private heartbeats that LLT uses. VCS provides the option to use LLT over Ethernet or over UDP (User Datagram Protocol). Symantec recommends that you configure heartbeat links that use LLT over Ethernet, unless hardware requirements force you to use LLT over UDP. If you want to configure LLT over UDP, make sure you meet the prerequisites.

The following procedure helps you configure LLT over Ethernet.

### To configure private heartbeat links

- 1 Choose one of the following options at the installer prompt based on whether you want to configure LLT over Ethernet or UDP.
  - Option 1: LLT over Ethernet (answer installer questions)  
 Enter the heartbeat link details at the installer prompt to configure LLT over Ethernet.  
 Skip to step 2.
  - Option 2: LLT over UDP (answer installer questions)  
 Make sure that each NIC you want to use as heartbeat link has an IP address configured. Enter the heartbeat link details at the installer prompt to configure LLT over UDP. If you had not already configured IP addresses to the NICs, the installer provides you an option to detect the IP address for a given NIC.  
 Skip to step 3.
  - Option 3: LLT over Ethernet (allow installer to detect)  
 Allow the installer to automatically detect the heartbeat link details to configure LLT over Ethernet. The installer tries to detect all connected links between all systems.  
 Skip to step 5.
- 2 If you chose option 1, enter the network interface card details for the private heartbeat links.

The installer discovers and lists the network interface cards. You can use either the standard interfaces or the aggregated interfaces (bonded NICs).

Answer the installer prompts. The following example shows different NICs based on architecture:

- For Solaris SPARC:  
 You must not enter the network interface card that is used for the public network (typically bge0.)

```
Enter the NIC for the first private heartbeat NIC on galaxy:
[b,q,?] bge1
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat NIC on galaxy:
[b,q,?] bge2
Would you like to configure a third private heartbeat link?
[y,n,q,b,?] (n)
Do you want to configure an additional low priority heartbeat
link? [y,n,q,b,?] (n)
```

■ For Solaris x64:

You must not enter the network interface card that is used for the public network (typically bge0.)

```
Enter the NIC for the first private heartbeat NIC on galaxy:
[b,q,?] e1000g1
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat NIC on galaxy:
[b,q,?] e1000g2
Would you like to configure a third private heartbeat link?
[y,n,q,b,?] (n)
Do you want to configure an additional low priority heartbeat
link? [y,n,q,b,?] (n)
```

**3** If you chose option 2, enter the NIC details for the private heartbeat links.  
This step uses examples such as *private\_NIC1* or *private\_NIC2* to refer to the available names of the NICs.

```
Enter the NIC for the first private heartbeat
NIC on galaxy: [b,q,?] private_NIC1
Do you want to use address 192.168.0.1 for the
first private heartbeat link on galaxy: [y,n,q,b,?] (y)
Enter the UDP port for the first private heartbeat
link on galaxy: [b,q,?] (50000) ?
Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat
NIC on galaxy: [b,q,?] private_NIC2
Do you want to use address 192.168.1.1 for the
second private heartbeat link on galaxy: [y,n,q,b,?] (y)
Enter the UDP port for the second private heartbeat
link on galaxy: [b,q,?] (50001) ?
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) y
Enter the NIC for the low priority heartbeat
link on galaxy: [b,q,?] (private_NIC0)
Do you want to use address 192.168.3.1 for
the low priority heartbeat link on galaxy: [y,n,q,b,?] (y)
Enter the UDP port for the low priority heartbeat
link on galaxy: [b,q,?] (50004)
```

- 4 Choose whether to use the same NIC details to configure private heartbeat links on other systems.

Are you using the same NICs for private heartbeat links on all systems? [y,n,q,b,?] (y)

If you want to use the NIC details that you entered for galaxy, make sure the same NICs are available on each system. Then, enter **y** at the prompt.

For LLT over UDP, if you want to use the same NICs on other systems, you still must enter unique IP addresses on each NIC for other systems.

If the NIC device names are different on some of the systems, enter **n**. Provide the NIC details for each system as the program prompts.

- 5 If you chose option 3, the installer detects NICs on each system and network links, and sets link priority.

If the installer fails to detect heartbeat links or fails to find any high-priority links, then choose option 1 or option 2 to manually configure the heartbeat links.

See step 2 for option 1, or step 3 for option 2.

- 6 Verify and confirm the information that the installer summarizes.

## Configuring the virtual IP of the cluster

You can configure the virtual IP of the cluster to use to connect to the Cluster Manager (Java Console) or to specify in the RemoteGroup resource.

See the *Veritas Cluster Server Administrator's Guide* for information on the Cluster Manager.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for information on the RemoteGroup agent.

### To configure the virtual IP of the cluster

- 1 Review the required information to configure the virtual IP of the cluster.
- 2 To configure virtual IP, enter **y** at the prompt.
- 3 Confirm whether you want to use the discovered public NIC on the first system.

Do one of the following:

- If the discovered NIC is the one to use, press **Enter**.
- If you want to use a different NIC, type the name of a NIC to use and press **Enter**.

```
Active NIC devices discovered on galaxy: bge0
Enter the NIC for Virtual IP of the Cluster to use on galaxy:
[b,q,?] (bge0)
```

**4** Confirm whether you want to use the same public NIC on all nodes.

Do one of the following:

- If all nodes use the same public NIC, enter *y*.
- If unique NICs are used, enter *n* and enter a NIC for each node.

```
Is bge0 to be the public NIC used by all systems
[y,n,q,b,?] (y)
```

**5** Enter the virtual IP address for the cluster.

```
Enter the Virtual IP address for the Cluster:
[b,q,?] 192.168.1.16
```

**6** Confirm the default netmask or enter another one:

```
Enter the netmask for IP 192.168.1.16: [b,q,?] (255.255.240.0)
```

**7** Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:
```

```
NIC: bge0
IP: 192.168.1.16
Netmask: 255.255.240.0
```

```
Is this information correct? [y,n,q] (y)
```

## Configuring the cluster in secure mode

If you want to configure the cluster in secure mode, make sure that you meet the prerequisites for secure cluster configuration.

The installer provides different configuration modes to configure a secure cluster. Make sure that you completed the pre-configuration tasks for the configuration mode that you want to choose.

See [“Preparing to configure the clusters in secure mode”](#) on page 119.



### To configure the cluster in secure mode

- 1 Choose whether to configure SF Oracle RAC to use Symantec Product Authentication Service.

```
Would you like to configure VCS to use Symantec Security
Services? [y,n,q] (n) y
```

- If you want to configure the cluster in secure mode, make sure you meet the prerequisites and enter **y**.
- If you do not want to configure the cluster in secure mode, enter **n**. You must add VCS users when the configuration program prompts. See [“Adding VCS users”](#) on page 156.

- 2 Select one of the options to enable security.

Before you choose any of the options, make sure that all the nodes in the cluster can successfully ping the root broker system.

```
Select the Security option you would like to perform [1-3,b,q,?] (1)
```

```
Security Menu
```

- ```

1) Configure security completely automatically
2) Provide AB credentials using BLOBs
3) Provide AB credentials without using BLOBs
b) Back to previous menu
```

Review the following configuration modes. Based on the configuration that you want to use, enter one of the following values:

Option 1.  
Automatic  
configuration

Based on the root broker you want to use, do one of the following:

■ To use an external root broker:

Enter the name of the root broker system when prompted.

Requires remote access to the root broker. Make sure that all the nodes in the cluster can successfully ping the root broker system.

Review the output as the installer verifies communication with the root broker system, checks vxatd process and version, and checks security domain.

■ To configure one of the nodes as root broker:

■ Press Enter at the following installer prompt:

```
If you already have an external
RB(Root Broker) installed and configured, enter
the RB name, or press Enter to skip: [b]
```

■ Choose the node that the installer must configure as root and authentication broker. The installer configures the other nodes as authentication brokers.

At the installer prompt, you can choose the first node in the cluster to configure as RAB, or you can enter n to configure another node as RAB. For example:

```
Do you want to configure <galaxy> as RAB,
and other nodes as AB? [y,n,q,b] (y) n
Enter the node name which you want to
configure as RAB: nebula
```

Option 2.  
Semiautomatic  
configuration

Enter the path of the encrypted file (BLOB file) for each node when prompted.

Option 3.  
Manual  
configuration

Enter the following Root Broker information as the installer prompts you:

```
Enter root broker name: [b]
east.symantecexample.com
Enter root broker FQDN: [b]
(symantecexample.com)
symantecexample.com
Enter the root broker domain name for the
Authentication Broker's identity: [b]
root@east.symantecexample.com
Enter root broker port: [b] 2821
Enter path to the locally accessible root hash [b]
(/var/tmp/installvcs-200910221810ROA/root_hash)
/var/tmp/installvcs-200910221810ROA/root_hash
```

Enter the following Authentication Broker information as the installer prompts you for each node:

```
Enter Authentication broker's identity on
galaxy [b]
(galaxy.symantecexample.com)
galaxy.symantecexample.com
Enter the password for the Authentication broker's
identity on galaxy:
Enter Authentication broker's identity on
nebula [b]
(nebula.symantecexample.com)
nebula.symantecexample.com
Enter the password for the Authentication broker's
identity on nebula:
```

- 3 After you provide the required information to configure the cluster in secure mode, the program prompts you to configure SMTP email notification.

Note that the installer does not prompt you to add VCS users if you configured the cluster in secure mode. However, you must add VCS users later.

See the *Veritas Cluster Server Administrator's Guide* for more information.

## Adding VCS users

If you have enabled Symantec Product Authentication Service, you do not need to add VCS users now. Otherwise, on systems operating under an English locale, you can add VCS users at this time.

### To add VCS users

- 1 Review the required information to add VCS users.
- 2 Reset the password for the Admin user, if necessary.

```
Do you want to set the username and/or password for the Admin user
(default username = 'admin', password='password')? [y,n,q] (n) y
Enter the user name: [b,q,?] (admin)
Enter the password:
Enter again:
```

- 3 To add a user, enter **y** at the prompt.

```
Do you want to add another user to the cluster? [y,n,q] (y)
```

- 4 Enter the user's name, password, and level of privileges.

```
Enter the user name: [b,q,?] smith
Enter New Password:*****

Enter Again:*****
Enter the privilege for user smith (A=Administrator, O=Operator,
G=Guest): [b,q,?] a
```

- 5 Enter **n** at the prompt if you have finished adding users.

```
Would you like to add another user? [y,n,q] (n)
```

- 6 Review the summary of the newly added users and confirm the information.

## Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP email services. You need to provide the SMTP server name and email addresses of people to be notified. Note that you can also configure the notification after installation.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information.

### To configure SMTP email notification

- 1 Review the required information to configure the SMTP email notification.
- 2 Specify whether you want to configure the SMTP notification.

Do you want to configure SMTP notification? [y,n,q,?] (n) **y**

If you do not want to configure the SMTP notification, you can skip to the next configuration option.

See [“Configuring SNMP trap notification”](#) on page 158.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: bge0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (bge0)
Is bge0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

- Enter the SMTP server's host name.

```
Enter the domain-based hostname of the SMTP server
(example: smtp.yourcompany.com): [b,q,?] smtp.example.com
```

- Enter the email address of each recipient.

```
Enter the full email address of the SMTP recipient
(example: user@yourcompany.com): [b,q,?] ozzie@example.com
```

- Enter the minimum security level of messages to be sent to each recipient.

```
Enter the minimum severity of events for which mail should be
sent to ozzie@example.com [I=Information, W=Warning,
E=Error, S=SevereError]: [b,q,?] w
```

- 4 Add more SMTP recipients, if necessary.

- If you want to add another SMTP recipient, enter **y** and provide the required information at the prompt.

Would you like to add another SMTP recipient? [y,n,q,b] (n) **y**

```
Enter the full email address of the SMTP recipient
```

```
(example: user@yourcompany.com): [b,q,?] harriet@example.com
```

```
Enter the minimum severity of events for which mail should be
sent to harriet@example.com [I=Information, W=Warning,
E=Error, S=SevereError]: [b,q,?] E
```

- If you do not want to add, answer **n**.

```
Would you like to add another SMTP recipient? [y,n,q,b] (n)
```

## 5 Verify and confirm the SMTP notification information.

```
NIC: bge0
```

```
SMTP Address: smtp.example.com
```

```
Recipient: ozzie@example.com receives email for Warning or
higher events
```

```
Recipient: harriet@example.com receives email for Error or
higher events
```

```
Is this information correct? [y,n,q] (y)
```

## Configuring SNMP trap notification

You can choose to configure VCS to send event notifications to SNMP management consoles. You need to provide the SNMP management console name to be notified and message severity levels.

Note that you can also configure the notification after installation.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information.

### To configure the SNMP trap notification

- 1 Review the required information to configure the SNMP notification feature of VCS.
- 2 Specify whether you want to configure the SNMP notification.

```
Do you want to configure SNMP notification? [y,n,q,?] (n) y
```

If you skip this option and if you had installed a valid HA/DR license, the installer presents you with an option to configure this cluster as global cluster. If you did not install an HA/DR license, the installer proceeds to configure SF Oracle RAC based on the configuration details you provided.

See [“Configuring global clusters”](#) on page 160.

### 3 Provide information to configure SNMP trap notification.

Provide the following information:

#### ■ Enter the NIC information.

```
Active NIC devices discovered on galaxy: bge0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (bge0)
Is bge0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

#### ■ Enter the SNMP trap daemon port.

```
Enter the SNMP trap daemon port: [b,q,?] (162)
```

#### ■ Enter the SNMP console system name.

```
Enter the SNMP console system name: [b,q,?] saturn
```

#### ■ Enter the minimum security level of messages to be sent to each console.

```
Enter the minimum severity of events for which SNMP traps
should be sent to saturn [I=Information, W=Warning, E=Error,
S=SevereError]: [b,q,?] E
```

### 4 Add more SNMP consoles, if necessary.

#### ■ If you want to add another SNMP console, enter *y* and provide the required information at the prompt.

```
Would you like to add another SNMP console? [y,n,q,b] (n) y
Enter the SNMP console system name: [b,q,?] jupiter
Enter the minimum severity of events for which SNMP traps
should be sent to jupiter [I=Information, W=Warning,
E=Error, S=SevereError]: [b,q,?] S
```

#### ■ If you do not want to add, answer *n*.

```
Would you like to add another SNMP console? [y,n,q,b] (n)
```

## 5 Verify and confirm the SNMP notification information.

```
NIC: bge0
```

```
SNMP Port: 162
```

```
Console: saturn receives SNMP traps for Error or  
higher events
```

```
Console: jupiter receives SNMP traps for SevereError or  
higher events
```

```
Is this information correct? [y,n,q] (y)
```

## Configuring global clusters

If you had installed a valid HA/DR license, the installer provides you an option to configure this cluster as global cluster. If not, the installer proceeds with other set of questions for CVM and CFS.

You can configure global clusters to link clusters at separate locations and enable wide-area failover and disaster recovery. The installer adds basic global cluster information to the VCS configuration file. You must perform additional configuration tasks to set up a global cluster.

See the *Veritas Cluster Server Administrator's Guide* for instructions to set up SF Oracle RAC global clusters.

---

**Note:** If you installed a HA/DR license to set up campus cluster, skip this installer option.

---

### To configure the global cluster option

- 1 Review the required information to configure the global cluster option.
- 2 Specify whether you want to configure the global cluster option.

```
Do you want to configure the Global Cluster Option? [y,n,q] (n) y
```

If you skip this option, the installer proceeds to configure VCS based on the configuration details you provided.



### 3 Provide information to configure this cluster as global cluster.

The installer prompts you for a NIC, a virtual IP address, and value for the netmask.

If you had entered virtual IP address details, the installer discovers the values you entered. You can use the same virtual IP address for global cluster configuration or enter different values.

### 4 Verify and confirm the configuration of the global cluster.

```
Global Cluster Option configuration verification:
```

```
NIC: bge0
IP: 192.168.1.16
Netmask: 255.255.240.0
```

```
Is this information correct? [y,n,q] (y)
```

On Solaris x64, an example for the NIC's port is bge0.

## Creation of SF Oracle RAC configuration files

The program consolidates all the information gathered in the preceding configuration tasks and creates configuration files.

If you chose to configure the cluster in secure mode, the installer also configures the Symantec Product Authentication Service. Depending on the mode you chose to set up Authentication Service, the installer creates security principal or executes the encrypted file to create security principal on each node in the cluster. The installer creates the VxSS service group, creates Authentication Server credentials on each node in the cluster, and Web credentials for VCS users, and sets up trust with the root broker. Then, the installer proceeds to start SF Oracle RAC in secure mode.

Review the output as the configuration program creates security principal, starts VCS, creates VCS configuration files, and copies the files to each node.

## Stopping and starting SF Oracle RAC processes

The installer stops and starts SF Oracle RAC processes and configures the SF Oracle RAC agents.

---

**Note:** To use aggregated interfaces that the installer has not discovered for private heartbeats, do not opt to start SF Oracle RAC.

---

#### To stop SF Oracle RAC processes

- 1 Enter **y** to stop SF Oracle RAC processes.

Do you want to stop SF Oracle RAC processes now? [y,n,q,?] (y)

- 2 Review the output as the installer stops and starts the SF Oracle RAC processes.

Note that SF Oracle RAC configuration program starts I/O fencing feature in disabled mode. SF Oracle RAC requires you to configure and enable I/O fencing feature.

## Configuring SF Oracle RAC using the Web-based installer

Before you begin to configure SF Oracle RAC using the Web-based installer, review the configuration requirements.

---

**Note:** If you want to reconfigure SF Oracle RAC, before you start the installer you must stop all the resources that are under VCS control using the `hasstop` command or the `hagrp -offline` command.

---

By default, the communication between the systems is selected as SSH. If SSH is used for communication between systems, the SSH commands execute without prompting for passwords or confirmations.

---

**Note:** To configure I/O fencing, you must either use the script-based installer or manually configure.

---

You can click **Quit** to quit the Web-installer at any time during the configuration process.

### To configure SF Oracle RAC on a cluster

- 1 Start the Web-based installer.  
See [“Starting the Veritas Web-based installer”](#) on page 97.
- 2 On the Select a task and a product page, select the task and the product as follows:

|                |                                           |
|----------------|-------------------------------------------|
| <b>Task</b>    | Configure a Product                       |
| <b>Product</b> | Veritas Storage Foundation for Oracle RAC |

Click **Next**.

- 3 On the Select Systems page, enter the system names where you want to configure SF Oracle RAC, and click **Validate**.

Example: **galaxy nebula**

The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

Click **Next** after the installer completes the system verification successfully.

- 4 On the Set Cluster Name/ID page, specify the following information for the cluster.

|                                  |                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Name                     | Enter a unique cluster name.                                                                                                                                                                                                                                                                                                                                                                  |
| Cluster ID                       | Enter a unique cluster ID.                                                                                                                                                                                                                                                                                                                                                                    |
| LLT Type                         | <p>Select an LLT type from the list. You can choose to configure LLT over UDP or over Ethernet.</p> <p>If you choose <b>Auto detect over Ethernet</b>, the installer auto-detects the LLT links over Ethernet. Verify the links and click <b>Yes</b> in the Confirmation dialog box. Skip to step 6. If you click No, you must manually enter the details to configure LLT over Ethernet.</p> |
| Number of Heartbeats             | Choose the number of heartbeat links you want to configure.                                                                                                                                                                                                                                                                                                                                   |
| Low Priority Heartbeat NIC       | Select the check box if you want to configure a low priority link. The installer configures one heartbeat link as low priority link.                                                                                                                                                                                                                                                          |
| Unique Heartbeat NICs per system | <p>For LLT over Ethernet, select the check box if you do not want to use the same NIC details to configure private heartbeat links on other systems.</p> <p>For LLT over UDP, this check box is selected by default.</p>                                                                                                                                                                      |

Click **Next**.

- 5 On the Set Cluster Heartbeat page, select the heartbeat link details for the LLT type you chose on the Set Cluster Name/ID page.

|                                |                                                                                                                                                                                                                                                                                                                               |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| For <b>LLT over Ethernet</b> : | <p>Do the following:</p> <ul style="list-style-type: none"><li>■ If you are using the same NICs on all the systems, select the NIC for each private heartbeat link.</li><li>■ If you had selected <b>Unique Heartbeat NICs per system</b> on the Set Cluster Name/ID page, provide the NIC details for each system.</li></ul> |
| For <b>LLT over UDP</b> :      | Select the NIC, Port, and IP address for each private heartbeat link. You must provide these details for each system.                                                                                                                                                                                                         |

Click **Next**.

- 6 In the Confirmation dialog box that appears, choose whether or not to configure the cluster in secure mode using Symantec Product Authentication Service (AT).

To configure the cluster in secure mode, click **Yes**.

If you want to perform this task later, click **No**. You can use the `installsf -security` command. Go to step 8.

- 7 On the Security Options page, choose an option to enable security and specify the required information.

|                                                       |                                                                                                                                                                                                                                          |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Do not configure security services</b>             | Choose this option if you do not want to enable security.<br><br>The installer takes you to the next page to configure optional features of SF Oracle RAC.                                                                               |
| <b>Configure security automatically</b>               | Choose this option to use an external root broker.<br><br>Enter the name of the root broker that is already configured for your enterprise environment, and click <b>Validate</b> . The installer configures the cluster in secure mode. |
| <b>Configure one node as RAB and the others as AB</b> | Select the system that you want to configure as RAB node.<br>The installer configures the cluster in secure mode.                                                                                                                        |

Click **Next**.

- 8 On the Optional Configuration page, decide the optional VCS features that you want to configure. Click the corresponding tab to specify the details for each option:

|                   |                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Virtual IP</b> | <ul style="list-style-type: none"> <li>■ Select the <b>Configure Virtual IP</b> check box.</li> <li>■ If each system uses a separate NIC, select the <b>Configure NICs for every system separately</b> check box.</li> <li>■ Select the interface on which you want to configure the virtual IP.</li> <li>■ Enter a virtual IP address and value for the netmask.</li> </ul> |
| <b>VCS Users</b>  | <ul style="list-style-type: none"> <li>■ Reset the password for the Admin user, if necessary.</li> <li>■ Click <b>Add</b> to add a new user.<br/>Specify the user name, password, and user privileges for this user.</li> </ul>                                                                                                                                              |

## SMTP

- Select the **Configure SMTP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- If all the systems use the same NIC, select the NIC for the VCS Notifier to be used on all systems. If not, select the NIC to be used by each system.
- In the **SMTP Server** box, enter the domain-based hostname of the SMTP server. Example: smtp.yourcompany.com
- In the **Recipient** box, enter the full email address of the SMTP recipient. Example: user@yourcompany.com.
- In the **Event** list box, select the minimum security level of messages to be sent to each recipient.
- Click **Add** to add more SMTP recipients, if necessary.

## SNMP

- Select the **Configure SNMP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- If all the systems use the same NIC, select the NIC for the VCS Notifier to be used on all systems. If not, select the NIC to be used by each system.
- In the **SNMP Port** box, enter the SNMP trap daemon port: (162).
- In the **Console System Name** box, enter the SNMP console system name.
- In the **Event** list box, select the minimum security level of messages to be sent to each console.
- Click **Add** to add more SNMP consoles, if necessary.

## GCO

If you installed a valid HA/DR license, you can now enter the wide-area heartbeat link details for the global cluster that you would set up later.

See the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide* for instructions to set up SF Oracle RAC global clusters.

- Select the **Configure GCO** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- Select a NIC.
- Enter a virtual IP address and value for the netmask.

Click **Next**.

- 9 On the Stop Processes page, click **Next** after the installer stops all the processes successfully.

- 10 On the Start Processes page, click **Next** after the installer performs the configuration based on the details you provided and starts all the processes successfully.
- 11 Click **Next** to complete the process of configuring SF Oracle RAC.  
On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.
- 12 Select the checkbox to specify whether you want to send your installation information to Symantec.  
Click **Finish**. The installer prompts you for another task.





# Configuring SF Oracle RAC clusters for data integrity

This chapter includes the following topics:

- [Setting up disk-based I/O fencing using installsfrac program](#)
- [Setting up disk-based I/O fencing manually](#)
- [Setting up server-based I/O fencing using installsfrac program](#)
- [Setting up server-based I/O fencing manually](#)

## Setting up disk-based I/O fencing using installsfrac program

You can configure I/O fencing using the `-fencing` option of the `installsfrac` program.

### Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

#### To initialize disks as VxVM disks

- 1 List the new external disks or the LUNs as recognized by the operating system.  
On each node, enter:  
  

```
# devfsadm
```
- 2 To initialize the disks as VxVM disks, use one of the following methods:
  - Use the interactive `vxdiskadm` utility to initialize the disks as VxVM disks.

For more information see the *Veritas Volume Manager Administrator's Guide*.

- Use the `vxdisksetup` command to initialize a disk as a VxVM disk.

```
vxdisksetup -i device_name
```

The example specifies the CDS format:

```
# vxdisksetup -i c2t13d0
```

Repeat this command for each disk you intend to use as a coordinator disk.

## Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See [“Initializing disks as VxVM disks”](#) on page 169.

Review the following procedure to identify disks to use as coordinator disks.

### To identify the coordinator disks

- 1 List the disks on each node.

For example, execute the following commands to list the disks:

```
# vxdisk -o alldgs list
```

- 2 Pick three SCSI-3 PR compliant shared disks as coordinator disks.

See [“Checking shared disks for I/O fencing”](#) on page 170.

## Checking shared disks for I/O fencing

Make sure that the shared storage you set up while preparing to configure SF Oracle RAC meets the I/O fencing requirements. You can test the shared disks using the `vxfcntlshdw` utility. The two nodes must have `ssh` (default) or `rsh` communication. To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxfenadm` command with the `-i` option. This command option verifies that the same serial number for the LUN is returned on all paths to the LUN.

Make sure to test the disks that serve as coordinator disks.

The `vxfsentsthdw` utility has additional options suitable for testing many disks. Review the options for testing the disk groups (`-g`) and the disks that are listed in a file (`-f`). You can also test disks without destroying data using the `-r` option.

See the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

Checking that disks support SCSI-3 involves the following tasks:

- Verifying the Array Support Library (ASL)  
See [“Verifying Array Support Library \(ASL\)”](#) on page 171.
- Verifying that nodes have access to the same disk  
See [“Verifying that the nodes have access to the same disk”](#) on page 172.
- Testing the shared disks for SCSI-3  
See [“Testing the disks using `vxfsentsthdw` utility”](#) on page 173.

## Verifying Array Support Library (ASL)

Make sure that the Array Support Library (ASL) for the array that you add is installed.

To verify Array Support Library (ASL)

- 1 If the Array Support Library (ASL) for the array that you add is not installed, obtain and install it on each node before proceeding.

The ASL for the supported storage device that you add is available from the disk array vendor or Symantec technical support.

- 2 Verify that the ASL for the disk array is installed on each of the nodes. Run the following command on each node and examine the output to verify the installation of ASL.

The following output is a sample:

```
# vxddladm listsupport all
```

| LIBNAME            | VID      | PID            |
|--------------------|----------|----------------|
| libvx3par.so       | 3PARdata | VV             |
| libvxCLARiiON.so   | DGC      | All            |
| libvxFUITSYe6k.so  | FUJITSU  | E6000          |
| libvxFUITSYe8k.so  | FUJITSU  | All            |
| libvxap.so         | SUN      | All            |
| libvxatf.so        | VERITAS  | ATFNODES       |
| libvxcompellent.so | COMPELNT | Compellent Vol |
| libvxcopan.so      | COPANSYS | 8814, 8818     |

- 3 Scan all disk drives and their attributes, update the VxVM device list, and reconfigure DMP with the new devices. Type:

```
# vxdisk scandisks
```

See the Veritas Volume Manager documentation for details on how to add and configure disks.

Verifying that the nodes have access to the same disk

Before you test the disks that you plan to use as shared data storage or as coordinator disks using the vxfcntl utility, you must verify that the systems see the same disk.

To verify that the nodes have access to the same disk

- 1 Verify the connection of the shared storage for data to two of the nodes on which you installed SF Oracle RAC.
- 2 Ensure that both nodes are connected to the same disk during the testing. Use the vxfsadm command to verify the disk serial number.

```
vxfenadm -i diskpath
```

Refer to the `vxfenadm` (1M) manual page.

For example, an EMC disk is accessible by the `/dev/rdisk/c1t1d0s2` path on node A and the `/dev/rdisk/c2t1d0s2` path on node B.

From node A, enter:

```
vxfenadm -i /dev/rdisk/c1t1d0s2
```

```
Vendor id : EMC  
Product id : SYMMETRIX  
Revision : 5567  
Serial Number : 42031000a
```

The same serial number information should appear when you enter the equivalent command on node B using the `/dev/rdisk/c2t1d0s2` path.

On a disk from another manufacturer, Hitachi Data Systems, the output is different and may resemble:

```
# vxfenadm -i /dev/rdisk/c3t1d2s2  
  
Vendor id      : HITACHI  
Product id     : OPEN-3      -SUN  
Revision       : 0117  
Serial Number  : 0401EB6F0002
```

## Testing the disks using vxfcntlsthaw utility

This procedure uses the `/dev/rdisk/c1t1d0s2` disk in the steps.

If the utility does not show a message that states a disk is ready, the verification has failed. Failure of verification can be the result of an improperly configured disk array. The failure can also be due to a bad disk.

If the failure is due to a bad disk, remove and replace it. The `vxfcntlsthaw` utility indicates a disk can be used for I/O fencing with a message resembling:

```
The disk /dev/rdisk/c1t1d0s2 is ready to be configured for I/O Fencing on  
node galaxy
```

For more information on how to replace coordinator disks, refer to the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

### To test the disks using vxfcntlshdw utility

- 1 Make sure system-to-system communication functions properly.  
See “[Setting up inter-system communication](#)” on page 702.
- 2 From one node, start the utility.  
Run the utility with the -n option if you use rsh for communication.  
  

```
# vxfcntlshdw [-n]
```
- 3 The script warns that the tests overwrite data on the disks. After you review the overview and the warning, confirm to continue the process and enter the node names.

---

**Warning:** The tests overwrite and destroy data on the disks unless you use the -r option.

---

```
***** WARNING!!!!!!!!!! *****
THIS UTILITY WILL DESTROY THE DATA ON THE DISK!!

Do you still want to continue : [y/n] (default: n) y
Enter the first node of the cluster: galaxy
Enter the second node of the cluster: nebula
```

- 4 Enter the names of the disks that you want to check. Each node may know the same disk by a different name.

```
Enter the disk name to be checked for SCSI-3 PGR on node
IP_adrs_of_galaxy in the format:
for dmp: /dev/vx/rdmp/cxtxdxsx
for raw: /dev/rdisk/cxtxdxsx
Make sure it's the same disk as seen by nodes
IP_adrs_ofgalaxy and IP_adrs_of_nebula
/dev/rdsk/c2t13d0s2
```

```
Enter the disk name to be checked for SCSI-3 PGR on node
IP_adrs_of_nebula in the format:
for dmp: /dev/vx/rdmp/cxtxdxsx
for raw: /dev/rdisk/cxtxdxsx
Make sure it's the same disk as seen by nodes
IP_adrs_ofgalaxy and IP_adrs_of_nebula
/dev/rdsk/c2t13d0s2
```

If the serial numbers of the disks are not identical, then the test terminates.

- 5 Review the output as the utility performs the checks and report its activities.
- 6 If a disk is ready for I/O fencing on each node, the utility reports success for each node. For example, the utility displays the following message for the node galaxy.

```
The disk is now ready to be configured for I/O Fencing on node
galaxy
```

```
ALL tests on the disk /dev/rdsk/c1t1d0s2 have PASSED
The disk is now ready to be configured for I/O Fencing on node
galaxy
```

- 7 Run the vxfcntlsthdw utility for each disk you intend to verify.

## Configuring disk-based I/O fencing using installsfrac program

---

**Note:** The installer stops and starts SF Oracle RAC to complete I/O fencing configuration. Make sure to unfreeze any frozen VCS service groups in the cluster for the installer to successfully stop SF Oracle RAC.

---

### To set up disk-based I/O fencing using the installsfrac program

- 1 Start the installsfrac program with `-fencing` option.

```
# /opt/VRTS/install/installsfrac -fencing
```

The installsfrac program starts with a copyright message and verifies the cluster information.

Note the location of log files which you can access in the event of any problem with the configuration process.

- 2 Confirm that you want to proceed with the I/O fencing configuration at the prompt.

The program checks that the local node running the script can communicate with remote nodes and checks whether SF Oracle RAC 5.1 Service Pack 1 is configured properly.

- 3 Review the I/O fencing configuration options that the program presents. Type **2** to configure disk-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-3,b,q] 2
```

- 4 Review the output as the configuration program checks whether VxVM is already started and is running.

- If the check fails, configure and enable VxVM before you repeat this procedure.
- If the check passes, then the program prompts you for the coordinator disk group information.

- 5 Choose whether to use an existing disk group or create a new disk group to configure as the coordinator disk group.

The program lists the available disk group names and provides an option to create a new disk group. Perform one of the following:

- To use an existing disk group, enter the number corresponding to the disk group at the prompt.  
 The program verifies whether the disk group you chose has an odd number of disks and that the disk group has a minimum of three disks.
- To create a new disk group, perform the following steps:
  - Enter the number corresponding to the **Create a new disk group** option.  
 The program lists the available disks that are in the CDS disk format in the cluster and asks you to choose an odd number of disks with at least three disks to be used as coordinator disks.



Symantec recommends that you use three disks as coordination points for disk-based I/O fencing.

- Enter the numbers corresponding to the disks that you want to use as coordinator disks.
- Enter the disk group name.

**6** Verify that the coordinator disks you chose meet the I/O fencing requirements.

You must verify that the disks are SCSI-3 PR compatible using the `vxfsntsthdw` utility and then return to this configuration program.

See [“Checking shared disks for I/O fencing”](#) on page 170.

**7** After you confirm the requirements, the program creates the coordinator disk group with the information you provided.

**8** Enter the I/O fencing disk policy that you chose to use. For example:

```
Enter fencing mechanism name (raw/dmp): [b,q,?] raw
```

The program also does the following:

- Populates the `/etc/vxfendg` file with this disk group information
- Populates the `/etc/vxfsenmode` file on each cluster node with the I/O fencing mode information and with the SCSI-3 disk policy information

**9** Verify and confirm the I/O fencing configuration information that the installer summarizes.

**10** Review the output as the configuration program does the following:

- Stops VCS and I/O fencing on each node.
- Configures disk-based I/O fencing and starts the I/O fencing process.
- Updates the VCS configuration file `main.cf` if necessary.
- Copies the `/etc/vxfsenmode` file to a date and time suffixed file `/etc/vxfsenmode-date-time`. This backup file is useful if any future fencing configuration fails.
- Starts VCS on each node to make sure that the SF Oracle RAC is cleanly configured to use the I/O fencing feature.

**11** Review the output as the configuration program displays the location of the log files, the summary files, and the response files.

**12** Configure the Coordination Point agent to monitor the coordinator disks.

See [“Configuring Coordination Point agent to monitor coordination points”](#) on page 203.

# Setting up disk-based I/O fencing manually

Table 9-1 lists the tasks that are involved in setting up I/O fencing.

Table 9-1 Tasks to set up I/O fencing manually

| Task                                                                | Reference                                                                                              |
|---------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| Initializing disks as VxVM disks                                    | See <a href="#">“Initializing disks as VxVM disks”</a> on page 169.                                    |
| Identifying disks to use as coordinator disks                       | See <a href="#">“Identifying disks to use as coordinator disks”</a> on page 170.                       |
| Checking shared disks for I/O fencing                               | See <a href="#">“Checking shared disks for I/O fencing”</a> on page 170.                               |
| Setting up coordinator disk groups                                  | See <a href="#">“Setting up coordinator disk groups”</a> on page 178.                                  |
| Creating I/O fencing configuration files                            | See <a href="#">“Creating I/O fencing configuration files”</a> on page 179.                            |
| Modifying SF Oracle RAC configuration to use I/O fencing            | See <a href="#">“Modifying VCS configuration to use I/O fencing”</a> on page 180.                      |
| Configuring Coordination Point agent to monitor coordination points | See <a href="#">“Configuring Coordination Point agent to monitor coordination points”</a> on page 203. |
| Starting SF Oracle RAC on all nodes                                 | See <a href="#">“Starting SF Oracle RAC on all nodes”</a> on page 181.                                 |
| Verifying I/O fencing configuration                                 | See <a href="#">“Verifying I/O fencing configuration”</a> on page 183.                                 |

## Setting up coordinator disk groups

From one node, create a disk group named vxencoorddg. This group must contain three disks or LUNs. You must also set the coordinator attribute for the coordinator disk group. VxVM uses this attribute to prevent the reassignment of coordinator disks to other disk groups.

Note that if you create a coordinator disk group as a regular disk group, you can turn on the coordinator attribute in Volume Manager.

Refer to the *Veritas Volume Manager Administrator’s Guide* for details on how to create disk groups.

The following example procedure assumes that the disks have the device names `c1t1d0s2`, `c2t1d0s2`, and `c3t1d0s2`.

### To create the `vxfencoorddg` disk group

- 1 On any node, create the disk group by specifying the device names:

```
# vxdg init vxfencoorddg c1t1d0s2 c2t1d0s2 c3t1d0s2
```

- 2 Set the coordinator attribute value as "on" for the coordinator disk group.

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

- 3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

- 4 Import the disk group with the `-t` option to avoid automatically importing it when the nodes restart:

```
# vxdg -t import vxfencoorddg
```

- 5 Deport the disk group. Deporting the disk group prevents the coordinator disks from serving other purposes:

```
# vxdg deport vxfencoorddg
```

## Creating I/O fencing configuration files

After you set up the coordinator disk group, you must do the following to configure I/O fencing:

- Create the I/O fencing configuration file `/etc/vxfendg`
- Update the I/O fencing configuration file `/etc/vxfenmode`

### To update the I/O fencing files and start I/O fencing

- 1 On each nodes, type:

```
# echo "vxfencoorddg" > /etc/vxfendg
```

Do not use spaces between the quotes in the "vxfencoorddg" text.

This command creates the `/etc/vxfendg` file, which includes the name of the coordinator disk group.

- 2 On all cluster nodes depending on the SCSI-3 mechanism, type one of the following selections:

- For DMP configuration:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode
```

- For raw device configuration:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_raw /etc/vxfenmode
```

- 3 To check the updated /etc/vxfenmode configuration, enter the following command on one of the nodes. For example:

```
# more /etc/vxfenmode
```

- 4 Edit the following file on each node in the cluster to change the values of the VXFEN\_START and the VXFEN\_STOP environment variables to 1:

```
/etc/default/vxfen
```

## Modifying VCS configuration to use I/O fencing

After you add coordination points and configure I/O fencing, add the UseFence = SCSI3 cluster attribute to the VCS configuration file /etc/VRTSvcs/conf/config/main.cf. If you reset this attribute to UseFence = None, VCS does not make use of I/O fencing abilities while failing over service groups. However, I/O fencing needs to be disabled separately.

### To modify VCS configuration to enable I/O fencing

- 1 Save the existing configuration:

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

- 2 Stop VCS on all nodes:

```
# hastop -all
```

- 3 If the I/O fencing driver vxfen is already running, stop the I/O fencing driver. Depending on the Solaris version on the cluster nodes, run the following command:

- Solaris 9:

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

- Solaris 10:

```
# svcadm disable -t vxfen
```

- 4 Make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

- 5 On one node, use vi or another text editor to edit the main.cf file. To modify the list of cluster attributes, add the UseFence attribute and assign its value as SCSI3.

```
cluster rac_cluster101(
  UserNames = { admin = "cDRpdxPmHpzS." }
  Administrators = { admin }
  HacliUserLevel = COMMANDROOT
  CounterInterval = 5
  UseFence = SCSI3
)
```

Regardless of whether the fencing configuration is disk-based or server-based, the value of the cluster-level attribute UseFence is set to SCSI3.

- 6 Save and close the file.
- 7 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

- 8 Using rcp or another utility, copy the VCS configuration file from a node (for example, galaxy) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp galaxy:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

## Starting SF Oracle RAC on all nodes

You must start SF Oracle RAC on each node to bring up the cluster configuration with I/O fencing enabled.

### Starting I/O fencing, VCS, CVM, and CFS

You must start I/O fencing, VCS, CVM, and CFS on all nodes in the cluster.

#### To start VCS, CVM, and CFS on a node

- 1 Start the I/O fencing driver. Run the following command on each node:

Depending on the Solaris version on the cluster nodes, run the following command:

■ Solaris 9:

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

■ Solaris 10:

```
# svcadm enable vxfen
```

The vxfen startup script also invokes the `vxfenconfig` command, which configures the vxfen driver to start and use the coordinator disks that are listed in `/etc/vxfentab`.

2 With the configuration file in place on each system, start VCS, CVM, and CFS:

```
# hastart
```

## Verifying GAB port membership

After setting up I/O fencing and starting VCS, CVM, and CFS on each node, verify GAB port membership.

See [“Verifying GAB”](#) on page 220.

### To verify GAB port membership

- ◆ Run the `gabconfig -a` command.

For example:

```
galaxy# gabconfig -a
GAB Port Memberships
=====
Port a gen ada401 membership 01
Port b gen ada40d membership 01
Port d gen ada409 membership 01
Port f gen ada41c membership 01
Port h gen ada40f membership 01
Port o gen ada406 membership 01

Port u gen ada41a membership 01

Port v gen ada416 membership 01
Port w gen ada418 membership 01
```

## Verifying the CVM group is online

Make sure the cvm group is in the online state.

### To verify CVM group

- 1 On any node, run the following command to verify that the cvm group is online:

```
# hagrps -state cvm
```

- 2 On any node, run the following command to make sure CVM has started properly:

```
# vxclustadm nidmap
```

## Verifying I/O fencing configuration

Verify from the vxfenadm output that the SCSI-3 disk policy reflects the configuration in the /etc/vxfenmode file.

### To verify I/O fencing configuration

- 1 On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the SCSI3 disk policy is dmp:

```
I/O Fencing Cluster Information:
=====
```

```
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
```

```
  * 0 (galaxy)
  1 (nebula)
```

```
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 2 Verify that the disk-based I/O fencing is using the specified disks.

```
# vxfenconfig -l
```

## Setting up server-based I/O fencing using installsfrac program

If SF Oracle RAC cluster is configured to run in secure mode, then verify that the configuration is correct before you configure CP server-based I/O fencing.

See [“Verifying the security configuration on the SF Oracle RAC cluster to use CP server coordination point”](#) on page 184.

See [“Configuring server-based I/O fencing using the installsfrac program”](#) on page 186.

## Verifying the security configuration on the SF Oracle RAC cluster to use CP server coordination point

After configuring security using the `installsfrac -security` command, follow the procedure below on each SF Oracle RAC cluster node to confirm that security is correctly configured.



**To verify the security configuration on SF Oracle RAC cluster to use CP server coordination point**

- 1** Run the following command:

```
# /opt/VRTScps/bin/cpsat listpd -t local
```

The following is an example of the command output:

```
Domain(s) Found 1

*****

Domain Name HA_SERVICES@galaxy.symantec.com

Expiry Interval 0

*****
```

- 2** There should be a domain name entry with the following format in the command output:

```
HA_SERVICES@hostname.domainname
```

or

```
HA_SERVICES@hostname
```

**3** There should not be duplicate entries for HA\_SERVICES domain.

The following is an example of an incorrect configuration:

```
showdomains

Domain(s) Found :          3

*****

Domain Name:      HA_SERVICES@galaxy.symantec.com

Domain Type:      vx

*****

Domain Name:      broker@galaxy.symantec.com

Domain Type:      vx

*****

Domain Name:      HA_SERVICES@galaxy

Domain Type:      vx

*****
```

Proceed to reconfigure security in case duplicate entries appear as shown in the above example.

## Configuring server-based I/O fencing using the installsfrac program

You can configure server-based I/O fencing for the SF Oracle RAC cluster using the installsfrac program.

With server-based fencing, you can have the coordination points in your configuration as follows:

- Combination of CP servers and SCSI-3 compliant coordinator disks
  - CP servers only
- Symantec also supports server-based fencing with a single highly available CP server that acts as a single coordination point.

See [“About planning to configure I/O fencing”](#) on page 67.

See [“Recommended CP server configurations”](#) on page 71.

This section covers the following example procedures:

|                                         |                                                                                                                                              |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| Mix of CP servers and coordinator disks | See <a href="#">“To configure server-based fencing for the SF Oracle RAC cluster (one CP server and two coordinator disks)”</a> on page 187. |
| Single CP server                        | See <a href="#">“To configure server-based fencing for the SF Oracle RAC cluster (single CP server)”</a> on page 192.                        |

### **To configure server-based fencing for the SF Oracle RAC cluster (one CP server and two coordinator disks)**

- 1 Depending on the server-based configuration model in your setup, make sure of the following:
  - CP servers are configured and are reachable from the SF Oracle RAC cluster. The SF Oracle RAC cluster is also referred to as the application cluster or the client cluster.  
See [“Setting up the CP server”](#) on page 128.
  - The coordination disks are verified for SCSI3-PR compliance.  
See [“Checking shared disks for I/O fencing”](#) on page 170.
- 2 Start the installsfrac program with `-fencing` option.

```
# /opt/VRTS/install/installsfrac -fencing
```

The installsfrac program starts with a copyright message and verifies the cluster information.

Note the location of log files which you can access in the event of any problem with the configuration process.

- 3 Confirm that you want to proceed with the I/O fencing configuration at the prompt.  
  
The program checks that the local node running the script can communicate with remote nodes and checks whether SF Oracle RAC 5.1 Service Pack 1 is configured properly.
- 4 Review the I/O fencing configuration options that the program presents.  
Type **1** to configure server-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-3,b,q] 1
```

- 5 Make sure that the storage supports SCSI3-PR, and answer y at the following prompt.

```
Does your storage environment support SCSI3 PR? [y,n,q] (y)
```

- 6 Provide the following details about the coordination points at the installer prompt:

- Enter the total number of coordination points including both servers and disks. This number should be at least 3.

```
Enter the total number of co-ordination points including both  
CP servers and disks: [b] (3)
```

- Enter the total number of coordinator disks among the coordination points.

```
Enter the total number of disks among these:  
[b] (0) 2
```

- 7 Provide the following CP server details at the installer prompt:

- Enter the virtual IP addresses or host names of the virtual IP address for each of the CP servers. The installer assumes these values to be identical as viewed from all the application cluster nodes.

```
Enter the Virtual IP address/fully qualified host name  
for the Co-ordination Point Server #1:  
[b] 10.209.80.197
```

- Enter the port that the CP server would be listening on.

```
Enter the port in the range [49152, 65535] which the  
Co-ordination Point Server 10.209.80.197  
would be listening on or simply accept the default port suggested:  
[b] (14250)
```

- 8 Provide the following coordinator disks-related details at the installer prompt:

- Enter the I/O fencing disk policy for the coordinator disks.

```
Enter fencing mechanism for the disk(s) (raw/dmp):  
[b,q,?] raw
```

- Choose the coordinator disks from the list of available disks that the installer displays. Ensure that the disk you choose is available from all the SF Oracle RAC (application cluster) nodes.

The number of times that the installer asks you to choose the disks depends on the information that you provided in step 6. For example, if you had chosen to configure two coordinator disks, the installer asks you to choose the first disk and then the second disk:

```
Select disk number 1 for co-ordination point
```

- 1) c1t1d0s2
- 2) c2t1d0s2
- 3) c3t1d0s2

```
Please enter a valid disk which is available from all the
cluster nodes for co-ordination point [1-3,q] 1
```

- If you have not already checked the disks for SCSI-3 PR compliance in step 1, check the disks now.  
 The installer displays a message that recommends you to verify the disks in another window and then return to this configuration procedure.  
 Press Enter to continue, and confirm your disk selection at the installer prompt.
- Enter a disk group name for the coordinator disks or accept the default.

```
Enter the disk group name for coordinating disk(s):
[b] (vxfencoorddg)
```

## 9 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 3
CP Server (Port):
  1. 10.209.80.197 (14250)
SCSI-3 disks:
  1. c1t1d0s2
  2. c2t1d0s2
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk mechanism used for customized fencing: raw
```

The installer initializes the disks and the disk group and depots the disk group on the SF Oracle RAC (application cluster) node.

- 10 If the CP server is configured for security, the installer sets up secure communication between the CP server and the SF Oracle RAC (application cluster):

- Make sure that the security configuration in the application cluster and the CP server is the same. If CP server is configured for security, ensure that the application cluster also runs in secure mode.
- If the CP server is configured for security, perform the following steps:
  - Review the output as the installer verifies if the SF Oracle RAC (application cluster) nodes have already established trust with an AT root broker.
  - If the SF Oracle RAC (application cluster) nodes and the CP server use different AT root brokers, enter y at the installer prompt and provide the following information:
    - Hostname for the authentication broker for any one of the CP servers
    - Port number where the authentication broker for the CP server is listening for establishing trust
    - Hostname for the authentication broker for any one of the SF Oracle RAC (application cluster) nodes
    - Port number where the authentication broker for the SF Oracle RAC (application cluster) is listening for establishing trust

After the installer establishes trust between the authentication brokers of the CP servers and the application cluster nodes, press Enter to continue.

## 11 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: rac_cluster101
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

- 12** Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the `/etc/vxfenmode` file with the appropriate details in each of the application cluster nodes.

```
Updating client cluster information on CP Server 10.210.80.199

Adding the client cluster to the CP Server 10.210.80.199 ..... Done

Registering client node galaxy with CP Server 10.210.80.199..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster rac_cluster101 to the CPClient user on CP Server 10.210.80.199 ... Done

Registering client node nebula with CP Server 10.210.80.199 ..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster rac_cluster101 to the CPClient user on CP Server 10.210.80.199 ... Done

Updating /etc/vxfenmode file on galaxy ..... Done
Updating /etc/vxfenmode file on nebula ..... Done
```

See [“About I/O fencing configuration files”](#) on page 672.

- 13** Configure the CP agent on the SF Oracle RAC (application cluster).

```
Do you want to configure CP Agent on the client cluster? [y,n,q]
(y)

Enter a non-existing name for the service group for CP Agent:
[b] (vxfen)

Adding CP Agent via galaxy ..... Done
```

- 14** Review the output as the installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.
- 15** Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

**To configure server-based fencing for the SF Oracle RAC cluster (single CP server)**

- 1** Make sure that the CP server is configured and is reachable from the SF Oracle RAC cluster. The SF Oracle RAC cluster is also referred to as the application cluster or the client cluster.

See “[Setting up the CP server](#)” on page 128.

- 2** Start the installsfrac program with `-fencing` option.

```
# /opt/VRTS/install/installsfrac -fencing
```

The installsfrac program starts with a copyright message and verifies the cluster information.

Note the location of log files which you can access in the event of any problem with the configuration process.

- 3** Confirm that you want to proceed with the I/O fencing configuration at the prompt.

The program checks that the local node running the script can communicate with remote nodes and checks whether SF Oracle RAC 5.1 Service Pack 1 is configured properly.

- 4** Review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

```
Select the fencing mechanism to be configured in this  
Application Cluster [1-3,b,q] 1
```

- 5** Make sure that the storage supports SCSI3-PR, and answer **y** at the following prompt.

```
Does your storage environment support SCSI3 PR? [y,n,q] (y)
```

- 6** Enter the total number of coordination points as **1**.

```
Enter the total number of co-ordination points including both  
CP servers and disks: [b] (3) 1
```

Read the installer warning carefully before you proceed with the configuration.

- 7** Provide the following CP server details at the installer prompt:

- Enter the virtual IP address or the host name of the virtual IP address for the CP server. The installer assumes these values to be identical as viewed from all the application cluster nodes.



```
Enter the Virtual IP address/fully qualified host name
for the Co-ordination Point Server #1:
[b] 10.209.80.197
```

- Enter the port that the CP server would be listening on.

```
Enter the port in the range [49152, 65535] which the
Co-ordination Point Server 10.209.80.197
would be listening on or simply accept the default port suggested:
[b] (14250)
```

## 8 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 1
CP Server (Port):
    1. 10.209.80.197 (14250)
```

## 9 If the CP server is configured for security, the installer sets up secure communication between the CP server and the SF Oracle RAC (application cluster):

- Make sure that the security configuration in the application cluster and the CP server is the same. If CP server is configured for security, ensure that the application cluster also runs in secure mode.
- If the CP server is configured for security, perform the following steps:
  - Review the output as the installer verifies if the SF Oracle RAC (application cluster) nodes have already established trust with an AT root broker.
  - If the SF Oracle RAC (application cluster) nodes and the CP server use different AT root brokers, enter y at the installer prompt and provide the following information:
    - Hostname for the authentication broker for any one of the CP servers
    - Port number where the authentication broker for the CP server is listening for establishing trust
    - Hostname for the authentication broker for any one of the SF Oracle RAC (application cluster) nodes
    - Port number where the authentication broker for the SF Oracle RAC (application cluster) is listening for establishing trust

After the installer establishes trust between the authentication brokers of the CP servers and the application cluster nodes, press Enter to continue.

## 10 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: rac_cluster101
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

## 11 Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the `/etc/vxfenmode` file with the appropriate details in each of the application cluster nodes.

The installer also populates the `/etc/vxfenmode` file with the entry `single_cp=1` for such single CP server fencing configuration.

```
Updating client cluster information on CP Server 10.210.80.199
```

```
Adding the client cluster to the CP Server 10.210.80.199 ..... Done
```

```
Registering client node galaxy with CP Server 10.210.80.199 ..... Done
```

```
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
```

```
Adding cluster rac_cluster101 to the CPClient user on CP Server 10.210.80.199 ... Done
```

```
Registering client node nebula with CP Server 10.210.80.199 ..... Done
```

```
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
```

```
Adding cluster rac_cluster101 to the CPClient user on CP Server 10.210.80.199 ... Done
```

```
Updating /etc/vxfenmode file on galaxy ..... Done
```

```
Updating /etc/vxfenmode file on nebula ..... Done
```

See [“About I/O fencing configuration files”](#) on page 672.

## 12 Configure the CP agent on the SF Oracle RAC (application cluster).

```
Do you want to configure CP Agent on the client cluster? [y,n,q]
(y)
```

```
Enter a non-existing name for the service group for CP Agent:
[b] (vxfen)
```

```
Adding CP Agent via galaxy ..... Done
```

- 13 Review the output as the installer stops and restarts VCS with the fencing processes on each application cluster node, and completes the I/O fencing configuration.
- 14 Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

## Setting up server-based I/O fencing manually

Tasks that are involved in setting up server-based I/O fencing manually include:

**Table 9-2** Tasks to set up server-based I/O fencing manually

| Action                                                                          | Description                                                                                               |
|---------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| Preparing the CP servers for use by the SF Oracle RAC cluster                   | See <a href="#">“Preparing the CP servers manually for use by the SF Oracle RAC cluster”</a> on page 195. |
| Modifying I/O fencing configuration files to configure server-based I/O fencing | See <a href="#">“Configuring server-based fencing on the SF Oracle RAC cluster manually”</a> on page 199. |
| Modifying SF Oracle RAC configuration to use I/O fencing                        | See <a href="#">“Modifying VCS configuration to use I/O fencing”</a> on page 180.                         |
| Configuring Coordination Point agent to monitor coordination points             | See <a href="#">“Configuring Coordination Point agent to monitor coordination points”</a> on page 203.    |
| Verifying the server-based I/O fencing configuration                            | See <a href="#">“Verifying server-based I/O fencing configuration”</a> on page 205.                       |

## Preparing the CP servers manually for use by the SF Oracle RAC cluster

Use this procedure to manually prepare the CP server for use by the SF Oracle RAC cluster or clusters.

[Table 9-3](#) displays the sample values used in this procedure.

**Table 9-3** Sample values in procedure

| CP server configuration component | Sample name                |
|-----------------------------------|----------------------------|
| CP server                         | mycps1.symantecexample.com |
| Node #1 - SF Oracle RAC cluster   | galaxy                     |

Table 9-3 Sample values in procedure (continued)

| CP server configuration component | Sample name          |
|-----------------------------------|----------------------|
| Node #2 - SF Oracle RAC cluster   | nebula               |
| Cluster name                      | rac_cluster101       |
| Cluster UUID                      | {f0735332-1dd1-11b2} |

To manually configure CP servers for use by the SF Oracle RAC cluster

- 1 Determine the cluster name and uuid on the SF Oracle RAC cluster.

For example, issue the following commands on one of the SF Oracle RAC cluster nodes (galaxy):

```
# grep cluster /etc/VRTSvcs/conf/config/main.cf

cluster rac_cluster101

# cat /etc/vx/.uuids/clusuuid

{f0735332-1dd1-11b2}
```

- 2 Use the `cpsadm` command to check whether the SF Oracle RAC cluster and nodes are present in the CP server.

For example:

```
# cpsadm -s mycps1.symantecexample.com -a list_nodes

ClusName  UUID                               Hostname(Node ID)  Registered
rac_cluster101 {f0735332-1dd1-11b2} galaxy(0)          0
rac_cluster101 {f0735332-1dd1-11b2} nebula(1)           0
```

If the output does not show the cluster and nodes, then add them as described in the next step.

For detailed information about the `cpsadm` command, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

### 3 Add the SF Oracle RAC cluster and nodes to each CP server.

For example, issue the following command on the CP server (mycps1.symantecexample.com) to add the cluster:

```
# cpsadm -s mycps1.symantecexample.com -a add_clus\  
-c rac_cluster101 -u {f0735332-1dd1-11b2}
```

```
Cluster rac_cluster101 added successfully
```

Issue the following command on the CP server (mycps1.symantecexample.com) to add the first node:

```
# cpsadm -s mycps1.symantecexample.com -a add_node\  
-c rac_cluster101 -u {f0735332-1dd1-11b2} -h galaxy -n0
```

```
Node 0 (galaxy) successfully added
```

Issue the following command on the CP server (mycps1.symantecexample.com) to add the second node:

```
# cpsadm -s mycps1.symantecexample.com -a add_node\  
-c rac_cluster101 -u {f0735332-1dd1-11b2} -h nebula -n1
```

```
Node 1 (nebula) successfully added
```

### 4 If security is to be enabled, check whether the \_HA\_VCS\_users are created in the CP server.

If the output below does not show the users, then add them as described in the next step.

```
# cpsadm -s mycps1.symantecexample.com -a list_users
```

| Username/Domain                                   | Type | Cluster Name / UUID                 | Role     |
|---------------------------------------------------|------|-------------------------------------|----------|
| _HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com/vx |      | rac_cluster101/{f0735332-1dd1-11b2} | Operator |
| _HA_VCS_nebula@HA_SERVICES@nebula.symantec.com/vx |      | rac_cluster101/{f0735332-1dd1-11b2} | Operator |

If security is to be disabled, then add the user name "cpsclient@hostname" to the server instead of the \_HA\_VCS\_users (for example, cpsclient@galaxy).

The CP server can only run in either secure mode or non-secure mode, both connections are not accepted at the same time.

**5** Add the users to the CP server.

First, determine the user@domain to be added on the SF Oracle RAC cluster (application cluster).

The user for fencing should be of the form `_HA_VCS_short-hostname` and domain name is that of `HA_SERVICES` user in the output of command:

```
# /opt/VRTScps/bin/cpsat listpd -t local
```

Next, issue the following commands on the CP server (mycps1.symantecexample.com):

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com  
successfully added
```

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_nebula@HA_SERVICES@nebula.symantec.com  
successfully added
```

- 6 Authorize the CP server user to administer the SF Oracle RAC cluster. You must perform this task for the CP server users corresponding to each node in the SF Oracle RAC cluster.

For example, issue the following command on the CP server (mycps1.symantecexample.com) for SF Oracle RAC cluster rac\_cluster101 with two nodes galaxy and nebula:

```
# cpsadm -s mycps1.symantecexample.com -a\  
add_clus_to_user -c rac_cluster101\  
-u {f0735332-1dd1-11b2}\  
-e _HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com\  
-f cps_operator -g vx
```

```
Cluster successfully added to user  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com privileges.
```

```
# cpsadm -s mycps1.symantecexample.com -a\  
add_clus_to_user -c rac_cluster101\  
-u {f0735332-1dd1-11b2}\  
-e _HA_VCS_nebula@HA_SERVICES@nebula.symantec.com\  
-f cps_operator -g vx
```

```
Cluster successfully added to user  
_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com privileges.
```

## Configuring server-based fencing on the SF Oracle RAC cluster manually

The configuration process for the client or SF Oracle RAC cluster to use CP server as a coordination point requires editing the `/etc/vxsfenmode` file.

You need to edit this file to specify the following information for your configuration:

- Fencing mode
- Fencing mechanism
- Fencing disk policy (if applicable to your I/O fencing configuration)
- Appropriate value for the security configuration
- CP server or CP servers
- Coordinator disk group (if applicable to your I/O fencing configuration)

---

**Note:** Whenever coordinator disks are used as coordination points in your I/O fencing configuration, you must create a disk group (vxfgndg). You must specify this disk group in the `/etc/vxfenmode` file.

See [“Setting up coordinator disk groups”](#) on page 178.

---

The customized fencing framework also generates the `/etc/vxfentab` file which has security setting and the coordination points (all the CP servers and disks from disk group specified in `/etc/vxfenmode` file).

### To configure server-based fencing on the SF Oracle RAC cluster manually

- 1 Use a text editor to edit the following file on each node in the cluster:

`/etc/default/vxfen`

You must change the values of the `VXFEN_START` and the `VXFEN_STOP` environment variables to 1.

- 2 Use a text editor to edit the `/etc/vxfenmode` file values to meet your configuration specifications.

If your server-based fencing configuration uses a single highly available CP server as its only coordination point, make sure to add the `single_cp=1` entry in the `/etc/vxfenmode` file.

The following sample file output displays what the `/etc/vxfenmode` file contains:

See [“Sample vxfenmode file output”](#) on page 200.

- 3 After editing the `/etc/vxfenmode` file, run the `vxfen init` script to start fencing.

For example:

On Solaris 9 systems:

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

On Solaris 10 systems:

```
# svcadm enable vxfen
```

## Sample vxfenmode file output

The following sample file output displays

```
#  
# vxfen_mode determines in what mode VCS I/O Fencing should work.  
#
```



```
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=customized

# vxfen_mechanism determines the mechanism for customized I/O
# fencing that should be used.
#
# available options:
# cps        - use a coordination point server with optional script
#              controlled scsi3 disks
#
vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
scsi3_disk_policy=dmp

# security when enabled uses secure communication to the cp server
# using VxAT (Veritas Authentication Service)
# available options:
# 0 - don't use Veritas Authentication Service for cp server
#    communication
# 1 - use Veritas Authentication Service for cp server
#    communication
security=1

#
# Specify 3 or more odd number of coordination points in this file,
# one in each row. They can be all-CP servers, all-SCSI-3 compliant
# coordinator disks, or a combination of CP servers and SCSI-3
# compliant coordinator disks. Please ensure that the CP server
# coordination points are numbered sequentially and in the same
# order on all the cluster nodes.
```

```
#
# Coordination Point Server(CPS) is specified as:
#
# cps<number>=<Virtual IP/ Virtual hostname of cp server> in
# square brackets ([]), followed by ":" and CPS port number.
#
# Examples:
# cps1=[192.168.0.23]:14250
# cps2=[mycps.company.com]:14250
#
# SCSI-3 compliant coordinator disks are specified as:
#
# vxfendg=<coordinator disk group name>
# Example:
# vxfendg=vxfencoorddg
#
# Examples of different configurations:
# 1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
#
```

Table 9-4 defines the vxfenmode parameters that must be edited.

Table 9-4 vxfenmode file parameters

| vxfenmode File Parameter | Description                                                            |
|--------------------------|------------------------------------------------------------------------|
| vxfen_mode               | Fencing mode of operation. This parameter must be set to “customized”. |

**Table 9-4** vxfenmode file parameters (*continued*)

| vxfenmode File Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| vxfen_mechanism              | Fencing mechanism. This parameter defines the mechanism that is used for fencing. If one of the three coordination points is a CP server, then this parameter must be set to "cps".                                                                                                                                                                                                                                                                                                                                                                  |
| scsi3_disk_policy            | Configure the vxfen module to use either DMP devices, "dmp" or the underlying raw character devices, "raw".<br><b>Note:</b> The configured disk policy is applied on all the nodes.                                                                                                                                                                                                                                                                                                                                                                  |
| security                     | Security parameter 1 indicates that Symantec Product Authentication Service is used for CP server communications.<br><br>Security parameter 0 indicates that communication with the CP server is made in non-secure mode.<br><br>The default security value is 1.<br><b>Note:</b> Symantec only supports a configuration where both the CP server and client sides have the same security setting. The security setting on both sides must be either enabled or disabled.                                                                            |
| cps1, cps2, cps3, or vxfendg | Coordination point parameters.<br><br>Enter either the Virtual IP address or FQHN (whichever is accessible) of the CP server.<br><br><b>Note:</b> Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxfendg) and specified in the /etc/vxfenmode file. Additionally, the customized fencing framework also generates the /etc/vxfentab file which specifies the security setting and the coordination points (all the CP servers and the disks from disk group specified in /etc/vxfenmode file). |
| single_cp                    | Value 1 for single_cp parameter indicates that the server-based fencing uses a single highly available CP server as its only coordination point.<br><br>Value 0 for single_cp parameter indicates that the server-based fencing uses at least three coordination points.                                                                                                                                                                                                                                                                             |

## Configuring Coordination Point agent to monitor coordination points

The following procedure describes how to manually configure the Coordination Point agent to monitor coordination points (CP server or SCSI-3 disks).

**To configure Configuration Point agent to monitor coordination points**

- 1** Ensure that your SF Oracle RAC cluster has been properly installed and configured with fencing enabled.
- 2** Create a parallel service group vxfen and add a coordpoint resource to the vxfen service group using the following commands:

```
# haconf -makerw
# hagrps -add vxfen
# hagrps -modify vxfen SystemList galaxy 0 nebula 1
# hagrps -modify vxfen AutoFailOver 0
# hagrps -modify vxfen Parallel 1
# hagrps -modify vxfen SourceFile "./main.cf"
# hares -add coordpoint CoordPoint vxfen
# hares -modify coordpoint FaultTolerance 1
# hares -modify coordpoint Enabled 1
# haconf -dump -makero
```

- 3 Verify the status of the agent on the SF Oracle RAC cluster using the `hares` commands. For example:

```
# hares -state coordpoint
```

The following is an example of the command and output:

```
# hares -state coordpoint

# Resource      Attribute    System      Value
coordpoint     State       galaxy      ONLINE
coordpoint     State       nebula      ONLINE
```

- 4 Access the engine log to view the agent log. The agent log is written to the engine log.

The agent log contains detailed Coordination Point agent monitoring information; including information about whether the Coordination Point agent is able to access all the coordination points, information to check on which coordination points the Coordination Point agent is reporting missing keys, etc.

To view all such information in the engine log, change the `dbg` level for that node using the following commands:

```
# haconf -makerw

# hatype -modify Coordpoint LogDbg 10

# haconf -dump -makero
```

The agent log can now be viewed at the following location:

`/var/VRTSvcs/log/engine_A.log`

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on the agent.

## Verifying server-based I/O fencing configuration

Follow the procedure described below to verify your server-based I/O fencing configuration.

### To verify the server-based I/O fencing configuration

- 1 Verify that the I/O fencing configuration was successful by running the `vxfenadm` command.

For example, run the following command:

```
# vxfenadm -d
```

---

**Note:** For troubleshooting any server-based I/O fencing configuration issues, refer to the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

---

- 2 Verify that I/O fencing is using the specified coordination points by running the `vxfenconfig` command.

For example, run the following command:

```
# vxfenconfig -l
```

If the output displays `single_cp=1`, it indicates that the application cluster uses a CP server as the single coordination point for server-based fencing.

# Performing post-installation and configuration tasks

This chapter includes the following topics:

- [Installing language packages](#)
- [About enabling LDAP authentication for clusters that run in secure mode](#)
- [Configuring Veritas Volume Replicator](#)
- [Verifying SF Oracle RAC installation using VCS configuration file](#)
- [Verifying LLT, GAB, and cluster operation](#)
- [Running VRTSexplorer to diagnose issues in the cluster](#)

## Installing language packages

To install SF Oracle RAC in a language other than English, install the required language packages after installing the English packages.

**To install the language packages on the server**

- 1 Make sure the VEA Service is not running.

```
# /opt/VRTS/bin/vxsvcctl status
Current state of server : RUNNING
```

- 2 If the VEA Service is running, stop it by using the `vxsvcctl stop` command.

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

- 3 Insert the "Language" disc into the DVD-ROM or CD-ROM drive. With Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.

- 4 Install the language packages using the `install_lp` command.

```
# cd /cdrom/cdrom0  
# ./install_lp
```

- 5 Restart the VEA Service.

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

## About enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the authentication broker. AT supports all common LDAP distributions such as Sun Directory Server, Netscape, OpenLDAP, and Windows Active Directory.

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.

See [“Enabling LDAP authentication for clusters that run in secure mode”](#) on page 210.


If you have not already added VCS users during installation, you can add the users later.

See the *Veritas Cluster Server Administrator's Guide* for instructions to add VCS users.

[Figure 10-1](#) depicts the SF Oracle RAC cluster communication with the LDAP servers when clusters run in secure mode.



**Figure 10-1** Client communication with LDAP servers



See the *Symantec Product Authentication Service Administrator's Guide*.

The LDAP schema and syntax for LDAP commands (such as, `ldapadd`, `ldapmodify`, and `ldapsearch`) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)
  - UserObjectClass (the default is `posixAccount`)
  - UserObject Attribute (the default is `uid`)
  - User Group Attribute (the default is `gidNumber`)
  - Group Object Class (the default is `posixGroup`)
  - GroupObject Attribute (the default is `cn`)
  - Group GID Attribute (the default is `gidNumber`)
  - Group Membership Attribute (the default is `memberUid`)
- URL to the LDAP Directory

- Distinguished name for the user container (for example, UserBaseDN=ou=people,dc=comp,dc=com)
- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

## Enabling LDAP authentication for clusters that run in secure mode

The following procedure shows how to enable the plug-in module for LDAP authentication. This section provides examples for OpenLDAP and Windows Active Directory LDAP distributions.

Before you enable the LDAP authentication, complete the following steps:

- Make sure that the cluster runs in secure mode.

```
# haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 5.0.32.0 or later.

```
# /opt/VRTSat/bin/vssat showversion  
vssat version: 5.0.32.0
```

See the `vssat.1m` and the `atldapconf.1m` manual pages.

## To enable OpenLDAP authentication for clusters that run in secure mode

- 1 Add the LDAP domain to the AT configuration using the `vssat` command.

The following example adds the LDAP domain, MYENTERPRISE:

```
# /opt/VRTSat/bin/vssat addldapdomain \
--domainname "MYENTERPRISE.symantecdomain.com"\
--server_url "ldap://my_openldap_host.symantecexample.com"\
--user_base_dn "ou=people,dc=symantecdomain,dc=myenterprise,dc=com"\
--user_attribute "cn" --user_object_class "account"\
--user_gid_attribute "gidNumber"\
--group_base_dn "ou=group,dc=symantecdomain,dc=myenterprise,dc=com"\
--group_attribute "cn" --group_object_class "posixGroup"\
--group_gid_attribute "member"\
--admin_user "cn=manager,dc=symantecdomain,dc=myenterprise,dc=com"\
--admin_user_password "password" --auth_type "FLAT"
```

- 2 Verify that you can successfully authenticate an LDAP user on the SF Oracle RAC nodes.

You must have a valid LDAP user ID and password to run the command. In the following example, authentication is verified for the MYENTERPRISE domain for the LDAP user, vcsadmin1.

```
galaxy# /opt/VRTSat/bin/vssat authenticate
--domain ldap:MYENTERPRISE.symantecdomain.com
--prplname vcsadmin1 --broker galaxy:2821
```

Enter password for vcsadmin1: #####

```
authenticate
-----
-----
```

```
Authenticated User vcsadmin1
-----
```

**3** Add the LDAP user to the main.cf file.

```
# haconf makerw
# hauser -add "CN=vcsadmin1/CN=people/\
DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" -priv Administrator
# haconf -dump -makero
```

If you want to enable group-level authentication, you must run the following command:

```
# hauser -addpriv \
ldap_group@ldap_domain AdministratorGroup
```

**4** Verify that the main.cf file has the following lines:

```
# cat /etc/VRTSvcs/conf/config/main.cf
...
...
cluster rac_cluster101 (
  SecureClus = 1
  Administrators = {
    "CN=vcsadmin1/CN=people/DC=symantecdomain/DC=myenterprise/
    DC=com@myenterprise.symantecdomain.com" }
  AdministratorGroups = {
    "CN=symantecusergroups/DC=symantecdomain/DC=myenterprise/
    DC=com@myenterprise.symantecdomain.com " }
)
...
...
```

**5** Set the VCS\_DOMAIN and VCS\_DOMAINTYPE environment variables as follows:

- VCS\_DOMAIN=myenterprise.symantecdomain.com
- VCS\_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
# export VCS_DOMAIN=myenterprise.symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

## 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute    Value
galaxy       Attribute    RUNNING
nebula       Attribute    RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SF Oracle RAC node using the VCS Cluster Manager (Java Console).

## 7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

**To enable Windows Active Directory authentication for clusters that run in secure mode**

- 1 Run the LDAP configuration tool `atldapconf` using the `-d` option. The `-d` option discovers and retrieves an LDAP properties file which is a prioritized attribute list.

```
# /opt/VRTSat/bin/atldapconf -d
-s domain_controller_name_or_ipaddress
-u domain_user -g domain_group
```

For example:

```
# /opt/VRTSat/bin/atldapconf -d -s 192.168.20.32 \
-u Administrator -g "Domain Admins"
```

Search User provided is invalid or Authentication is required to proceed further.

Please provide authentication information for LDAP server.

Username/Common Name: **symantecdomain\administrator**

Password:

Attribute file created.

- 2 Run the LDAP configuration tool `atldapconf` using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSat/bin/atldapconf -c -d windows_domain_name
```

For example:

```
# /opt/VRTSat/bin/atldapconf -c -d symantecdomain.com
```

Attribute list file not provided, using default AttributeList.txt.

CLI file name not provided, using default CLI.txt.

CLI for addldapdomain generated.

- 3 Run the LDAP configuration tool `atldapconf` using the `-x` option. The `-x` option reads the CLI file and executes the commands to add a domain to the AT.

```
# /opt/VRTSat/bin/atldapconf -x
```

- 4 List the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSat/bin/vssat listldapdomains
```

```
Domain Name :          symantecdomain.com
Server URL :          ldap://192.168.20.32:389
SSL Enabled :          No
User Base DN :         CN=people,DC=symantecdomain,DC=com
User Object Class :    account
User Attribute :       cn
User GID Attribute :   gidNumber
Group Base DN :        CN=group,DC=symantecdomain,DC=com
Group Object Class :   group
Group Attribute :      cn
Group GID Attribute :  cn
Auth Type :            FLAT
Admin User :
Admin User Password :
Search Scope :         SUB
```

- 5 Set the VCS\_DOMAIN and VCS\_DOMAINTYPE environment variables as follows:

- VCS\_DOMAIN=symantecdomain.com

- VCS\_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
# export VCS_DOMAIN=symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute      Value
galaxy       Attribute      RUNNING
nebula       Attribute      RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SF Oracle RAC node using the VCS Cluster Manager (Java Console).

- 7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

## Configuring Veritas Volume Replicator

Perform this step only if you have not already configured VVR during the installation.

By default, the installer installs the required VVR configuration files irrespective of whether or not you choose to enable VVR. To configure VVR manually in SF Oracle RAC, simply start VVR using the `vxstart_vvr` command . The command starts the VVR daemons and configures the ports. You may change the default settings at any time.

For instructions on changing the default settings, see the *Veritas Volume Replicator Administrator's Guide*.

### To configure VVR

- 1 Log into each node in the cluster as the root user.
- 2 Start VVR:

```
# vxstart_vvr start
VxVM VVR INFO V-5-2-3935 Using following ports:
heartbeat: 4145
vradmind: 8199
vxrsyncd: 8989
data: Anonymous-Ports
To change, see vrport(1M) command
VxVM VVR V-5-2-3728 Starting Communication daemon: [OK]
VxVM VVR V-5-4-2415 Starting VRAS daemon: [OK]
VxVM VVR V-5-4-2415 Starting VxRSyncd daemon: [OK]
```



## Verifying SF Oracle RAC installation using VCS configuration file

The configuration file, `main.cf`, is created on each node at `/etc/VRTSvcs/conf/config/`. Review the `main.cf` configuration file after the SF Oracle RAC installation and before the Oracle installation.

Verify the following information in the `main.cf` file:

- The cluster definition within the `main.cf` includes the cluster information that was provided during the configuration. The cluster information includes the cluster name, cluster address, and the names of cluster users and administrators.
- The `UseFence = SCSI3` attribute is not automatically present; you must manually add it after the installation.
- If you configured the cluster in secure mode, the `main.cf` includes the `VxSS` service group and “`SecureClus = 1`” cluster attribute.

For more information on the configuration file:

See [“About VCS configuration file”](#) on page 669.

## Verifying LLT, GAB, and cluster operation

Verify the operation of LLT, GAB, and the cluster using the VCS commands.

**To verify LLT, GAB, and cluster operation**

- 1 Log in to any node in the cluster as superuser.
- 2 Make sure that the `PATH` environment variable is set to run the VCS commands.
- 3 On Solaris 9, if you use Sun SCI adapters for your private network, move the scripts `s70llt` and `s92gab` from the directory `/etc/rc2.d` to directory `/etc/rc3.d`, so that they are run after the `s19sci` and `s23scid` scripts.
- 4 Verify LLT operation.  
See [“Verifying LLT”](#) on page 218.
- 5 Verify GAB operation.  
See [“Verifying GAB”](#) on page 220.
- 6 Verify the cluster operation.  
See [“Verifying the cluster”](#) on page 221.

## Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

### To verify LLT

- 1 Log in as superuser on the node galaxy.
- 2 Run the `lltstat` command on the node galaxy to view the status of LLT.

```
lltstat -n
```

The output on galaxy resembles:

```
LLT node information:
Node           State      Links
*0 galaxy      OPEN      2
 1 nebula      OPEN      2
```

Each node has two links and each node is in the OPEN state. The asterisk (\*) denotes the node on which you typed the command.

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

```
LLT node information:
Node           State      Links
* 0 galaxy      OPEN      2
 1 nebula      OPEN      2
 2 saturn       OPEN      1
```

- 3 Log in as superuser on the node nebula.
- 4 Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

```
LLT node information:
Node           State      Links
 0 galaxy      OPEN      2
*1 nebula      OPEN      2
```

- 5 To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node `galaxy` in a two-node cluster:

```
lltstat -nvv active
```

The output on `galaxy` resembles the following:

■ For Solaris SPARC:

| Node      | State | Link | Status | Address           |
|-----------|-------|------|--------|-------------------|
| *0 galaxy | OPEN  |      |        |                   |
|           |       | bge1 | UP     | 08:00:20:93:0E:34 |
|           |       | bge2 | UP     | 08:00:20:93:0E:34 |
| 1 nebula  | OPEN  |      |        |                   |
|           |       | bge1 | UP     | 08:00:20:8F:D1:F2 |
|           |       | bge2 | DOWN   |                   |

■ For Solaris x64:

| Node      | State | Link     | Status | Address           |
|-----------|-------|----------|--------|-------------------|
| *0 galaxy | OPEN  |          |        |                   |
|           |       | e1000g:1 | UP     | 08:00:20:93:0E:34 |
|           |       | e1000g:2 | UP     | 08:00:20:93:0E:38 |
| 1 nebula  | OPEN  |          |        |                   |
|           |       | e1000g:1 | UP     | 08:00:20:8F:D1:F2 |
|           |       | e1000g:2 | DOWN   |                   |

The command reports the status on the two active nodes in the cluster, `galaxy` and `nebula`.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- A MAC address for each link

However, the output in the example shows different details for the node `nebula`. The private network connection is possibly broken or the information in the `/etc/llttab` file may be incorrect.

- 6 To obtain information about the ports open for LLT, type `lltstat -p` on any node.

For example, type `lltstat -p` on the node `galaxy` in a two-node cluster:

```
lltstat -p
```

The output resembles:

```
LLT port information:
Port  Usage      Cookie
0     gab        0x0
      opens:    0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects: 0 1
7     gab        0x7
      opens:    0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects: 0 1
63    gab        0x1F
      opens:    0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects: 0 1
```

## Verifying GAB

Verify the GAB operation using the `gabconfig -a` command. This command returns the GAB port membership information. The output displays the nodes that have membership with the modules you installed and configured. You can use GAB port membership as a method of determining if a specific component of the SF Oracle RAC stack communicates with its peers.

[Table 10-1](#) lists the different ports that the software configures for different functions.

**Table 10-1** GAB port description

| Port | Function                                               |
|------|--------------------------------------------------------|
| a    | GAB                                                    |
| b    | I/O fencing                                            |
| d    | Oracle Disk Manager (ODM)                              |
| f    | Cluster File System (CFS)                              |
| h    | Veritas Cluster Server (VCS: High Availability Daemon) |
| o    | VCSMM driver                                           |

**Table 10-1** GAB port description (*continued*)

| Port | Function                                                                                                                                                                               |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| u    | Cluster Volume Manager (CVM)<br>(to ship commands from slave node to master node)<br><br>Port u in the <code>gabconfig</code> output is visible with CVM protocol version $\geq 100$ . |
| v    | Cluster Volume Manager (CVM)                                                                                                                                                           |
| w    | vxconfigd (module for CVM)                                                                                                                                                             |

For more information on GAB, refer to the *Veritas Cluster Server Administrator's Guide*.

### To verify GAB

- ◆ To verify the GAB operation, type the following command on each node:

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

For example, the command returns the following output:

```
GAB Port Memberships
=====
Port a gen  ada401 membership 01

Port b gen  ada40d membership 01

Port d gen  ada409 membership 01
Port f gen  ada41c membership 01
Port h gen  ada40f membership 01

Port o gen  ada406 membership 01

Port u gen  ada41a membership 01

Port v gen  ada416 membership 01
Port w gen  ada418 membership 01
```

## Verifying the cluster

Verify the status of the cluster using the `hastatus` command. This command returns the system state and the group state.

Refer to the `hastatus(1M)` manual page.

Refer to the *Veritas Cluster Server Administrator's Guide* for a description of system states and the transitions between them.

To verify the cluster

- 1 To verify the status of the cluster, type the following command:

```
hasstatus -summary
```

The output resembles:

|                 |        |         |        |              |        |
|-----------------|--------|---------|--------|--------------|--------|
| -- SYSTEM STATE |        |         |        |              |        |
| -- System       |        | State   |        | Frozen       |        |
| A               | galaxy | RUNNING |        | 0            |        |
| A               | nebula | RUNNING |        | 0            |        |
|                 |        |         |        |              |        |
| -- GROUP STATE  |        |         |        |              |        |
| -- Group        |        | System  | Probed | AutoDisabled | State  |
| B               | cvm    | galaxy  | Y      | N            | ONLINE |
| B               | cvm    | nebula  | Y      | N            | ONLINE |
|                 |        |         |        |              |        |
| B               | VxSS   | galaxy  | Y      | N            | ONLINE |
| B               | VxSS   | nebula  | Y      | N            | ONLINE |

Note that the VxSS service group is displayed only if you have configured the cluster in secure mode.

- 2 Review the command output for the following information:
- The system state  
If the value of the system state is RUNNING, the cluster is successfully started.
  - The cvm group state  
In the sample output, the group state lists the cvm group, which is ONLINE on both the nodes galaxy and nebula.

## Verifying the cluster nodes

Verify the information of the cluster systems using the `hasys -display` command. The information for each node in the output should be similar.

Refer to the `hasys(1M)` manual page.

Refer to the *Veritas Cluster Server Administrator's Guide* for information about the system attributes for VCS.

---

**Note:** The example in the following procedure is for SPARC. x64 clusters have different command output.

---

### To verify the cluster nodes

- ◆ On one of the nodes, type the `hasys -display` command:

```
hasys -display
```

The example shows the output when the command is run on the node galaxy. The list continues with similar information for nebula (not shown) and any other nodes in the cluster.

| #System | Attribute          | Value                                                                                               |
|---------|--------------------|-----------------------------------------------------------------------------------------------------|
| galaxy  | AgentsStopped      | 0                                                                                                   |
| galaxy  | AvailableCapacity  | 100                                                                                                 |
| galaxy  | CPUBinding         | BindTo None CPUNumber 0                                                                             |
| galaxy  | CPUThresholdLevel  | Critical 90 Warning 80 Note 70<br>Info 60                                                           |
| galaxy  | CPUUsage           | 0                                                                                                   |
| galaxy  | CPUUsageMonitoring | Enabled 0 ActionThreshold 0<br>ActionTimeLimit 0 Action NONE<br>NotifyThreshold 0 NotifyTimeLimit 0 |
| galaxy  | Capacity           | 100                                                                                                 |
| galaxy  | ConfigBlockCount   | 130                                                                                                 |
| galaxy  | ConfigChecksum     | 46688                                                                                               |
| galaxy  | ConfigDiskState    | CURRENT                                                                                             |
| galaxy  | ConfigFile         | /etc/VRTSvcscs/conf/config                                                                          |
| galaxy  | ConfigInfoCnt      | 0                                                                                                   |
| galaxy  | ConfigModDate      | Wed 14 Oct 2009 17:22:48                                                                            |
| galaxy  | ConnectorState     | Down                                                                                                |
| galaxy  | CurrentLimits      |                                                                                                     |
| galaxy  | DiskHbStatus       |                                                                                                     |

|        |                    |                                                |
|--------|--------------------|------------------------------------------------|
| galaxy | DynamicLoad        | 0                                              |
| galaxy | EngineRestarted    | 0                                              |
| galaxy | EngineVersion      | 5.1.10.0                                       |
| galaxy | FencingWeight      | 0                                              |
| galaxy | Frozen             | 0                                              |
| galaxy | GUIIPAddr          |                                                |
| galaxy | HostUtilization    | CPU 0 Swap 0                                   |
| galaxy | LLTNodeId          | 0                                              |
| galaxy | LicenseType        | DEMO                                           |
| galaxy | Limits             |                                                |
| galaxy | LinkHbStatus       |                                                |
| galaxy | LoadTimeCounter    | 0                                              |
| galaxy | LoadTimeThreshold  | 600                                            |
| galaxy | LoadWarningLevel   | 80                                             |
| galaxy | NoAutoDisable      | 0                                              |
| galaxy | NodeId             | 0                                              |
| galaxy | OnGrpCnt           | 1                                              |
| galaxy | ShutdownTimeout    |                                                |
| galaxy | SourceFile         | ./main.cf                                      |
| galaxy | SwapThresholdLevel | Critical 90 Warning 80 Note 70<br>Info 60      |
| galaxy | SysInfo            | Solaris:galaxy,Generic_<br>118558-11,5.9,sun4u |
| galaxy | SysName            | galaxy                                         |
| galaxy | SysState           | RUNNING                                        |
| galaxy | SystemLocation     |                                                |
| galaxy | SystemOwner        |                                                |



|        |             |         |
|--------|-------------|---------|
| galaxy | TFrozen     | 0       |
| galaxy | TRSE        | 0       |
| galaxy | UpDownState | Up      |
| galaxy | UserInt     | 0       |
| galaxy | UserStr     |         |
| galaxy | VCSFeatures | DR      |
| galaxy | VCSMode     | VCS_RAC |

## Running VRTSexplorer to diagnose issues in the cluster

Perform this step if you encounter issues in your cluster environment that require professional support.

VRTSexplorer, also known as Veritas Explorer, is a tool provided by Symantec to gather system and configuration information from a node to diagnose or analyze issues in the cluster. The utility is located at `/opt/VRTSspt/VRTSexplorer`.

---

**Note:** The utility collects only technical information such as the server configuration and installed Veritas products information required to troubleshoot issues. No user data is collected by the utility. For more details on the information that is collected by the utility, refer to the README located in the `/opt/VRTSspt/VRTSexplorer` directory.

---

To run the utility:

```
# cd /opt/VRTSspt/VRTSexplorer
# ./VRTSexplorer
```

For information on various command options, see `VRTSexplorer (1M)`.

If you find issues in the cluster that require professional help, run the utility and send the compressed tar file output to Symantec Technical Support to resolve the issue.



## Upgrade of SF Oracle RAC

- [Chapter 11. Preparing to upgrade to SF Oracle RAC 5.1 SP1](#)
- [Chapter 12. Performing a full upgrade to SF Oracle RAC 5.1 SP1](#)
- [Chapter 13. Performing a phased upgrade to SF Oracle RAC 5.1 SP1](#)
- [Chapter 14. Performing a rolling upgrade to SF Oracle RAC 5.1 SP1](#)
- [Chapter 15. Upgrading to SF Oracle RAC 5.1 SP1 using Live Upgrade](#)
- [Chapter 16. Upgrading from Storage Foundation products to SF Oracle RAC 5.1 SP1](#)
- [Chapter 17. Migrating from single instance Storage Foundation for Oracle HA to SF Oracle RAC](#)
- [Chapter 18. Performing post-upgrade tasks](#)



# Preparing to upgrade to SF Oracle RAC 5.1 SP1

This chapter includes the following topics:

- [About upgrading SF Oracle RAC](#)
- [About types of upgrade](#)
- [Supported upgrade paths](#)
- [Preparing to upgrade to SF Oracle RAC 5.1 SP1](#)
- [Pre-upgrade tasks for migrating the SFDB repository database](#)

## About upgrading SF Oracle RAC

SF Oracle RAC supports various ways of upgrading your cluster to the latest version. Choose a method that best suits your environment and planned upgrade path.

For the supported types of upgrade:

See [“About types of upgrade”](#) on page 230.

For the supported upgrade paths:

See [“Supported upgrade paths”](#) on page 230.

---

**Note:** SF Oracle RAC software must be at the same version across all nodes in an SF Oracle RAC cluster after the upgrade, that is 5.1 Service Pack 1.

---

# About types of upgrade

Table 11-1 lists the supported types of upgrade.

Table 11-1           Types of upgrade

| Type of upgrade      | Method of upgrade                                                                                                                                                 | Procedures                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Full upgrade         | Veritas script-based installation programs <ul style="list-style-type: none"><li>■ Interactive mode</li><li>■ Non-interactive mode using response files</li></ul> | Complete the following steps: <ul style="list-style-type: none"><li>■ Preparing to upgrade<br/>See “Preparing to upgrade to SF Oracle RAC 5.1 SP1” on page 231.</li><li>■ Upgrading to SF Oracle RAC 5.1 SP1<br/>See the chapter <i>Performing a full upgrade to SF Oracle RAC 5.1 SP1</i>.</li><li>■ Completing post-upgrade tasks<br/>See the chapter "Performing post-upgrade tasks".</li></ul> |
| Phased upgrade       | Combination of manual steps and the Veritas script-based installation programs                                                                                    | Complete the steps in the chapter <i>Performing a phased upgrade to SF Oracle RAC 5.1 SP1</i> .                                                                                                                                                                                                                                                                                                    |
| Rolling upgrade      | Veritas script-based installation programs                                                                                                                        | Complete the steps in the chapter <i>Performing a rolling upgrade to SF Oracle RAC 5.1 SP1</i> .                                                                                                                                                                                                                                                                                                   |
| Solaris Live Upgrade | Combination of native operating system upgrade mechanism and the Veritas script-based installation programs                                                       | Complete the following steps: <ul style="list-style-type: none"><li>■ Upgrading to SF Oracle RAC 5.1 SP1<br/>See the chapter <i>Upgrading to SF Oracle RAC 5.1 SP1 using Live Upgrade</i>.</li><li>■ Completing post-upgrade tasks<br/>See the chapter "Performing post-upgrade tasks".</li></ul>                                                                                                  |

# Supported upgrade paths

Table 11-2 lists the supported upgrade paths.

Table 11-2           Supported upgrade paths

| From SF Oracle RAC version   | To SF Oracle RAC version | Supported upgrade type |
|------------------------------|--------------------------|------------------------|
| 4.1 MP2 (Solaris SPARC only) | 5.1 SP1                  | Full or phased upgrade |

**Table 11-2** Supported upgrade paths (*continued*)

| From SF Oracle RAC version                                              | To SF Oracle RAC version | Supported upgrade type  |
|-------------------------------------------------------------------------|--------------------------|-------------------------|
| 5.0 and later (including maintenance packs and rolling patches on 5.0 ) | 5.1 SP1                  | Full or phased upgrade  |
| 5.1 and later (including maintenance packs and rolling patches on 5.1 ) | 5.1 SP1                  | Full or rolling upgrade |

## Preparing to upgrade to SF Oracle RAC 5.1 SP1

Perform the preparatory steps in this section if you are performing a full upgrade of the cluster. Before you upgrade, make sure that your systems meet the hardware and software requirements for this release.

### To prepare to upgrade SF Oracle RAC

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```
# 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
# cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
/etc/VRTSvcs/conf/config/OracleTypes.cf.save
# cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
/var/tmp/PrivNIC.cf.save
# cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
/var/tmp/MultiPrivNIC.cf.save
```

- 3 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.

If the applications are under VCS control:

```
# hagrps -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

- 4 For Oracle RAC 9i, log in as Oracle user on each node and stop `gsd`:

```
$ $ORACLE_HOME/bin/psedctl stop
```

- 5 For Oracle RAC 10g and Oracle RAC 11g:

Stop all Oracle RAC resources.

- If the database instances are not managed by VCS, then run the following on one node:

```
$ srvctl stop database -d db_name
```

- If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

```
# hagrps -offline group_name -any
```

- 6 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw
```

```
# hagrps -modify oracle_group AutoStart 0
```

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

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

```
$ srvctl modify database -d db_name -y manual
```



- 7 Unmount the VxFS file system, which is not under VCS control.

```
# mount -v |grep vxfs

# fuser -c /mount_point

# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

- 8 Stop VCS on all nodes:

```
# hastop -all
```

- 9 If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to SF Oracle RAC 5.1 SP1.

## Pre-upgrade tasks for migrating the SFDB repository database

If you plan to continue using checkpoints or SmartTier for Oracle policies you created with a 5.0x or earlier version of Storage Foundation for Oracle RAC, you must migrate the SFDB repository database to 5.1 Service Pack 1.

If you are upgrading from 5.1 to 5.1 Service Pack 1, no upgrade steps are required for the SFDB tools.

---

**Note:** The Sfua\_Base repository resource group will be removed from the main.cf file. It is not required as a separate service group for SF Oracle RAC 5.1 Service Pack 1.

---

Perform the following before upgrading SF Oracle RAC.

### To prepare to migrate the repository database

- ◆ Resynchronize all existing snapshots before upgrading. As Oracle user, enter:

```
$ /opt/VRTS/bin/dbed_vmsnap -S $ORACLE_SID \  
-f SNAPPLAN -o resync
```

---

**Warning:** The Database Flashsnap clone database will not be able to be carried over after upgrading. You must create a new Database Flashsnap clone database after upgrading to 5.1 Service Pack 1.

---

# Performing a full upgrade to SF Oracle RAC 5.1 SP1

This chapter includes the following topics:

- [About full upgrades](#)
- [Upgrading to SF Oracle RAC 5.1 SP1](#)

## About full upgrades

A full upgrade involves upgrading all the nodes in the cluster at the same time. The cluster remains unavailable for the duration of the upgrade.

You can perform the upgrade using one of the following Veritas script-based installation programs:

- **Common product installer** (`installer`)  
The common product installer provides menu options for installing and configuring multiple Veritas products.
- **SF Oracle RAC installation script** (`installsfprac`)  
The SF Oracle RAC installation script provides menu options for installing and configuring SF Oracle RAC.

---

**Note:** If you obtained SF Oracle RAC from an electronic download site, you must use the product installer (`installsfprac`) instead of the common product installer (`installer`).

---

You can also perform a full upgrade using a response file. You can create a response file using the response file template or customize an installer-generated response file.

---

**Note:** After you upgrade to version 5.1 SP1 from version 5.1 or related 5.1 rolling patches, Symantec recommends that you do not roll back to version 5.1 or corresponding rolling patches.

---

For more information about response files:

See [“About response files”](#) on page 425.

## Upgrading to SF Oracle RAC 5.1 SP1

Perform the steps in the following procedure to upgrade to SF Oracle RAC 5.1 SP1.

---

**Note:** For Solaris SPARC: If you are upgrading from SF Oracle RAC version 4.1 MP1 or earlier, first upgrade to SF Oracle RAC version 4.1 MP2, then upgrade to SF Oracle RAC 5.1 SP1.

For instructions, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide (4.1 Maintenance Pack 2)*.

Contact Technical Support to obtain the software media or to download the software for version 4.1 MP2.

---

### To upgrade to SF Oracle RAC 5.1 SP1

- 1 Upgrade the operating system, if required.  
For instructions, see the operating system documentation.
- 2 Upgrade to SF Oracle RAC 5.1 SP1.  
See [“Upgrading SF Oracle RAC using the Veritas script-based installation program”](#) on page 238.  
You can also perform a silent upgrade:  
See [“Upgrading SF Oracle RAC using a response file”](#) on page 240.
- 3 Relink the SF Oracle RAC libraries with Oracle.  
See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 301.
- 4 Bring the Oracle database service group online.
  - If the Oracle database is managed by VCS:  

```
# hagrps -online oracle_group -sys node_name
```

- If the Oracle database is not managed by VCS:

```
# srvctl start database -d db_name
```

5 Manually mount the VxFS and CFS file systems that are not managed by VCS.

6 Start all applications that are not managed by VCS. Use native application commands to start the applications.

- 7 ■ If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

```
$ srvctl modify database -d db_name -y AUTOMATIC
```

8 Complete other post-upgrade steps.

For instructions, see the chapter *Performing post-upgrade tasks* in this document.

9 For upgrade scenarios that involve Oracle RAC 9i, start `gsd` as the Oracle user:

```
$ $ORACLE_HOME/bin/psctl start
```

10 Upgrade Oracle RAC.

For instructions, see the chapter *Upgrading Oracle RAC* in this document.

11 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

## Upgrading SF Oracle RAC using the Veritas script-based installation program

Use one of the following Veritas script-based installation programs to upgrade SF Oracle RAC: `installer` or `installsf rac`

The installer performs the following tasks to upgrade SF Oracle RAC:

- Verifies the compatibility of the systems before the upgrade.
- Stops the SF Oracle RAC processes before the upgrade.
- Uninstalls SF Oracle RAC.
- Installs the SF Oracle RAC 5.1 SP1 packages on the nodes.
- Starts SF Oracle RAC 5.1 SP1 on all the nodes.
- Displays the location of the log files, summary file, and response file.

---

**Note:** If you upgrade from version 5.0x, you need not manually restart the nodes in the cluster. The SF Oracle RAC processes are started automatically after the upgrade completes successfully.

If you upgrade to version 5.1 SP1 from version 5.1 or related 5.1 rolling patches, Symantec recommends that you do not roll back to version 5.1 or corresponding rolling patches.

---

## To upgrade to SF Oracle RAC 5.1 SP1 using the installsfrac program

### 1 Start the installation program using one of the following ways:

SF Oracle RAC installer      Navigate to the directory that contains the installation program.  
The program is located in the product directory.

Run the program:

```
# ./installsfrac galaxy nebula
```

Common product installer      Navigate to the directory that contains the installation program.  
Run the program:

```
# ./installer galaxy nebula
```

From the opening Selection Menu, choose **G** for "**Upgrade a Product.**"

From the displayed list of products, choose **Veritas Storage Foundation for Oracle RAC.**

The installer displays the copyright message and specifies the directory where the running logs are created.

### 2 Enter **3** to install all the SF Oracle RAC packages.

Before you select an option, you can view the list of packages that are installed for each option by choosing option 4.

```
1) Install minimal required Veritas Storage Foundation for Oracle RAC
   packages
2) Install recommended Veritas Storage Foundation for Oracle RAC
   packages
3) Install all Veritas Storage Foundation for Oracle RAC
   packages
4) Display packages to be installed for each option
Select the packages to be installed on all systems? [1-4,q,?] (2) 3
```

The installer verifies the systems for compatibility.

Review the messages displayed and make sure that you meet the requirements before proceeding with the upgrade.

### 3 Press **Enter** to continue with the upgrade.

Depending on the installation option selected, the installer displays the list of packages that will be installed and prompts for the stoppage of SF Oracle RAC processes before the upgrade.

- 4 Enter **y** to stop the SF Oracle RAC processes.

```
Do you want to stop SF Oracle RAC processes now? [y,n,q,?] (y)
```

The installer stops the processes and uninstalls SF Oracle RAC. After the uninstallation, the installer installs SF Oracle RAC 5.1 SP1 and starts SF Oracle RAC 5.1 SP1 on all the nodes.

- 5 Install the language packages and patches if you would like to run SF Oracle RAC in a language other than English.

See [“Installing language packages”](#) on page 207.

- 6 If you are upgrading from versions earlier than 5.0, restart each system in the cluster.

---

**Note:** You need not restart the systems manually if you are upgrading from version 5.0 and later. However, if any of the processes fail to start after the upgrade, you need to restart the systems.

---

```
# shutdown -g0 -y -i6
```

- 7 Complete the remaining tasks to finish the upgrade:

See [“Upgrading to SF Oracle RAC 5.1 SP1”](#) on page 236.

## Upgrading SF Oracle RAC using a response file

You can upgrade from SF Oracle RAC version 5.0 and later using a response file. Perform the steps in the following procedure to upgrade to SF Oracle RAC 5.1 SP1 using a response file.



To upgrade SF Oracle RAC using a response file

- 1 Upgrade the operating system, if required.  
For instructions, see the operating system documentation.
- 2 Create a response file using one of the available options.

---

**Note:** Make sure that you replace the host names in the response file with the names of the systems that you plan to upgrade.

---

For information on various options available for creating a response file:

See [“About response files”](#) on page 425.

For response file variable definitions:

See [“Response file variables to upgrade SF Oracle RAC”](#) on page 241.

For a sample response file:

See [“Sample response file for upgrading SF Oracle RAC”](#) on page 243.

- 3 Navigate to the directory containing the SF Oracle RAC installation program.
- 4 Start the installation:

```
# ./installsf rac -responsefile /tmp/response_file
```

Where /tmp/response\_file is the full path name of the response file.

- 5 Complete the post-upgrade steps.  
See the chapter "Performing post-upgrade tasks" in this document.

Response file variables to upgrade SF Oracle RAC

[Table 12-1](#) lists the response file variables that you can define to upgrade SF Oracle RAC.

**Table 12-1** Response file variables specific to upgrading SF Oracle RAC

| Variable          | List or Scalar | Description                                                           |
|-------------------|----------------|-----------------------------------------------------------------------|
| CFG{opt}{upgrade} | Scalar         | Upgrades SF Oracle RAC packages.<br>(Required)                        |
| CFG{accepteula}   | Scalar         | Specifies whether you agree with EULA.pdf on the media.<br>(Required) |

**Table 12-1** Response file variables specific to upgrading SF Oracle RAC  
*(continued)*

| Variable            | List or Scalar | Description                                                                                                                                                                             |
|---------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{systems}   | List           | List of systems on which the product is to be upgraded.<br>(Required)                                                                                                                   |
| CFG{prod}           | Scalar         | Defines the product to be upgraded.<br>(Required)                                                                                                                                       |
| CFG{vcs_allowcomms} | Scalar         | Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start).<br>(Required)                                      |
| CFG{opt}{keyfile}   | Scalar         | Defines the location of an ssh keyfile that is used to communicate with all remote systems.<br>(Optional)                                                                               |
| CFG{opt}{patchpath} | Scalar         | Defines a location, typically an NFS mount, from which all remote systems can install product patches. The location must be accessible from all target systems.<br>(Optional)           |
| CFG{opt}{pkgpath}   | Scalar         | Defines a location, typically an NFS mount, from which all remote systems can install product packages. The location must be accessible from all target systems.<br>(Optional)          |
| CFG{opt}{tmppath}   | Scalar         | Defines the location where a working directory is created to store temporary files and the packages that are needed during the install. The default location is /var/tmp.<br>(Optional) |

**Table 12-1** Response file variables specific to upgrading SF Oracle RAC  
(continued)

| Variable          | List or Scalar | Description                                                                                                                                                                                                                                            |
|-------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{logpath} | Scalar         | Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.<br><b>Note:</b> The installer copies the response files and summary files also to the specified <i>logpath</i> location.<br><br>(Optional) |
| CFG{opt}{rsh}     | Scalar         | Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems.<br><br>(Optional)                                                                                                                              |

## Sample response file for upgrading SF Oracle RAC

The following sample response file upgrades SF Oracle RAC to version 5.1 Service Pack 1 on nodes, galaxy and nebula.

```
our %CFG;  
  
$CFG{accepteula}=1;  
$CFG{opt}{installallpkgs}=1;  
$CFG{opt}{patchupgrade}=1;  
$CFG{prod}="SFRAC51";  
$CFG{systems}=[ qw(galaxy nebula) ];  
$CFG{vcs_allowcomms}=1;  
  
1;
```



# Performing a phased upgrade to SF Oracle RAC 5.1 SP1

This chapter includes the following topics:

- [About phased upgrade](#)
- [Performing phased upgrade of SF Oracle RAC from version 5.0](#)

## About phased upgrade

The phased upgrade methodology involves upgrading half of the nodes in the cluster at a time.

---

**Note:** You can perform a phased upgrade only if the Oracle RAC binaries are present on the local file system.

---

---

**Caution:** There is a potential for dependency problems between product components that no longer match when upgrading part of a cluster at a time. Follow the phased upgrade procedures carefully to avoid these problems.

---

---

**Note:** There will be some downtime involved. Review the procedures and carefully plan your downtime before proceeding with any steps. The sample procedures assume that Oracle RAC binaries are installed on local file systems for each node in the cluster.

---

The examples in the procedures assume a four-node SF Oracle RAC cluster with the nodes *galaxy* and *nebula* constituting the first half of the cluster and the nodes *jupiter* and *mercury* constituting the second half of the cluster.

# Performing phased upgrade of SF Oracle RAC from version 5.0

**Table 13-1** illustrates the phased upgrade process. Each column describes the steps to be performed on the corresponding subcluster and the status of the subcluster when operations are performed on the other subcluster.

**Table 13-1** Summary of phased upgrade



| First half of the cluster                                                                                                                                                                                                                                                                                                                                                                                                                       | Second half of the cluster                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC cluster before the upgrade:                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                     |
| <div><div>First half of the cluster</div><div></div><div>Second half of the cluster</div><div></div></div>                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                     |
| <p><b>STEP 1:</b> Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"><li>■ Switch failover applications.</li><li>■ Stop all parallel applications.</li></ul> <p>See “<a href="#">Step 1: Performing pre-upgrade tasks on the first half of the cluster</a>” on page 247.</p> <p><b>STEP 2:</b> Upgrade SF Oracle RAC.</p> <p>See “<a href="#">Step 2: Upgrading the first half of the cluster</a>” on page 250.</p> | <p>The second half of the cluster is up and running.</p> <div></div>                                                                                                                                                                                            |
| <p>The first half of the cluster is not running.</p> <div></div>                                                                                                                                                                                                                                                                                             | <p><b>STEP 3:</b> Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"><li>■ Stop all parallel and failover applications.</li><li>■ Stop SF Oracle RAC.</li></ul> <p>See “<a href="#">Step 3: Performing pre-upgrade tasks on the second half of the cluster</a>” on page 251.</p> <p><b>The downtime starts now.</b></p> |

Table 13-1 Summary of phased upgrade (continued)

| First half of the cluster                                                                                                                                                                                                                                                                                                       | Second half of the cluster                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>STEP 4:</b> Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"><li>■ Start SF Oracle RAC.</li><li>■ Start all applications.</li></ul> <p>See <a href="#">“Step 4: Performing post-upgrade tasks on the first half of the cluster”</a> on page 252.</p> <p><b>The downtime ends here.</b></p> | <p>The second half of the cluster is not running.</p> <div><p>jupiter                      mercury</p></div>                                                                                                                                                                                                                                     |
| <p>The first half of the cluster is up and running.</p> <div><p>galaxy                      nebula</p></div>                                                                                                                                   | <p><b>STEP 5:</b> Upgrade SF Oracle RAC.</p> <p>See <a href="#">“Step 5: Upgrading the second half of the cluster”</a> on page 254.</p> <p><b>STEP 6:</b> Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"><li>■ Start SF Oracle RAC.</li><li>■ Start all applications.</li></ul> <p>See <a href="#">“Step 6: Performing post-upgrade tasks on the second half of the cluster”</a> on page 254.</p> |
| <p>The phased upgrade is complete and both the first and the second half of the cluster are running.</p> <div><p>galaxy                      nebula                      jupiter                      mercury</p></div>                      |                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Step 1: Performing pre-upgrade tasks on the first half of the cluster

Perform the following pre-upgrade steps on the first half of the cluster.

### To perform the pre-upgrade tasks on the first half of the cluster

- 1 Back up the following configuration files: `main.cf`, `types.cf`, `CVMTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```
# 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
# cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
/etc/VRTSvcs/conf/config/OracleTypes.cf.save
# cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
/var/tmp/PrivNIC.cf.save
# cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
/var/tmp/MultiPrivNIC.cf.save
```

- 2 Stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 3 If you plan to continue using checkpoints or storage tiering policies you created with a 5.0x or earlier version of Storage Foundation for Oracle RAC, complete the following preparatory step before migrating the SFDB repository database to 5.1 Service Pack 1.

See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 233.

- 4 Stop the Oracle database:
  - If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the first half of the cluster and shut down the instances:

```
$ srvctl stop instance -d database_name \
-i instance_name
```

- If the Oracle RAC instance is managed by VCS:

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



- 5 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw
# hagrpl -modify oracle_group AutoStart 0
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

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

- 6 Unmount the non-system VxFS and CFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system or shared volumes. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs
# fuser -cu /mount_point
```

- Unmount the non-system VxFS or CFS file system:

```
# umount /mount_point
```

- 7 Stop the parallel service groups and switch over failover service groups:

```
# hastop -local -evacuate
```

- 8 Verify that only ports a, b, d, and o are open:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 79c302 membership 0123
Port b gen 79c307 membership 0123
Port d gen 79c306 membership 0123
Port o gen 79c304 membership 0123
```

## Step 2: Upgrading the first half of the cluster

Perform the following steps to upgrade the first half of the cluster.

### To upgrade the first half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
$ mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.  
For instructions, see the operating system documentation.

- 3 Rename the `/etc/llttab` file to its original name:

```
$ mv /etc/llttab.save /etc/llttab
```

- 4 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.

- 5 Upgrade SF Oracle RAC:

```
# cd /dvd_mount/storage_foundation_for_oracle_rac
```

```
# ./installsfrac galaxy nebula
```

---

**Note:** After you complete the upgrade of the first half of the cluster, no GAB ports will be shown in the output when you run the `gabconfig -a` command.

---

- 6 Change the `/etc/default/llt` file to prevent LLT from starting automatically after reboot by setting the `LLT_START` attribute to 0:

```
LLT_START=0
```

- 7 Restart the nodes:

```
# shutdown -g0 -y -i6
```

You may see some errors in the system log file when the nodes restart. This is because LLT is disabled. Ignore these messages.

```
svc.startd[7]: [ID 652011 daemon.warning] svc:/system/llt:default:
Method "/lib/svc/method/llt start" failed with exit status 2.
gab: [ID 438192 kern.notice] GAB WARNING V-15-1-20115
Port d registration failed, GAB not configured
```

## Step 3: Performing pre-upgrade tasks on the second half of the cluster

Perform the following pre-upgrade steps on the second half of the cluster.

### To perform the pre-upgrade tasks on the second half of the cluster

- 1 Stop all applications that are not configured under VCS. Use native application commands to stop the application.
- 2 If you plan to continue using checkpoints or storage tiering policies you created with a 5.0x or earlier version of Storage Foundation for Oracle RAC, complete the following preparatory step before migrating the SFDB repository database to 5.1 Service Pack 1.

See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 233.

- 3 Stop the Oracle database:

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the second half of the cluster and shut down the instances:

```
$ srvctl stop instance -d database_name \
-i instance_name
```

- If the Oracle RAC instance is managed by VCS:

```
# hagrps -offline oracle_group -sys jupiter
# hagrps -offline oracle_group -sys mercury
```

---

**Note:** The downtime starts now.

---

- 4 Unmount the non-system VxFS or CFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system or shared volumes. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs

# fuser -cu /mount_point
```

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

- 5 Stop VCS:

```
# hastop -local
```

- 6 Stop the following SF Oracle RAC modules: Fencing, ODM, VCSMM, VCS, GAB, LMX, and LLT

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

# /etc/init.d/lmx stop

# /etc/init.d/vcsmm stop

# /etc/init.d/vxfen stop

# /etc/init.d/gab stop

# modinfo |grep lmx
# modunload -i module_no

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

## Step 4: Performing post-upgrade tasks on the first half of the cluster

Perform the following post-upgrade steps on the first half of the cluster.

### To perform the post-upgrade tasks on the first half of the cluster

- 1 Change `/etc/default/llt` to start LLT on the nodes by setting the `LLT_START` attribute to 1:

```
LLT_START=1
```

- 2 On the first half of the cluster, force GAB to form a cluster.

```
# gabconfig -x
```

- 3 On the first half of the cluster, start SF Oracle RAC:

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

```
# ./installsfrc -start galaxy nebula
```

Verify that the GAB ports a, b, d, o and h appear in the `gabconfig -a` command output.

- 4 Upgrade VxVM disk group version.

See [“Upgrading VxVM disk group version”](#) on page 305.

- 5 Upgrade disk layout version.

See [“Upgrading disk layout versions”](#) on page 305.

- 6 Relink the SF Oracle RAC libraries with Oracle.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 301.

- 7 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys node_name
```

If the Oracle database is not managed by VCS:

```
$ srvctl start instance -d db_name \
-i instance_name
```

---

**Note:** The downtime ends here.

---

- 8 On the first half of the cluster, manually mount the VxFS or CFS file systems that are not managed by VCS.
- 9 On the first half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

## Step 5: Upgrading the second half of the cluster

Perform the following steps to upgrade the second half of the cluster.

### To upgrade the second half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
$ mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.  
For instructions, see the operating system documentation.

- 3 Rename the `/etc/llttab` file to its original name:

```
$ mv /etc/llttab.save /etc/llttab
```

- 4 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.

- 5 On the second half of the cluster, upgrade SF Oracle RAC:

```
# cd /dvd_mount/storage_foundation_for_oracle_rac
```

```
# ./installsffrac jupiter mercury
```

- 6 Restart the nodes:

When the nodes in the second half of the cluster come up, all the GAB ports a, b, d, o, h, u, v, w and f will be online. All the CFS mount service groups also come online automatically.

```
# shutdown -g0 -y -i6
```

## Step 6: Performing post-upgrade tasks on the second half of the cluster

Perform the following post-upgrade steps on the second half of the cluster.

## To perform the post-upgrade tasks on the second half of the cluster

- 1 Relink the SF Oracle RAC libraries with Oracle.  
See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 301.
- 2 Bring the Oracle database service group online.  
If the Oracle database is managed by VCS:  

```
# hagrpl -online oracle_group -sys node_name
```

  
If the Oracle database is not managed by VCS:  

```
$ srvctl start instance -d db_name \
-i instance_name
```
- 3 ■ If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:  

```
# haconf -makerw
# hagrpl -modify oracle_group AutoStart 1
# haconf -dump -makero
```

  
■ If the Oracle database is not managed by VCS, change the management policy for the database to automatic:  

```
$ srvctl modify database -d db-name -y AUTOMATIC
```
- 4 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 5 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 6 Set or change the product license level, if required.  
See [“Setting or changing the product license level”](#) on page 304.
- 7 Migrate the SFDB repository database.  
See [“Post upgrade tasks for migrating the SFDB repository database”](#) on page 305.

**8** Verify the cluster.

See [“Verifying the cluster”](#) on page 323.

**9** If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.



# Performing a rolling upgrade to SF Oracle RAC 5.1 SP1

This chapter includes the following topics:

- [About rolling upgrades](#)
- [Prerequisites for a rolling upgrade](#)
- [Preparing to perform a rolling upgrade to SF Oracle RAC 5.1 SP1](#)
- [Performing a rolling upgrade using the installer](#)

## About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime. Rolling upgrades are not compatible with phased upgrades. Do not perform "mixed" rolling upgrades with phased upgrades.

Rolling upgrades take two discrete phases. In the first, you upgrade the product kernel . In the second, you upgrade the non-kernel such as VCS and agent .

You can perform a rolling upgrade from 5.1, 5.1 P1, 5.1 RP1, or 5.1 RP2 to 5.1 SP1.

## Prerequisites for a rolling upgrade

Meet the following prerequisites before performing a rolling upgrade:

- Make sure the product that you want to upgrade supports rolling upgrades.

- Split your clusters into sub-clusters for the upgrade to keep the service groups available during upgrade.
- Make sure you logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.

## Preparing to perform a rolling upgrade to SF Oracle RAC 5.1 SP1

Perform the preparatory steps in this section if you are performing a rolling upgrade of the cluster. Before you upgrade, make sure that your systems meet the hardware and software requirements for this release.

---

**Note:** Perform the steps on the first subcluster.

---

### To prepare to upgrade SF Oracle RAC

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```
# 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
# cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
/etc/VRTSvcs/conf/config/OracleTypes.cf.save
# cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
/var/tmp/PrivNIC.cf.save
# cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
/var/tmp/MultiPrivNIC.cf.save
```

- 3 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

If the applications are under VCS control:

```
# hagrps -offline grp_name -sys node_name
```

If the applications are not under VCS control:

Use native application commands to stop the application.

- 4 For Oracle RAC 10g and Oracle RAC 11g:

Stop the Oracle RAC resources on each node.

- If the database instances are not managed by VCS, then run the following on one node:

```
$ srvctl stop instance -d db_name -i instance_name
```

- If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

```
# hagrps -offline grp_name -sys node_name
```

- 5 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 0
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

```
$ srvctl modify database -d db_name -y manual
```

- 6 Take all the VCS service groups offline:

```
# hagrps -offline grp_name -sys sys_name
```

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

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

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

- 8 If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to SF Oracle RAC 5.1 SP1.

See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 233.

## Performing a rolling upgrade using the installer

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

### Performing a rolling upgrade on kernel : phase 1

Note that in the following instructions that a sub-cluster can represent one or more nodes in a full cluster, but is represented by nodeA.

#### To perform the rolling upgrade on kernel : phase 1

- 1 On the first sub-cluster, start the installer for the rolling upgrade with the `-upgrade_kernelpkgs` option.

```
./installer -upgrade_kernelpkgs nodeA
```

- 2 Note that if the boot disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
- 3 The installer checks system communications, package versions, product versions, and completes prechecks.

It then upgrades applicable product kernel .

- 4 The installer loads new kernel modules.

- 5 The installer starts all the relevant processes and brings all the service groups online.

In case the startup of some of the processes fails, you may need to reboot the nodes and manually check the cluster's status.

---

**Note:** The Oracle service group at this point will be offline as the AutoStart attribute is set to 0 to prevent the service group from starting automatically. The service group is started later in the process.

---

- 6 Relink the SF Oracle RAC libraries with Oracle by choosing the option **Relink Oracle Database Binary** from the program menu.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 301.

- 7 Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys node_name
```

- If the Oracle database is not managed by VCS:

```
# srvctl start database -d db_name
```

- 8 Manually mount the VxFS and CFS file systems that are not managed by VCS.

- 9 Start all applications that are not managed by VCS. Use native application commands to start the applications.

- 10 If the boot disk is encapsulated, reboot the first sub-cluster's system.

- 11 Before you proceed to phase 2, complete step 1 to 10 on the second subcluster.

- 12 ■ If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

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

**13** Migrate the SFDB repository database.

See [“Post upgrade tasks for migrating the SFDB repository database”](#) on page 305.

## Performing a rolling upgrade on non-kernel : phase 2

In this phase installer installs all non-kernel packages on all the nodes in cluster and restarts VCS cluster.

### To perform the rolling upgrade on non-kernel : phase 2

- 1** Start the installer for the rolling upgrade with the `-upgrade_nonkernelpkgs` option. Specify all the nodes in the cluster:

```
./installer -upgrade_nonkernelpkgs nodeA nodeB nodeC...
```

- 2** The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1.
- 3** The installer upgrades non-kernel packages.
- 4** The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1. The installer loads the new kernel modules. It then starts all relevant processes and brings all the service groups online.
- 5** Verify the cluster's status:

```
# hastatus -sum
```

- 6** If you want to upgrade VCS or SFHA 5.1 on the CP server systems to version 5.1 SP1 PR1, make sure you upgraded all application clusters to 5.1 SP1 PR1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

# Upgrading to SF Oracle RAC 5.1 SP1 using Live Upgrade

This chapter includes the following topics:

- [About Live Upgrade](#)
- [Supported upgrade paths for Live Upgrade](#)
- [Before you upgrade SF Oracle RAC using Solaris Live Upgrade](#)
- [Upgrading the operating system and SF Oracle RAC using Live Upgrade](#)
- [Upgrading SF Oracle RAC only using Live Upgrade](#)
- [Upgrading Solaris only using Live Upgrade](#)
- [Creating a new boot environment on the alternate boot disk](#)
- [Upgrading SF Oracle RAC manually](#)
- [Upgrading SF Oracle RAC using the installer for a Live Upgrade](#)
- [Completing the Live Upgrade](#)
- [Verifying Live Upgrade of SF Oracle RAC](#)
- [Reverting to the primary boot environment](#)

## About Live Upgrade


You can use Live Upgrade to perform the following types of upgrade:

- Upgrade the operating system and SF Oracle RAC.
- Upgrade the operating system.

- Upgrade SF Oracle RAC.

Figure 15-1 illustrates an example of an upgrade of Veritas products from 5.0MP3 to 5.1 Service Pack 1, and the operating system from Solaris 9 to Solaris 10.

Figure 15-1 Live Upgrade process



## About Live Upgrade in a Veritas Volume Replicator (VVR) environment

In a SF Oracle RAC environment that uses Veritas Volume Replicator, the following scripts provide the means to upgrade the VVR configuration:

- `vvr_upgrade_lu_start`
- `vvr_upgrade_lu_finish`

This section provides an overview of the VVR upgrade process. See the Live Upgrade procedures for SF Oracle RAC for the complete procedure.

- Use the `vxlustart` script to perform upgrade steps for SF Oracle RAC.
- Immediately before rebooting the system to switch over to the alternate boot environment, run the `vvr_upgrade_lu_start` script.

**Note:** Use the `vvr_upgrade_lu_start` script only when the applications are stopped and the next step is to switch over to the alternate boot environment.



- After the `vvr_upgrade_lu_start` script completes successfully, reboot the system. This reboot results in the system booting from the alternate boot environment.
- After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.

## Supported upgrade paths for Live Upgrade

The systems where you plan to use Live Upgrade must run Solaris 9 or Solaris 10. For Solaris 10, make sure that all non-global zones are booted and in the running state before you use the Veritas product installer to upgrade the Storage Foundation products in the global zone. If the non-global zones are not mounted and running at the time of the upgrade, you must upgrade each package in each non-global zone manually.

For Live Upgrade, if the alternative root environment also has a zone, you cannot install `VRTSodm`. You must remove the `VRTSodm` package first then install the Storage Foundation product. After you reboot the alternative root, you can install `VRTSodm`.

SF Oracle RAC version must be at least 4.1.

---

**Note:** If you use Live Upgrade on a system where non-global zones are configured, make sure that all the zones are in the `installed` state before you start Live Upgrade.

---

You can use Live Upgrade in the following virtualized environments:

**Table 15-1** Live Upgrade support in virtualized environments

| Environment                | Procedure                                                                                                                                                                                   |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oracle VM Server for SPARC | Perform Live Upgrade on the Domain controller only.<br><br>Perform Live Upgrade on the Guest domain only.<br><br>Use the standard Live Upgrade procedure for both types of logical domains. |

# Before you upgrade SF Oracle RAC using Solaris Live Upgrade

Before you upgrade, perform the following procedure.

## To prepare for the Live Upgrade

- 1 Make sure that the SF Oracle RAC installation media and the operating system installation images are available and on hand.
- 2 On the nodes to be upgraded, select an alternate boot disk that is at least the same size as the root partition of the primary boot disk.  
  
If the primary boot disk is mirrored, you need to break off the mirror for the alternate boot disk.
- 3 Before you perform the Live Upgrade, take offline any services that involve non-root file systems. This prevents file systems from being copied to the alternate boot environment that could potentially cause a root file system to run out of space.
- 4 On the primary boot disk, patch the operating system for Live Upgrade.

---

**Note:** Live Upgrade may fail on a Solaris 9 host if a VxFS file system is in `/etc/vfstab`. On the Solaris 9 host, install the Live Upgrade packages `SUNWlucfg`, `SUNWluu`, and `SUNWlur` from a Solaris 10 image. After you install the packages, install the latest Live Upgrade patch.

---

For more information on required packages and patches, visit the following site and search on "Live Upgrade requirements."

<http://wikis.sun.com>

- 5 The version of the Live Upgrade packages must match the version of the operating system to which you want to upgrade on the alternate boot disk. If you are upgrading the Solaris operating system, do the following steps:
  - Remove the installed Live Upgrade packages for the current operating system version:  
All Solaris versions: `SUNWluu`, `SUNWlur` packages.  
Solaris 10 update 5 or later also requires: `SUNWlucfg` package.
  - From the new Solaris installation image, install the new versions of the Live Upgrade packages:  
All Solaris versions: `SUNWluu`, `SUNWlur` packages.

Solaris 10 update 5 or later also requires: SUNWlucfg package.

- 6 If the specified image is missing patches that are installed on the primary boot disk, note the patch numbers. To ensure that the alternate boot disk is the same as the primary boot disk, you need to install any missing patches on the alternate boot disk.

In the procedure examples, the primary or current boot environment resides on Disk0 (c0t0d0) and the alternate or inactive boot environment resides on Disk1 (c0t1d0).

## Upgrading the operating system and SF Oracle RAC using Live Upgrade

Perform the following steps to upgrade both the operating system and SF Oracle RAC using Live Upgrade.

### To upgrade the operating system and SF Oracle RAC using Live Upgrade

- 1 Prepare to upgrade using Solaris Live Upgrade.  
See [“Before you upgrade SF Oracle RAC using Solaris Live Upgrade”](#) on page 266.
- 2 Create a new boot environment on the alternate boot disk.  
See [“Creating a new boot environment on the alternate boot disk”](#) on page 270.
- 3 Upgrade SF Oracle RAC using the installer or manually.  
See [“Upgrading SF Oracle RAC using the installer for a Live Upgrade”](#) on page 277.  
See [“Upgrading SF Oracle RAC manually”](#) on page 272.
- 4 Complete the Live Upgrade.  
See [“Completing the Live Upgrade”](#) on page 278.
- 5 Verify Live Upgrade of SF Oracle RAC.  
See [“Verifying Live Upgrade of SF Oracle RAC”](#) on page 281.
- 6 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

## Upgrading SF Oracle RAC only using Live Upgrade

Perform the following steps to upgrade only SF Oracle RAC using Live Upgrade.

### To upgrade only SF Oracle RAC using Live Upgrade

- 1 Prepare to upgrade using Solaris Live Upgrade.  
See [“Before you upgrade SF Oracle RAC using Solaris Live Upgrade”](#) on page 266.
- 2 Create a new boot environment on the alternate boot disk.  
See [“Creating a new boot environment on the alternate boot disk”](#) on page 270.
- 3 Upgrade SF Oracle RAC using the installer or manually.  
See [“Upgrading SF Oracle RAC using the installer for a Live Upgrade”](#) on page 277.  
See [“Upgrading SF Oracle RAC manually”](#) on page 272.
- 4 Complete the Live Upgrade.  
See [“Completing the Live Upgrade”](#) on page 278.
- 5 Verify Live Upgrade of SF Oracle RAC.  
See [“Verifying Live Upgrade of SF Oracle RAC”](#) on page 281.
- 6 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

## Upgrading Solaris only using Live Upgrade

If you are upgrading Solaris only, you must remove and reinstall SF Oracle RAC from the alternate boot environment before completing the Live Upgrade. SF Oracle RAC has internal components that are OS-specific and therefore needs to be re-installed after the operating system upgrade.

Perform the following steps to upgrade only Solaris using Live Upgrade.

### To upgrade only Solaris using Live Upgrade

- 1 Prepare to upgrade using Solaris Live Upgrade.  
 See [“Before you upgrade SF Oracle RAC using Solaris Live Upgrade”](#) on page 266.
- 2 Create a new boot environment on the alternate boot disk.  
 See [“Creating a new boot environment on the alternate boot disk”](#) on page 270.
- 3 Remove and reinstall SF Oracle RAC 5.1 SP1 on the alternate boot environment.  
 See [“Removing and reinstalling SF Oracle RAC using the installer”](#) on page 269.
- 4 Complete the Live Upgrade.
- 5 Verify Live Upgrade of SF Oracle RAC.  
 See [“Verifying Live Upgrade of SF Oracle RAC”](#) on page 281.
- 6 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.  
 For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

## Removing and reinstalling SF Oracle RAC using the installer

SF Oracle RAC has kernel components that are specific for Solaris operating system versions. When you use Solaris Live Upgrade to upgrade the Solaris operating system, you must complete these steps to ensure the correct version of SF Oracle RAC components are installed.

On a node in the cluster, run the installer on the alternate boot disk to remove and reinstall SF Oracle RAC 5.1 SP1 on all the nodes in the cluster.

At the end of the process the following occurs:

- SF Oracle RAC 5.1 SP1 is installed on the alternate boot disk, with the correct binaries for the new operating system version

### To remove and reinstall SF Oracle RAC using the installer

- 1 Insert the product disc with SF Oracle RAC 5.1 SP1 or access your copy of the software on the network.
- 2 Complete the pre-uninstallation tasks for the installed version of SF Oracle RAC.

For instructions, see the "Preparing to uninstall SF Oracle RAC" chapter in this document.

- 3 Uninstall SF Oracle RAC, specifying the alternate boot disk as the root path:

```
# /dvd_mount/uninstallsfrac -rootpath /altroot.5.10
```

- 4 Enter the names of the nodes that you want to uninstall.

The installer displays the list of packages that will be uninstalled.

- 5 Press **Return** to continue.

- 6 Install using the installer script, specifying the root path as the alternate boot disk:

```
# cd /dvd_mount
```

```
# ./installsfrac -install -rootpath /altroot.5.10
```

The installer displays the list of packages that will be installed.

- 7 Press **Return** to continue with the installation.

- 8 When the installer prompts you to configure the SF Oracle RAC, select **n**.

---

**Note:** To maintain the existing product configurations, you must select No. If you select yes, the existing product configurations are overwritten.

---

- 9 Verify that the version of the Veritas packages on the alternate boot disk is 5.1 Service Pack 1.

```
# pkginfo -R /altroot.5.10 -l VRTSpkgname
```

For example:

```
# pkginfo -R /altroot.5.10 -l VRTSdbac
```

Review the installation logs at `/altroot.5.10/opt/VRTS/install/log`.

## Creating a new boot environment on the alternate boot disk

Run the `vxlustart` command on each node in the cluster to create a new boot environment on the alternate boot disk.

---

**Note:** This step can take several hours to complete. Do not interrupt the session as it may leave the boot environment unstable.

---

At the end of the process:

- The Solaris operating system on the alternate boot disk is upgraded, if you have chosen to upgrade the operating system.
- A new boot environment is created on the alternate boot disk by cloning the primary boot environment.

#### To create a new boot environment on the alternate boot disk

- 1 Navigate to the install media for the Symantec products:

```
# cd /dvd_mount/scripts
```

- 2 On each node, run one of the following commands:

To upgrade the operating system, by itself or together with upgrading the Veritas products:

```
# ./vxlustart -v -u targetos_version \  
-s osimage_path -d disk_name
```

To upgrade the Veritas product only:

```
# ./vxlustart -v -u 5.10 -U -d disk_name
```

For example, to upgrade to Solaris 10 update 6:

```
# ./vxlustart -v -u 5.10 -s /mnt/Solaris_10u6
```

- 3 Review the output and note the new mount points. If the system is rebooted before completion of the upgrade or if the mounts become unmounted, you may need to remount the disks.

If you need to remount, run the command:

```
# vxlustart -r -u targetos_version -d disk_name
```

- 4 After the alternate boot disk is created, install any operating system patches that are required for the Veritas product installation.

# Upgrading SF Oracle RAC manually

This section provides instructions for performing a manual upgrade of SF Oracle RAC using Live Upgrade. On each node, remove and install the SF Oracle RAC packages as described in the procedure.

At the end of the process:

- SF Oracle RAC 5.1 SP1 is installed on the alternate boot disk.
- The boot environment on the alternate disk is activated.

Depending on the SF Oracle RAC version you are upgrading from, perform the steps in one of the following procedures:

|                   |                                                                                                        |
|-------------------|--------------------------------------------------------------------------------------------------------|
| From version 5.1x | See <a href="#">“To perform Live Upgrade of SF Oracle RAC manually from version 5.1x”</a> on page 272. |
| From version 5.0x | See <a href="#">“To perform Live Upgrade of SF Oracle RAC manually from version 5.0x”</a> on page 274. |

## To perform Live Upgrade of SF Oracle RAC manually from version 5.1x

- 1 Remove the SF Oracle RAC packages on the alternate boot disk in the listed order.

The `-R` option removes the packages from the root path `/altroot.5.10` on the alternate boot disk.

```
# pkgrm -R /altroot.5.10 \
VRTScutil VRTSspt VRTSvlic
```

- 2 Install the following 5.1 Service Pack 1 packages one at a time to the alternate boot disk using the `pkgadd` command.

```
# pkgadd -R /altroot.5.10 VRTSvlic.pkg
# pkgadd -R /altroot.5.10 VRTSspt.pkg
# pkgadd -R /altroot.5.10 VRTSamf.pkg
```

- 3 Perform an in-place upgrade of the following packages:

- Copy the admin file from the directory `/var/sadm/install/admin/default` to the directory `/var/tmp/admin_file`:

```
# cp -p /var/sadm/install/admin/default \
/var/tmp/admin_file
```



- Edit the file `/var/tmp/admin_file` to set the value of the `instance` attribute to `overwrite`. This ensures that the existing version of the packages is replaced by the new version.

```
instance=overwrite
```

- Install the following packages:

```
# pkgadd -a /var/tmp/admin_file -R /altroot.5.10 VRTSaslapm.pkg
# pkgadd -a /var/tmp/admin_file -R /altroot.5.10 VRTSsfmh.pkg
```

- 4 In the `/media` directory, list the patches for each platform. Enter the following:

```
# ./installer -listpatches
```

- 5 Unzip and extract the tar file for each patch.

```
# gunzip patch_name.tar.gz
# tar vxf patch_name.tar
```

where *patch\_name* is the name of a patch listed in step 4.

For example:

```
# gunzip 1427629-06.tar.gz
# tar vxf 143282-01.tar
```

Repeat this step for each patch.

- 6 Install the patches on the alternative boot disk using the `patchadd` command.

```
# patchadd -R /altroot.5.10 patch_name
```

For example:

```
# patchadd -R /altroot.5.10 1427629-06
```

**7** Copy the following configuration files to the directory

`/altroot.5.10/etc/VRTSvcs/conf/config:`

```
# cp -p /altroot.5.10/etc/VRTSvcs/conf/types.cf \
/altroot.5.10/etc/VRTSvcs/conf/config
# cp -p /altroot.5.10/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf \
/altroot.5.10/etc/VRTSvcs/conf/config
# cp -p /altroot.5.10/etc/VRTSagents/ha/conf/Oracle/OracleTypes.cf \
/altroot.5.10/etc/VRTSvcs/conf/config
# cp -p /altroot.5.10/etc/VRTSagents/ha/conf\
/OracleASM/OracleASMTypes.cf \
/altroot.5.10/etc/VRTSvcs/conf/config
# cp -p /altroot.5.10/etc/VRTSagents/ha/conf/Sybase/SybaseTypes.cf \
/altroot.5.10/etc/VRTSvcs/conf/config
```

**8** Install the language packages if you would like to run SF Oracle RAC in a language other than English. Insert the language disc into the disc drive and run the following commands:

```
# cd /dvd_mount
# ./install_lp -rootpath /altroot.5.10 galaxy nebula
```

**9** Verify that the version of the SF Oracle RAC packages on the alternate boot disk is 5.1 Service Pack 1.

```
# pkginfo -R /altroot.5.10 -l VRTSdbac
```

You may also review the installation logs at `/opt/VRTS/install/log`.

**To perform Live Upgrade of SF Oracle RAC manually from version 5.0x**

**1** Remove the SF Oracle RAC packages on the alternate boot disk in the listed order.

The `-R` option removes the packages from the root path `/altroot.5.10` on the alternate boot disk.

```

For upgrades # pkgrm -R /altroot.5.10 \
from 5.0 and VRTScsocw VRTSdbac VRTSodm VRTSgms VRTScavf VRTSglm \
related VRTSgapms VRTSvail VRTSvxmsa VRTSdbed VRTSdbcom \
patches VRTSvcSor VRTScmccc VRTScmcs VRTSacclib VRTScscm \
VRTScscw VRTScssim VRTScutil VRTSvcsmn VRTSvcsmg \
VRTSvcSag VRTSvcS VRTSvxfen VRTSgab VRTSllt VRTSfsmnd \
VRTSfssdk VRTSfsman VRTSvrw VRTSweb VRTSvcsvr VRTSvrpro \
VRTSddlpr VRTSvdid VRTSvmpro VRTSalloC VRTSdcli VRTSvman \
VRTSspt VRTSaa VRTSmh VRTSccg VRTSobgui VRTSfspro VRTSdsa \
VRTSob VRTSobc33 VRTSat VRTSspbX VRTSicsco VRTSvxfs \
VRTSvxvm VRTSjre15 VRTSperl

```

- 2 Install SF Oracle RAC 5.1 Service Pack 1 packages and patches in the following order one at a time to the alternate boot disk using the `pkgadd` command.

---

**Note:** The patches are in compressed format in the `/media/patches` directory. Uncompress the patches and install them on the alternate boot disk.

---

#### Packages:

```

VRTSvlic.pkg VRTSperl.pkg VRTSspt.pkg VRTSvxvm.pkg VRTSaslapm.pkg
VRTSob.pkg VRTSsfmh.pkg VRTSvxfs.pkg VRTSfssdk.pkg VRTSat.pkg
VRTSllt.pkg VRTSgab.pkg VRTSvxfen.pkg VRTSamf.pkg VRTSvcS.pkg
VRTScps.pkg VRTSvcSag.pkg VRTSvcsea.pkg VRTSdbed.pkg VRTSglm.pkg
VRTScavf.pkg VRTSgms.pkg VRTSodm.pkg VRTSdbac.pkg

```

#### Patches (Solaris 9):

```

142633-05 143270-05 143273-05 143281-01
143283-01 143677-01 143680-01 145450-01
145454-01 142629-06 142631-04 143287-01
143288-01 143289-01 143290-01 143687-02

```

#### Patches (Solaris 10 SPARC):

```

142634-05 143271-05 143274-05 143282-01
143284-01 143678-01 143681-01 145451-01
145455-01 142629-06 142631-04 143287-01
143288-01 143289-01 143290-01 143687-02

```

#### Patches (Solaris 10 x86-64):

```

142632-04 143693-02 143291-01 145456-01
143296-01 143295-01 143294-01 143292-01

```

```
143297-01 145452-01 142630-06 143682-01
143679-01 143275-05 143272-05 142635-05
```

For example:

```
# pkgadd -R /altroot.5.10 -d package_name.pkg

# patchadd -R /altroot.5.10 -d patch_id
```

Where:

*package\_name.pkg* is the name of the package, for example *VRTSvxvm.pkg*.

*patch\_id* is the patch ID, for example 142629-04.

```
# pkgadd -R /altroot.5.10 -d VRTSvxvm.pkg

# patchadd -R /altroot.5.10 -d 142629-04
```

- 3 Install the language packages if you would like to run SF Oracle RAC in a language other than English. Insert the language disc into the disc drive and run the following commands:

```
# cd /dvd_mount
# ./install_lp -rootpath /altroot.5.10 galaxy nebula
```

- 4 Verify that the version of the SF Oracle RAC packages on the alternate boot disk is 5.1 Service Pack 1.

```
# pkginfo -R /altroot.5.10 -l VRTSdbac
```

You may also review the installation logs at `/opt/VRTS/install/log`.

- 5 On the alternate root path of any one node in the cluster, set the environment variable `INSTALL_ROOT_PATH` and configure a unique VCS cluster ID:

```
# export INSTALL_ROOT_PATH=/altroot.5.10
# /altroot.5.10/opt/VRTSvcs/bin/uuidconfig.pl \
-clus -configure -use_llthost
```

The `-use_llthost` option indicates that the `/etc/llthost` file is used to determine the names of the nodes in the cluster. Alternatively, you may specify the node names instead of the file name.

# Upgrading SF Oracle RAC using the installer for a Live Upgrade

You can use the Veritas product installer to upgrade SF Oracle RAC as part of the Live Upgrade.

On a node in the cluster, run the installer on the alternate boot disk to upgrade SF Oracle RAC on all the nodes in the cluster. The program uninstalls the existing version of SF Oracle RAC on the alternate boot disk during the process.

At the end of the process the following occurs:

- SF Oracle RAC 5.1 SP1 is installed on the alternate boot disk.

## To perform Live Upgrade of SF Oracle RAC using the installer

- 1 Insert the product disc with SF Oracle RAC 5.1 SP1 or access your copy of the software on the network.
- 2 Run the installer script specifying the root path as the alternate boot disk:

```
# ./installer -upgrade -rootpath /altroot.5.10
```

If you are upgrading from Solaris 9 to Solaris 10, and you are upgrading SF Oracle RAC 4.1MP2, 5.0, or 5.0MPx to 5.1SP1, upgrade SF Oracle RAC using the `installer -upgrade` command.

If you are upgrading from Solaris 9 to Solaris 10 and you are upgrading SF Oracle RAC from 5.1 or 5.1RPx to 5.1SP1, uninstall and reinstall SF Oracle RAC.

See [“Removing and reinstalling SF Oracle RAC using the installer”](#) on page 269.

- 3 Enter the names of the nodes that you want to upgrade to SF Oracle RAC 5.1 SP1.

---

**Note:** Make sure that the installed version of VxFS uses the disk layout version 6 or later. If you are on a previous disk layout version, upgrade the version before you proceed with the SF Oracle RAC installation.

---

The installer displays the list of packages to be installed or upgraded on the nodes.

- 4 Press **Return** to continue with the installation.
- 5 Verify that the version of the Veritas packages on the alternate boot disk is 5.1 Service Pack 1.

```
# pkginfo -R /altroot.5.10 -l VRTSpkgname
```

For example:

```
# pkginfo -R /altroot.5.10 -l VRTSdbac
```

Review the installation logs at `/altroot.5.10/opt/VRTS/install/log`.

## Completing the Live Upgrade

At the end of the process:

- If the original primary boot disk was encapsulated, the alternate boot disk is encapsulated.
- The alternate boot environment is activated.
- The system is booted from the alternate boot disk.

### To complete the Live Upgrade

- 1 Complete the Live upgrade process using one of the following commands:

If the primary root disk is not encapsulated, run the following command:

```
# ./vxlufinish -u target_os_version
Live Upgrade finish on the Solaris release <5.10>
```

If the primary root disk is encapsulated by VxVM, run the following command:

```
# ./vxlufinish -u target_os_version -g diskgroup
Live Upgrade finish on the Solaris release <5.10>
```

The Live Upgrade process encapsulates the alternate root disk if the primary root disk was encapsulated.

- 2 If the system crashes or reboots before Live Upgrade completes successfully, you may remount the alternate disk using the following command:

```
# ./vxlustart -r -u target_os_version
```

Then, rerun the `vxlufinish` command:

```
# ./vxlufinish -u target_os_version
```

- 3 On Solaris x64 systems: Restore the original `PrivNIC.cf` file:

```
# cp /var/tmp/PrivNIC.cf.save \
/altroot.5.10/etc/VRTSvcs/config/config/PrivNIC.cf
```

- 4 If the Oracle database is managed by VCS, modify the VCS configuration file on the alternate root disk (/altroot.5.10/etc/VRTSvcs/conf/config/main.cf) to set the **AutoStart** value to 0. This prevents the database service group from starting automatically when VCS starts:

```
group oradb_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoStart = 0
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

.
.
```

If the database is not managed by VCS, change the management policy for the database to manual on the primary boot disk:

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

- 5 Stop the SF Oracle RAC components and running applications that are not controlled by VCS:

- Stop Oracle Clusterware on each node in the cluster:

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

- Stop the applications using native application commands.
- Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu mount-point
```

- Take offline all VCS groups that contain CFSMount and CVMVolDg:

```
# hagrp -offline group -sys galaxy
# hagrp -offline group -sys nebula
```

- Unmount the VxFS file systems:

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

- Deport CVM disk groups:

```
# vxkg deport diskgroup_name
```

- Make sure that no disk groups are imported:

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

- 6 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

---

**Note:** DO NOT use the reboot, halt, or uadmin commands to reboot the system. Use either the init or the shutdown commands to enable the system to boot using the alternate boot environment.

---

```
# shutdown -g0 -y -i6
```

- 7 If the Oracle database is managed by VCS, modify the VCS configuration file on the alternate root disk (/altroot.5.10/etc/VRTSvcs/conf/config/main.cf) to set the AutoStart value to 1.

```
group oradb_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoStart = 1
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)
.
.
```

If the database is not managed by VCS, change the management policy for the database to automatic on the primary boot disk:

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



- 8 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA Installation Guide.

- 9 Complete the post-upgrade tasks.  
See the chapter "Performing post-upgrade tasks" in this document.
- 10 If you are on an unsupported version of Oracle RAC, upgrade Oracle RAC.  
For instructions, see the chapter *Upgrading Oracle RAC* in this document.

## Verifying Live Upgrade of SF Oracle RAC

To ensure that Live Upgrade has completed successfully, verify that all the nodes have booted from the alternate boot environment and joined the cluster.

### To verify that Live Upgrade completed successfully

- 1 Verify that the alternate boot environment is active.

```
# lustatus
```

If the alternate boot environment is not active, you can revert to the primary boot environment.

See [“Reverting to the primary boot environment”](#) on page 282.

- 2 In a cluster environment, make sure that all the GAB ports are up. Note different ports appear for different products.

```
# gabconfig -a
Port a gen 39d901 membership 01
Port b gen 39d905 membership 01
Port d gen 39d904 membership 01
Port f gen 39d90f membership 01
Port h gen 39d909 membership 01
Port o gen 39d903 membership 01
Port u gen 71520b membership 01
Port v gen 39d90b membership 01
Port w gen 39d90d membership 01
```

- 3 Perform other verification as required to ensure that the new boot environment is configured correctly.
- 4 In a zone environment, verify the zone configuration.

## Reverting to the primary boot environment

If the alternate boot environment fails to start, you can revert to the primary boot environment.

On each node, start the system from the primary boot environment in the PROM monitor mode.

```
ok> boot disk0
```

# Upgrading from Storage Foundation products to SF Oracle RAC 5.1 SP1

This chapter includes the following topics:

- [About upgrading from Storage Foundation products to SF Oracle RAC 5.1 SP1](#)
- [Upgrading from Storage Foundation products to SF Oracle RAC 5.1 SP1](#)

## About upgrading from Storage Foundation products to SF Oracle RAC 5.1 SP1

You can upgrade from earlier versions of the following Storage Foundation products to SF Oracle RAC 5.1 SP1:

- Storage Foundation High Availability
- Storage Foundation Cluster File System

---

**Note:** The installer does not support direct upgrade from earlier versions of Storage Foundation products, such as Storage Foundation Cluster File System 5.0, to SF Oracle RAC 5.1 SP1. You must first upgrade to version 5.1 Service Pack 1 of the installed product, then install SF Oracle RAC 5.1 SP1.

---

# Upgrading from Storage Foundation products to SF Oracle RAC 5.1 SP1

Perform the steps in the following procedure to upgrade Storage Foundation products to SF Oracle RAC 5.1 SP1.

## To upgrade Storage Foundation products to SF Oracle RAC 5.1 SP1

- 1 If you are running an earlier version of the product, upgrade to version 5.1 Service Pack 1 of the product.

For instructions, see the installation guide of the product.

- 2 Install SF Oracle RAC 5.1 SP1.

See “[About installing SF Oracle RAC](#)” on page 91.

- 3 When you are prompted to configure SF Oracle RAC, enter **y** at the prompt:

```
Would you like to configure SF Oracle RAC on galaxy? [y,n,q] (n) y
```

The program menu is displayed.

- 4 Select the option **Configure SF Oracle RAC sub-components**.

```
1) Configure SF Oracle RAC sub-components
2) SF Oracle RAC Installation and Configuration Checks
3) Prepare to Install Oracle
4) Install Oracle Clusterware/Grid Infrastructure and Database
5) Post Oracle Installation Tasks
6) Exit SF Oracle RAC Configuration
```

```
Choose option: [1-6,q] (1)
```

If VCS is configured in your environment, the installation program detects the VCS configuration.

---

**Note:** Do not reconfigure VCS when you are prompted by the installation program:

```
Do you want to re-configure VCS? [y,n,q] (n)
```

---

The installation program creates the VCS configuration file and the `/etc/vcsmmtab` file and starts all processes.

## 5 Verify that all GAB ports are up and running.

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen      29d01  membership 01
Port b gen      29d0f  membership 01
Port d gen      29d07  membership 01
Port f gen      29d17  membership 01
Port h gen      29d11  membership 01
Port o gen      29d12  membership 01
Port u gen      71520b membership 01
Port v gen      29d13  membership 01
Port w gen      29d15  membership 01
```

## 6 Restart ODM on all nodes. This will start ODM in clustered mode.

For Solaris 9:

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

For Solaris 10:

```
# svcadm disable -t vxodm
# svcadm enable vxodm
```

## 7 If fencing is not configured, configure fencing.

For instructions, see the chapter *Configuring SF Oracle RAC clusters for data integrity* in this document.

## 8 Install the language packages, if required.

See [“Installing language packages”](#) on page 207.

## 9 ■ If you have neither single-instance Oracle nor Oracle RAC running in your environment, install Oracle RAC.

For instructions, see the chapter *Installing and configuring Oracle RAC* in this document.

## ■ If you have a single-instance Oracle database in your environment, migrate the database to Oracle RAC.

For instructions, see the chapter *Migrating from single instance Storage Foundation for Oracle HA to SF Oracle RAC* in this document.



# Migrating from single instance Storage Foundation for Oracle HA to SF Oracle RAC

This chapter includes the following topics:


- [Migration overview](#)
- [Migration requirements](#)
- [Before you migrate](#)
- [Migrating to SF Oracle RAC 5.1 SP1](#)
- [Sample configuration files](#)

## Migration overview

This chapter provides instructions for migrating from Veritas Storage Foundation for Oracle High Availability to Veritas Storage Foundation for Oracle RAC (SF Oracle RAC).

[Figure 17-1](#) illustrates the migration from Storage Foundation for Oracle HA to SF Oracle RAC.

Figure 17-1 Migration from Storage Foundation for Oracle HA to SF Oracle RAC



Sample configuration before and after migration

Table 17-1 describes a sample existing configuration in a Storage Foundation for Oracle HA database environment.

Table 17-1 Configuration in a Storage Foundation for Oracle HA database environment

| Host name | Instance name | Database name | Oracle home                                                                                                                  | Database file storage     |
|-----------|---------------|---------------|------------------------------------------------------------------------------------------------------------------------------|---------------------------|
| galaxy    | vrts          | vrts          | /u01/app/oracle/product/release/dbhome_1<br><br>where <i>release</i> is 10.2.0 or 11.1.0 depending on the Oracle RAC version | Shared storage using VxFS |
| nebula    | vrts          | vrts          | /u01/app/oracle/product/release/dbhome_1<br><br>where <i>release</i> is 10.2.0 or 11.1.0 depending on the Oracle RAC version | Shared storage using VxFS |

Table 17-2 describes a sample configuration in an SF Oracle RAC database environment.



Table 17-2 Configuration in an SF Oracle RAC environment

| Host name | Instance name | Database name | Oracle home                                                                                                                                        | Database file storage    | OCR and voting disk  |
|-----------|---------------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|----------------------|
| galaxy    | vrts1         | vrts          | /u01/app/oracle/<br>product/ <i>release</i> /dbhome_1<br><br>where <i>release</i> is 10.2.0, 11.1.0, or 11.2.0 depending on the Oracle RAC version | Shared storage using CFS | Reside on CVM or CFS |
| nebula    | vrts2         | vrts          | /u01/app/oracle/<br>product/ <i>release</i> /dbhome_1<br><br>where <i>release</i> is 10.2.0, 11.1.0, or 11.2.0 depending on the Oracle RAC version | Shared storage using CFS | Reside on CVM or CFS |

Migration requirements

Make sure that you meet the following requirements before migrating to SF Oracle RAC:

- Storage Foundation for Oracle HA version 5.1 Service Pack 1 is installed on the systems.
- Oracle 10g Release 2, Oracle 11g Release 1, or Oracle 11g Release 2 is installed on the systems.

Before you migrate

Before you migrate to SF Oracle RAC, complete the following tasks:

- Set up the single instance Storage Foundation for Oracle database with the following configuration:
  - Storage Foundation for Oracle HA installed on all nodes
  - Oracle binaries installed on each node
  - Oracle database created on shared storage using Veritas File System (VxFS)
- Back up the existing database before the migration.

## Migrating to SF Oracle RAC 5.1 SP1

Perform the following steps to migrate from Storage Foundation for Oracle HA to SF Oracle RAC.

1. Migrate Storage Foundation for Oracle HA to SF Oracle RAC.  
See [“Migrating Storage Foundation for Oracle HA to SF Oracle RAC”](#) on page 290.
2. Migrate single instance Oracle database to Oracle RAC database.  
See [“Migrating a single instance Oracle database to Oracle RAC database”](#) on page 291.
3. Complete the post-migration tasks.  
See [“Completing post-migration tasks”](#) on page 294.

## Migrating Storage Foundation for Oracle HA to SF Oracle RAC

Perform the following steps to migrate from Storage Foundation for Oracle HA to SF Oracle RAC.

### To migrate from Storage Foundation for Oracle HA to SF Oracle RAC

- 1 Log in as a superuser.
- 2 Back up the existing Storage Foundation for Oracle database HA resource configuration:

```
# cp -rp /etc/VRTSvcs/conf/config config.old
```

- 3 Take the database service groups offline:

```
# hagrps -offline group_name -any
```

- 4 If the VxFS file systems are not managed by VCS, unmount them :

```
# umount mount_point
```

- 5 Stop all the other VCS service groups.

To view the current state of the service groups:

```
# hagrps -state
```

To stop each group:

```
# hagrps -offline servicegroup -sys node_name
```

## 6 Freeze the VCS service groups:

```
# haconf -makerw
# hagr -freeze servicegroup -persistent
# haconf -dump -makero
```

## 7 Stop VCS on all nodes:

```
# hastop -all -force
```

## 8 Install and configure SF Oracle RAC.

For information on installing and configuring SF Oracle RAC, see *Section 2: Installation and configuration of SF Oracle RAC* in this guide.

## 9 Unfreeze the VCS service groups:

```
# haconf -makerw
# hagr -unfreeze servicegroup -persistent
# haconf -dump -makero
```

# Migrating a single instance Oracle database to Oracle RAC database

Complete the steps in the following procedure for migrating a single instance Oracle database to the Oracle RAC database.

## To migrate a single instance Oracle database to Oracle RAC database

### 1 Install Oracle Clusterware.

For information on installing Oracle Clusterware, see *Section 4: Installation and upgrade of Oracle RAC* in this guide.

### 2 Import the single instance Storage Foundation for Oracle database HA storage disk group in shared mode:

```
# vxvg -s import oradatadg
```

where oradatadg is a disk group in the Storage Foundation for Oracle database HA environment.

In an Oracle RAC environment, all instances concurrently access a single database. Thus, all datafiles, control files, SPFILE, redo log files and archive log files must reside on shared storage.

- 3 Mount the file system in cluster mode on all nodes:

```
# mount -F vxfs -o cluster,largefiles\  
/dev/vx/dsk/oradatadg/oradatavol /oradata
```

where oradata is a file system in the Storage Foundation for Oracle database HA environment.

- 4 Relink Oracle binaries with Veritas libraries.

See [“Relinking Oracle RAC libraries with the SF Oracle RAC libraries”](#) on page 301.

- 5 Start the database as an Oracle user from any one of the nodes in the cluster using SQLPLUS:

```
$ export ORACLE_SID=vrts  
$ export ORACLE_HOME=/oracle/dbhome  
$ /oracle/dbhome/bin/sqlplus "/as sysdba"
```

where vrts is a database in the Storage Foundation for Oracle database HA environment.

- 6 Add redo logs. Each instance requires its own redo thread. The following example assumes a two node configuration, with galaxy and nebula as the nodes in the cluster. The existing redo thread will be used by galaxy.

To add a new redo thread, log on to nebula as Oracle user and run the following commands:

```
SQL> alter database add logfile thread 2  
group 4 ('/oradata/vrts/redo04.log') size 50M REUSE;  
Database altered
```

```
SQL> alter database add logfile thread 2  
group 5 ('/oradata/vrts/redo05.log') size 50M REUSE;  
Database altered.
```

```
SQL> alter database add logfile thread 2  
group 6 ('/oradata/vrts/redo06.log') size 50M REUSE;  
Database altered
```

- 7 Enable redo log thread. While enabling, it may be designated as a public thread or a private thread:

```
SQL> ALTER DATABASE ENABLE PUBLIC THREAD 2;  
OR  
SQL> ALTER DATABASE ENABLE PRIVATE THREAD 2;
```

## 8 Add UNDO tablespaces for each additional instance:

```
SQL> create undo tablespace UNDOTBS2 datafile\
'/oradata/vrts/undotbs02.dbf' size 500M autoextend on;
```

## 9 Create the cluster views needed for Oracle RAC:

```
SQL> @?/rdbms/admin/catclust
```

## 10 If you are using an SPFILE, create a PFILE from it:

```
SQL> create pfile='/tmp/initORA.ora' from spfile
```

## 11 Edit the Oracle initialization parameter file /tmp/initORA.ora to include cluster parameters:

```
*.cluster_database_instances=2
*.cluster_database=TRUE
vrts1.instance_name=vrts1
vrts2.instance_name=vrts2
vrts1.instance_number=1
vrts2.instance_number=2
vrts1.thread=1
vrts2.thread=2
vrts1.undo_tablespace='UNDOTBS1'
vrts2.undo_tablespace='UNDOTBS2'
vrts1.local_listener= 'LISTENER_GALAXY'
vrts2.local_listener= 'LISTENER_NEBULA'
vrts1.remote_listener= 'LISTENERS_VRTS'
vrts2.remote_listener= 'LISTENERS_VRTS'
```

## 12 Start the Oracle initialization parameter file:

```
SQL> connect / as sysdba
SQL> shutdown immediate
SQL> startup pfile='/tmp/initORA.ora'
```

## 13 If the database starts successfully using the Oracle initialization parameter file, create an spfile in a shared location:

```
SQL> create spfile='sharedlocation/spfiledbrac.ora'\
from pfile='/tmp/initORA.ora'
```

- 14 On each node create a link in \$ORACLE\_HOME/dbs to the shared spfile directory.

For galaxy:

```
cd $ORACLE_HOME/dbs
ln -s sharedlocation/spfiledbrac.ora spfiledbrac1.ora
```

For nebula:

```
cd $ORACLE_HOME/dbs
ln -s sharedlocation/spfiledbrac.ora spfiledbrac2.ora
```

- 15 Add the database and instances to the cluster registry:

```
# srvctl add database -d vrts -o $ORACLE_HOME
# srvctl add instance -d vrts -i vrts1 -n galaxy
# srvctl add instance -d vrts -i vrts2 -n nebula
```

- 16 Stop and start the database:

```
$ srvctl stop database -d vrts
$ srvctl start database -d vrts
```

- 17 For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

Configure Oracle listener using Oracle Net Configuration Assistant (NETCA).

---

**Note:** The NETCA utility requires xterm capability.

---

```
$ export DISPLAY=localhost:0.0
$ ORACLE_HOME/bin/netca
```

The NETCA utility starts Oracle listener on both nodes.

This completes the migration of the single-instance Oracle database to an Oracle RAC configuration.

## Completing post-migration tasks

Perform the steps in the following procedure to complete the migration process. The examples in the procedures assume a two-node SF Oracle RAC cluster with the nodes galaxy and nebula.

## To complete the post-migration tasks:

### 1 For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

Copy the file `tnsnames.ora` to the directory `$ORACLE_HOME/network/admin` on all nodes. A sample `tnsnames.ora` file is given below:

```

VRTS =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = galaxy-vip) (PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP) (HOST = nebula-vip) (PORT = 1521))
    (LOAD_BALANCE = yes)
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = vrts)
    )
  )
VRTS1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = galaxy-vip) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = vrts)
      (INSTANCE_NAME = vrts1)
    )
  )
VRTS2 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = nebula-vip) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = vrts)
      (INSTANCE_NAME = vrts2)
    )
  )
LISTENERS_VRTS =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP) (HOST = galaxy-vip) (PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP) (HOST = nebula-vip) (PORT = 1521))
  )

```

For Oracle RAC 11g Release 2, see the Oracle documentation to update the `tnsnames.ora` file.

- 2 Optional: Add services to a database and assign them to instances. You can create services manually or by using the Database Configuration Assistant (DBCA) utility.

In Real Application Clusters (RAC), a service can span one or more instances and facilitate real workload balancing based on real transaction performance. You can control the instances in your cluster that are allocated to different services at different times.

To create the services manually:

```
# srvctl add service -d vrts -s ORA_TAF -r vrts1, vrts2
```

- 3 Configure the CSSD, PrivNIC or MultiPrivNIC and Oracle resource agents in the VCS configuration file.

For more information see *Section 4: Installation and upgrade of Oracle RAC* in this guide.

- 4 Depending on your deployment needs you may perform the following optional tasks:
  - Update the application (database client) configuration to connect to Oracle RAC database (if required) and to take advantage of Oracle RAC features like load balancing and scalability.
  - Update the VCS configuration file for any additional resources, which were present in the VCS configuration prior to the migration process.

## Migrating Storage Foundation for Databases (SFDB) objects from single instance Oracle to Oracle RAC

After migrating from Storage Foundation HA to Storage Foundation for Oracle RAC, the Storage Foundation for Databases (SFDB) objects created in a single instance Oracle environment will not be available in the Oracle RAC environment.

### To re-create the SFDB repository

- 1 After migrating from Storage Foundation HA to Storage Foundation for Oracle RAC, run `dbed_update`.
- 2 Re-create the SFDB objects that existed in the Storage Foundation HA environment. They will now be supported for SFRAC.

For information on configuring SFDB objects, see the *Storage Foundation: Storage and Availability Management for Oracle Databases* guide.



# Sample configuration files

This section illustrates sample configurations for the following files:

|                                      |                                                                                                              |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------|
| VCS configuration file               | See <a href="#">“VCS configuration file for Storage Foundation for Oracle HA”</a> on page 297.               |
| Oracle initialization parameter file | See <a href="#">“Oracle initialization parameter file for Storage Foundation for Oracle HA”</a> on page 298. |
| tnsnames.ora                         | See <a href="#">“tnsnames.ora for Storage Foundation for Oracle HA”</a> on page 299.                         |
| listener.ora                         | See <a href="#">“listener.ora for Storage Foundation for Oracle HA”</a> on page 299.                         |

## VCS configuration file for Storage Foundation for Oracle HA

The sample VCS configuration file for Storage Foundation for Oracle HA is as follows:

```
include "types.cf"
include "OracleTypes.cf"

cluster tyur (
  UserNames = { admin = bQRjQLqNRmRRpZRlQO }
  Administrators = { admin }
)

system galaxy (
)

system nebula (
)

group sfora (
  SystemList = { galaxy = 0, nebula = 1 }
  AutoStartList = { galaxy, nebula }
)

DiskGroup oradatadg (
  Critical = 0
```

```

        DiskGroup = oradatadg
    )

Mount oradatamnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/datavol"
    FSType = vxfs
    FsckOpt = "-y"
)

Oracle oradb (
    Critical = 0
    Sid = vrts
    Owner = oracle
    Home = "/oracle/dbhome"
)
oradatamnt requires oradatadg
oradb requires oradatamnt

```

## Oracle initialization parameter file for Storage Foundation for Oracle HA

The sample Oracle initialization parameter file for Storage Foundation for Oracle HA is as follows:

```

vrts.__db_cache_size=1375731712
vrts.__java_pool_size=16777216
vrts.__large_pool_size=16777216
vrts.__oracle_base='/oracle/dbhome'\
#ORACLE_BASE set from environment
vrts.__pga_aggregate_target=1224736768
vrts.__sga_target=1811939328
vrts.__shared_io_pool_size=0
vrts.__shared_pool_size=385875968
vrts.__streams_pool_size=0
*.audit_file_dest='/oracle/dbhome/admin/vrts/adump'
*.audit_trail='none'
*.compatible='11.1.0.0.0'
*.control_files='/oradata/vrts/control01.ctl',
'/oradata/vrts/control02.ctl',
'/oradata/vrts/control03.ctl'
*.db_block_size=8192

```

```
*.db_domain=''
*.db_name='vrts'
*.diagnostic_dest='/oracle/dbhome'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=dbracXDB)'
*.log_archive_dest_1='LOCATION=/oradata/archive'
*.log_archive_format='%t_%s_%r.dbf'
*.memory_target=3036676096
*.open_cursors=300
*.processes=150
*.remote_login_passwordfile='EXCLUSIVE'
*.undo_tablespace='UNDOTBS1'
```

## tnsnames.ora for Storage Foundation for Oracle HA

The sample tnsnames.ora for Storage Foundation for Oracle HA is as follows:

```
Host : galaxy
VRTS =
    (DESCRIPTION =
        (ADDRESS_LIST =
            (ADDRESS = (PROTOCOL = TCP) (HOST = galaxy) (PORT = 1521))
        )
        (CONNECT_DATA =
            (SERVICE_NAME = vrts)
        )
    )

Host : Nebula
VRTS =
    (DESCRIPTION =
        (ADDRESS_LIST =
            (ADDRESS = (PROTOCOL = TCP) (HOST = nebula) (PORT = 1521))
        )
        (CONNECT_DATA =
            (SERVICE_NAME = vrts)
        )
    )
```

## listener.ora for Storage Foundation for Oracle HA

The sample listener.ora for Storage Foundation for Oracle HA is as follows:

Host : galaxy

```
LISTENER =  
  (DESCRIPTION_LIST =  
    (DESCRIPTION =  
      (ADDRESS = (PROTOCOL = TCP) (HOST = galaxy) (PORT = 1521))  
      (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1521))  
    )  
  )
```

Host : Nebula

```
LISTENER =  
  (DESCRIPTION_LIST =  
    (DESCRIPTION =  
      (ADDRESS = (PROTOCOL = TCP) (HOST = nebula) (PORT = 1521))  
      (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1521))  
    )  
  )
```

# Performing post-upgrade tasks

This chapter includes the following topics:

- [Relinking Oracle RAC libraries with the SF Oracle RAC libraries](#)
- [Setting or changing the product license level](#)
- [Changing permissions for Storage Foundation for Databases](#)
- [Upgrading disk layout versions](#)
- [Upgrading VxVM disk group version](#)
- [Post upgrade tasks for migrating the SFDB repository database](#)
- [Upgrading the Array Support Library](#)
- [Converting from QuickLog to Multi-Volume support](#)
- [Verifying the cluster](#)

## Relinking Oracle RAC libraries with the SF Oracle RAC libraries

You must relink the Oracle RAC libraries with the SF Oracle RAC libraries after upgrading SF Oracle RAC.

The steps vary depending on the version of Oracle RAC in use before the upgrade:

- If you upgraded from an SF Oracle RAC version running Oracle RAC 10g Release 2 or Oracle RAC 11g:

See [“To relink Oracle RAC 10g Release 2 or Oracle RAC 11g using the installer”](#) on page 302.

- If you upgraded from an SF Oracle RAC version running Oracle RAC 9i or Oracle RAC 10g Release 1 or if you want to relink the libraries manually for later versions of Oracle RAC:

See [“To relink Oracle RAC 9i or Oracle RAC 10g Release 1”](#) on page 303.

## To relink Oracle RAC 10g Release 2 or Oracle RAC 11g using the installer

### 1 Run the `installsf` installer:

```
# cd /opt/VRTS/install
# ./installsf -configure galaxy nebula
```

### 2 Enter 5 to select the option **Post Oracle Installation Tasks**.

```
1) Configure SF Oracle RAC sub-components
2) SF Oracle RAC Installation and Configuration Checks
3) Prepare to Install Oracle
4) Install Oracle Clusterware/Grid Infrastructure and Database
5) Post Oracle Installation Tasks
6) Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 5
```

### 3 Select the option **Relink Oracle Database Binary**.

```
1) Configure CSSD agent
2) Relink Oracle Database Binary
3) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-3,b,q] (1) 2
```

### 4 Provide the following information:

```
Enter Oracle UNIX user name: [b] (oracle)
Enter Oracle UNIX group name: [b] (oinstall)
Enter absolute path of Oracle Clusterware/Grid Infrastructure
Home directory: [b]
Enter absolute path of Oracle Database Home directory: [b]
```

The installer detects the Oracle version.

**5** Enter **y** to proceed with relinking.

```
Do you want to continue? [y,n,q] (y)
```

**6** After the relinking completes, confirm that the correct ODM library is used.

```
$ ldd $ORACLE_HOME/bin/oracle | grep odm
```

If the library is not linked correctly, no output is displayed. Sometimes, the Oracle ODM library may not be correctly linked with the Veritas ODM library because of the presence of a static ODM library at `$ORACLE_HOME/rdbms/lib/libodm.a`. This is a known issue with Oracle.

To resolve this issue, run the following commands:

```
$ cd $ORACLE_HOME/rdbms/lib
$ mv libodm10.a libodm10.a.backup
$ /usr/ccs/bin/make -f ins_rdbms.mk ioracle
```

Verify that the correct ODM library is used:

```
$ ldd $ORACLE_HOME/bin/oracle | grep odm
libodm10.so => /app/oracle/orahome/lib/libodm10.so
```

For more information, see the Oracle metalink document: 725903.1

**To relink Oracle RAC 9i or Oracle RAC 10g Release 1**

**1** Replace the Oracle `libskgxp` libraries with the corresponding Veritas libraries:

■ For Oracle RAC 10g Release 1:

```
$ cp /opt/VRTSvcs/rac/lib/libskgxp10_ver23_64.so \
$ORACLE_HOME/lib/libskgxp10.so
```

■ For Oracle RAC 9i:

```
$ cp /opt/VRTSvcs/rac/lib/libskgxp10_ver24_64.so \
$ORACLE_HOME/lib/libskgxp9.so
```

**2** Create a soft link for the `libodm` library as follows:

■ For Oracle RAC 10g Release 1 (Solaris SPARC):

```
$ $ cp /usr/lib/sparcv9/libodm.so \
$ORACLE_HOME/lib/libodm10.so
```

For Oracle RAC 10g Release 1 (Solaris x64):

```
$ cp /usr/lib/amd64/libodm.so \  
$ORACLE_HOME/lib/libodm10.so
```

- For Oracle RAC 9i:

```
$ cp /usr/lib/sparcv9/libodm.so \  
$ORACLE_HOME/lib/libodm9.so
```

## Setting or changing the product license level

If you upgrade to this release from a previous release of the Veritas software, the product installer does not change the license keys that are already installed. The existing license keys may not activate new features in this release.

After you upgrade, perform one of the following steps:

- Obtain a valid license key and run the `vxlicinst` command to add it to your system.
- Use the `vxkeyless` command to update the license keys to the keyless license model.

For more information and instructions, see the chapter *Licensing SF Oracle RAC*.

## Changing permissions for Storage Foundation for Databases

Change the permissions on the `/opt/VRTSdbed` directory and reset owner and group settings as appropriate. This is necessary to ensure that the Veritas Storage Foundation for Oracle commands work correctly.

### To change permissions for Storage Foundation for Databases

- 1 Change permissions for the `/opt/VRTSdbed` directory:

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

- 2 Reset the owner to user "oracle" and the group to "dba":

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



## Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout Version 6, Version 7, and Version 8. No prior versions can be created or mounted.

Use the `vxfsconvert` or `vxupgrade` utilities to upgrade older disk layout versions to disk layout Version 8.

See the `vxfsconvert` or `vxupgrade` man pages.

For more information about disk layouts, see the *Veritas File System Administrator's Guide*.

## Upgrading VxVM disk group version

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks, you need to upgrade existing disk group version to 160.

Run the following command on the master node to upgrade the disk group version:

```
# vxdg -T 160 upgrade diskgroup_name
```

The default Cluster Volume Manager protocol version is 100 and does not need to be upgraded.

Run the following command to verify the CVM protocol version:

```
# /opt/VRTS/bin/vxdctl protocolversion
```

If the protocol version is not 100, run the following command to upgrade the version:

```
# /opt/VRTS/bin/vxdctl upgrade
```

## Post upgrade tasks for migrating the SFDB repository database

To continue using the checkpoints or tering policies you created with a 5.0.x or earlier version of Storage Foundation for Oracle RAC, you must perform one of the following procedures after upgrading SF Oracle RAC to 5.1 Service Pack 1:

- After you upgrade from 5.0.x and before you migrate SFDB:  
See [“After upgrading from 5.0.x and before migrating SFDB”](#) on page 311.

- Migrating from a 5.0 SFDB repository database
- Migrating from a 4.x SFDB repository database
- Upgrading without migrating existing Database Storage Checkpoints and SmartTier parameters

## Migrating from a 5.0 repository database to 5.1 Service Pack 1

Perform the following on one node only.

### To migrate from a 5.0 repository database to 5.1 Service Pack 1

- 1 Rename the startup script NO\_S\*vxdbms3 to S\*vxdbms3.  
See [“After upgrading from 5.0.x and before migrating SFDB”](#) on page 311.
- 2 As root, set the Oracle group permission for various directories used by Oracle on one node:

```
# /opt/VRTSdbed/common/bin/sfua_db_config
```

- 3 As root, dump out the old Sybase ASA repository. If you are using SFHA or SF Oracle RAC, you only need to this on one node.

```
# /opt/VRTSdbed/migrate/sfua_rept_migrate
```

- 4 On the same node that you ran `sfua_rept_migrate` run the following command as Oracle user. For each Oracle instance, migrate the old repository data to the SQLite repository.

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

- 5 By default, the repository is created on the filesystem which contains the Oracle SYSTEM tablespace. If you need an alternative repository path, first verify the following requirements:

- Repository path has to be a directory writable by Oracle user.
- The repository must be accessible by all nodes. You can put it in a CFS file system, or put it in a resource group under VCS control so it can be failed over together with the Oracle database.
- The update commands will not be able to verify accessibility of the repository path and will fail if you have not set up the path correctly.

To create an alternate repository path:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME -R \  
Alternate_path
```

- 6 If you are using Database Flashsnap for off-host processing, and if you have a repository on the secondary host that you need to migrate: perform the previous steps on the secondary host.

If you do not have a repository that you need to migrate from 5.0:

As root, set the Oracle group permission for various directories used by Oracle:

```
# /opt/VRTSdbed/common/bin/sfua_db_config
```

- 7 On the primary host, edit your snapplans to remove the "SNAPSHOT\_DG=SNAP\_\*" parameter and add "SNAPSHOT\_DG\_PREFIX=SNAP\_\*. The parameter can be any PREFIX value and not necessarily "SNAP\_".

For example:

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.pdx.symantec.com
PRIMARY_DG=system1_data
SNAPSHOT_DG=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.pdx.symantec.com
PRIMARY_DG=judge_data
SNAPSHOT_DG_PREFIX=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

- 8 On the primary host, revalidate your snapshots using the following command:

```
$ /opt/VRTS/bin/dbed_vmchecksnap -S $ORACLE_SID \
-H $ORACLE_HOME -f SNAPPLAN -o validate
```

This completes the migration of the repository for Database Storage Checkpoints and Database Tiered Storage parameters.

To begin using the Storage Foundation for Databases (SFDB) tools:

See *Storage Foundation: Storage and Availability Management for Oracle Databases*

## Migrating from a 4.x repository database to 5.1 Service Pack 1

If you are upgrading Database Storage Checkpoints or Database FlashSnap information created in SF Oracle RAC 4.x, you can migrate to `/var/vx/vxdba` to save space under the root partition. Migrating to `/var/vx/vxdba` is optional. However, if you do not perform this migration, you cannot remove any file or directory from `/etc/vx/vxdba` to ensure proper operation.

To migrate from `/etc/vx/vxdba` to `/var/vx/vxdba`

- 1 Copy the `/etc/vx/vxdba` directory and contents to `/var/vx/vxdba`.

```
# cp -rp /etc/vx/vxdba /var/vx/vxdba
```

- 2 Remove `/etc/vx/vxdba`.

```
# rm -rf /etc/vx/vxdba
```

- 3 Link the two directories.

```
# ln -s /var/vx/vxdba /etc/vx/vxdba
```

To upgrade the SFDB tools from 4.x to 5.1 Service Pack 1

- 1 # `/opt/VRTSdbed/common/bin/sfua_db_config`

- 2 On one node, as Oracle user, for each Oracle instance, migrate the old repository data to SQLite repository.

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

- 3 By default, the repository is created on the filesystem which contains the Oracle SYSTEM tablespace. If you need an alternative repository path, first verify the following requirements:

- The SFDB repository path has to be a directory writable by Oracle user.

- The repository must be accessible by all nodes. You can put it in a resource group under VCS control so it can be failed over together with the Oracle database.
- The update commands will not be able to verify accessibility of the repository path and will fail if you have not set up the path correctly.

To create an alternate repository path:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME -R \
Alternate_path
```

- 4 On the primary host, edit your snapplans to remove the "SNAPSHOT\_DG=SNAP\_\*" parameter and add "SNAPSHOT\_DG\_PREFIX=SNAP\_\*. The parameter can be any PREFIX value and not necessarily "SNAP\_".

For example:

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=4.0
PRIMARY_HOST=host1
SECONDARY_HOST=host1
PRIMARY_DG=PRODDg
SNAPSHOT_DG=SNAP_PRODDg
ORACLE_SID=PROD
ARCHIVELOG_DEST=/prod_ar
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
```

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=4.0
PRIMARY_HOST=host1
SECONDARY_HOST=host1
PRIMARY_DG=PRODDg
SNAPSHOT_DG_PREFIX=SNAP_PRODDg
ORACLE_SID=PROD
ARCHIVELOG_DEST=/prod_ar
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
```

- 5 If you are using Database Flashsnap for off-host processing, and if you have a repository on the secondary host that you need to migrate: perform steps 1-4 on the secondary host.

If you do not have a repository that you need to migrate from 4.x:

As root, set the Oracle group permission for various directories used by Oracle on one node:

```
# /opt/VRTSdbed/common/bin/sfua_db_config
```

- 6 On the primary host, revalidate your snapshots using the following command:

```
$ /opt/VRTS/bin/dbed_vmchecksnap -S $ORACLE_SID \
-H $ORACLE_HOME -f SNAPPLAN -o validate
```

This completes the migration of the SFDB repository.

To begin using the Storage Foundation for Databases (SFDB) tools:

See *Storage Foundation: Storage and Availability Management for Oracle Databases*

## After upgrading from 5.0.x and before migrating SFDB

When upgrading from SF Oracle RAC version 5.0 to SF Oracle RAC 5.1 Service Pack 1 the S\*vxdbs3 startup script is renamed to NO\_S\*vxdbs3. The S\*vxdbs3 startup script is required by sfua\_rept\_migrate. Thus when sfua\_rept\_migrate is run, it is unable to find the S\*vxdbs3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbs3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

**To prevent S\*vxdbs3 startup script error**

- ◆ Rename the startup script NO\_S\*vxdbs3 to S\*vxdbs3.

## Upgrading the Array Support Library

VxVM provides support for new disk arrays in the form of Array Support Library (ASL) software package.

## Adding JBOD support for storage arrays for which there is not an ASL available

If an array is of type A/A-A, A/P or ALUA and a suitable ASL is not available, the array must be claimed as an JBOD of type A/P. This is to prevent path delays and I/O failures arising. As JBODs are assumed to be type A/A by default, you must create appropriate JBOD entries for such arrays.

**To configure an A/A-A, A/P or ALUA array as a JBOD**

- 1
- Stop all applications, such as databases, from accessing VxVM volumes that are configured on the array, and unmount all VxFS file systems and checkpoints that are configured on the array.
- 2
- Add the array as a JBOD of type A/P:

```
# vxddladm addjbod vid=SUN pid=T300 policy=ap
```

- 3
- If you have not already done so, upgrade the Storage Foundation or VxVM software to 5.1 Service Pack 1. Device discovery will be performed during the upgrade, and the array will be claimed as a JBOD of appropriate type.

If you have already upgraded your system to 5.1 Service Pack 1, run the following command to perform device discovery:

```
# vxdctl enable
```

- 4
- Verify that the array has been added with the policy set to APdisk:

```
# vxddladm listjbod
VID      PID      Opcode Page Code Page Offset SNO length Policy
=====
SUN      T300      18       -1        36          12          APdisk
```

- 5
- Check that the correct devices are listed for the array:

```
# vxdisk list
DEVICE      TYPE          DISK      GROUP      STATUS
APdisk_0    auto:cdsdisk  -         -          online invalid
APdisk_1    auto:cdsdisk  -         -          online invalid
APdisk_2    auto:cdsdisk  -         -          online invalid
...
```



## Unsuppressing DMP for EMC PowerPath disks

This section is only applicable if you are upgrading a system that includes EMC PowerPath disks.

In releases of VxVM before 4.1, a combination of DMP subpaths and the controllers of DMP subpaths were usually suppressed to prevent interference between DMP and the EMC PowerPath multipathing driver. Suppression has the effect of hiding these subpaths and their controllers from DMP, and as a result the disks on these subpaths and controllers cannot be seen by VxVM.

VxVM 4.1 and later releases have the ability to discover EMCpower disks, and configure them as autodiscovered disks that DMP recognizes are under the control of a separate multipathing driver. This has the benefit of allowing such disks to be reconfigured in cluster-shareable disk groups. Before upgrading to VxVM 5.1 Service Pack 1, you must remove the suppression of the subpaths and controllers so that DMP can determine the association between EMCpower metadevices and `c#t#d#` disk devices.

In the following scenarios, you may need to unsuppress DMP subpaths and controllers:

- Converting a foreign disk  
See [“Converting a foreign disk to auto:simple”](#) on page 313.
- Converting a defined disk  
See [“Converting a defined disk to auto:simple”](#) on page 316.
- Converting a powervxvm disk  
See [“Converting a powervxvm disk to auto:simple”](#) on page 319.

Because emcpower disks are auto-discovered, the powervxvm script should be disabled and removed from the startup script. To remove the powervxvm script, use the command:

```
# powervxvm remove
```

### Converting a foreign disk to auto:simple

Release 4.0 of VxVM provided the `vxddladm addforeign` command to configure foreign disks with default disk offsets for the private and public regions, and to define them as simple disks. A foreign disk must be manually converted to `auto:simple` format before upgrading to VxVM 5.1 Service Pack 1.

If the foreign disk is defined on a slice other than `s2`, you must copy the partition entry for that slice to that for `s0` and change the tag. If the tag of the original slice is changed, the status of the disk is seen as `online:aliased` after the upgrade.

The following example is used to illustrate the procedure. The `vxdisk list` command can be used to display the EMCpower disks that are known to VxVM:

```
# vxdisk list
DEVICE          TYPE          DISK    GROUP    STATUS
c6t0d12s2       auto:sliced    -       -         online
emcpower10c     simple         fdisk   fdg       online
...
```

The `vxprint` command is used to display information about the disk group, `fdg`:

```
# vxprint
Disk group: fdg
TY NAME    ASSOC          KSTATE  LENGTH    PLOFFS    STATE  TUTILO  PUTILO
dg fdg     fdg            -       -         -         -       -       -
dm fdisk   emcpower10c    -       17673456  -         -       -       -
...
```

### To convert a foreign disk to `auto:simple` format

- 1 Stop all the volumes in the disk group, and then deport it:

```
# vxvol -g fdg stopall
# vxdg deport fdg
```

- 2 Use the `vxddladm` command to remove definitions for the foreign devices:

```
# vxddladm rmforeign blockpath=/dev/dsk/emcpower10c \
  charpath=/dev/rdisk/emcpower10c
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
# vxdisk list
DEVICE          TYPE          DISK    GROUP    STATUS
c6t0d12s2       auto:sliced    -       -         online
...
```

- 3 Run the `vxprtvtoc` command to retrieve the partition table entry for the device:

```
# /etc/vx/bin/vxpvtvtoc -f /tmp/vtoc /dev/rdisk/emcpower10c
```

#### 4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
# /etc/vx/bin/vxedvtoc -f /tmp/vtoc /dev/rdisk/emcpower10c

# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
# SLICE      TAG  FLAGS   START  SIZE
# 0          0x0  0x201   0      0
# 1          0x0  0x200   0      0
# 2          0x5  0x201   0      17675520

# THE NEW PARTITIONING WILL BE AS FOLLOWS:
# SLICE      TAG  FLAGS   START  SIZE
# 0          0xf  0x201   0      17675520
# 1          0x0  0x200   0      0
# 2          0x5  0x201   0      17675520

DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

#### 5 Upgrade to VxVM 5.1 Service Pack 1 using the appropriate upgrade procedure.

- 6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to `auto:simple` format:

```
# vxdisk list
```

| DEVICE       | TYPE        | DISK | GROUP | STATUS |
|--------------|-------------|------|-------|--------|
| c6t0d12s2    | auto:sliced | -    | -     | online |
| emcpower10s2 | auto:simple | -    | -     | online |
| ...          |             |      |       |        |

To display the physical device that is associated with the metadvice, `emcpower10s2`, enter the following command:

```
# vxdmpadm getsubpaths dmpnodename=emcpower10s2
```

- 7 Import the disk group and start the volumes:

```
# vxdg import fdg
# vxvol -g fdg startall
```

You can use the `vxdisk list` command to confirm that the disk status is displayed as `online:simple`:

```
# vxdisk list
```

| DEVICE       | TYPE        | DISK  | GROUP | STATUS |
|--------------|-------------|-------|-------|--------|
| c6t0d12s2    | auto:sliced | -     | -     | online |
| emcpower10s2 | auto:simple | fdisk | fdg   | online |

## Converting a defined disk to auto:simple

In VxVM 4.0, and particularly in prior releases, EMCpower disks could be defined by a persistent disk access record (`darec`), and identified as simple disks. If an EMCpower disk is defined with a persistent `darec`, it must be manually converted to `auto:simple` format before upgrading to VxVM 5.1 Service Pack 1.

If the defined disk is defined on a slice other than `s2`, you must copy the partition entry for that slice to that for `s0` and change the tag. If the tag of the original slice is changed, the status of the disk is seen as `online:aliased` after the upgrade.

The following example is used to illustrate the procedure. The `ls` command shows the mapping of the EMC disks to persistent disk access records:

```
# ls -l /dev/vx/dmp/emcdisk1
lrwxrwxrwx 1 root other 36 Sep 24 17:59 /dev/vx/dmp/emcdisk1->
/dev/dsk/c6t0d11s5
# ls -l /dev/vx/rdmp/emcdisk1
```

```
lrwxrwxrwx 1 root other 408Sep 24 17:59 /dev/vx/rdmp/emcdisk1->
/dev/dsk/c6t0d11s5
```

Here the fifth partition of `c6t0d11s5` is defined as the persistent disk access record `emcdisk1`.

The `vxdisk list` command can be used to display the EMCpower disks that are known to VxVM:

```
# vxdisk list
```

| DEVICE    | TYPE        | DISK  | GROUP | STATUS |
|-----------|-------------|-------|-------|--------|
| c6t0d12s2 | auto:sliced | -     | -     | online |
| emcdisk1  | simple      | fdisk | fdg   | online |
| ...       |             |       |       |        |

The `vxprint` command is used to display information about the disk group, `fdg`:

```
# vxprint
```

Disk group: `fdg`

| TY  | NAME  | ASSOC    | KSTATE | LENGTH   | PLOFFS | STATE | TUTILO | PUTILO |
|-----|-------|----------|--------|----------|--------|-------|--------|--------|
| dg  | fdg   | fdg      | -      | -        | -      | -     | -      | -      |
| dm  | fdisk | emcdisk1 | -      | 17673456 | -      | -     | -      | -      |
| ... |       |          |        |          |        |       |        |        |

To convert a disk with a persistent disk access record to `auto:simple` format

- 1 Stop all the volumes in the disk group, and then deport it:

```
# vxvol -g fdg stopall
# vxdg deport fdg
```

- 2 Use the `vxdisk rm` command to remove the persistent record definitions:

```
# vxdisk rm emcdisk1
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
# vxdisk list
```

| DEVICE    | TYPE        | DISK | GROUP | STATUS |
|-----------|-------------|------|-------|--------|
| c6t0d12s2 | auto:sliced | -    | -     | online |
| ...       |             |      |       |        |

- 3 Use the `vxprtvtoc` command to retrieve the partition table entry for the device:

```
# /etc/vx/bin/vxprtvtoc -f /tmp/hdisk /dev/rdsk/c6t0d11s2
```

- 4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
# /etc/vx/bin/vxedvtoc -f /tmp/hdisk /dev/rdisk/c6t0d11s2

# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
# SLICE      TAG  FLAGS      START    SIZE
# 4          0x0  0x200      0         0
# 5          0x0  0x200    3591000 2100375
# 6          0x0  0x200      0         0

# THE NEW PARTITIONING WILL BE AS FOLLOWS:
# SLICE      TAG  FLAGS      START    SIZE
# 4          0x0  0x200      0         0
# 5          0xf  0x200    3591000 2100375
# 6          0x0  0x200      0         0

DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

- 5 Upgrade to VxVM 5.1 Service Pack 1 using the appropriate upgrade procedure.

- 6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to `auto:simple` format:

```
# vxdisk list
```

| DEVICE       | TYPE        | DISK | GROUP | STATUS         |
|--------------|-------------|------|-------|----------------|
| c6t0d12s2    | auto:sliced | -    | -     | online         |
| emcpower10s2 | auto:simple | -    | -     | online:aliased |
| ...          |             |      |       |                |

To display the physical device that is associated with the metadevice, `emcpower10s2`, enter the following command:

```
# vxdmpadm getsubpaths dmpnodename=emcpower10s2
```

- 7 Import the disk group and start the volumes:

```
# vxdg import fdg
# vxvol -g fdg startall
```

You can use the `vxdisk list` command to confirm that the disk status is displayed as `online:simple`:

```
# vxdisk list
```

| DEVICE       | TYPE        | DISK  | GROUP | STATUS         |
|--------------|-------------|-------|-------|----------------|
| c6t0d12s2    | auto:sliced | -     | -     | online         |
| emcpower10s2 | auto:simple | fdisk | fdg   | online:aliased |

To allow DMP to receive correct enquiry data, the common Serial Number (C-bit) Symmetrix Director parameter must be set to enabled.

## Converting a powervxvm disk to auto:simple

In VxVM 4.0, and particularly in prior releases, EMCpower disks could be defined by a persistent disk access record (darec) using `powervxvm` script, and identified as simple disks. If an EMCpower disk is used using `powervxvm`, it must be manually converted to `auto:simple` format before upgrading to VxVM 5.1 Service Pack 1.

If there are any controllers or devices that are suppressed from VxVM as `powervxvm` requirement, then such controllers/disks must be unsuppressed. This is required for Veritas DMP to determine the association between PowerPath metanodes and their subpaths. After the conversion to `auto:simple` is complete, the `powervxvm` script is no longer useful, and should be disabled from startup script.

The following example is used to illustrate the procedure. The `ls` command shows the mapping of the EMC disks to persistent disk access records:

```
# ls -l /dev/vx/rdmp/
crw----- 1 root      root      260, 76 Feb  7 02:36 emcpower0c

# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
c6t0d12s2      auto:sliced   -             -              online
emcpower0c     simple       ppdisk01      ppdg           online

# vxprint
Disk group: fdg
TY NAME        ASSOC        KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg ppdg        ppdg        -        -        -        -        -        -
dm ppdisk01    emcpower0c -      2094960 -        -        -        -
```

To convert an EMCpower disk (defined using `powervxvm`) to `auto:simple` format

- 1 Stop all the volumes in the disk group, and then deport it:

```
# vxvol -g ppdg stopall
# vxdg deport ppdg
```

- 2 Use the `vxdisk rm` command to remove all emcpower disks from VxVM:

```
# vxdisk rm emcpower0c
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
c6t0d12s2      auto:sliced   -             -              online
```

- 3 Use the `vxprvtoc` command to retrieve the partition table entry for this device:

```
# /etc/vx/bin/vxprvtoc -f /tmp/vtoc /dev/vx/rdmp/emcpower0c
```



#### 4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
# /etc/vx/bin/vxedvtoc -f /tmp/vtoc /dev/vx/rdmp/emcpower0c
# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
# SLICE      TAG  FLAGS  START  SIZE
# 0          0x0  0x201  0      0
# 1          0x0  0x200  0      0
# 2          0x5  0x201  0      17675520

# THE NEW PARTITIONING WILL BE AS FOLLOWS:
# SLICE      TAG  FLAGS  START  SIZE
# 0          0xf  0x201  0      17675520
# 1          0x0  0x200  0      0
# 2          0x5  0x201  0      17675520

DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

#### 5 Upgrade to VxVM 5.1 Service Pack 1 using the appropriate upgrade procedure.

#### 6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to auto:simple format:

```
# vxdisk list
```

| DEVICE      | TYPE        | DISK | GROUP | STATUS |
|-------------|-------------|------|-------|--------|
| c6t0d12s2   | auto:sliced | -    | -     | online |
| emcpower0s2 | auto:simple | -    | -     | online |

#### 7 Import the disk group and start the volumes.

```
# vxdg import ppdg
# vxvol -g ppdg startall
# vxdisk list
```

| DEVICE      | TYPE        | DISK    | GROUP | STATUS |
|-------------|-------------|---------|-------|--------|
| c6t0d12s2   | auto:sliced | -       | -     | online |
| emcpower0s2 | auto:simple | ppdsk01 | ppdg  | online |

## Converting from QuickLog to Multi-Volume support

The 4.1 release of the Veritas File System is the last major release to support QuickLog. The Version 6 or Version 7 disk layout does not support QuickLog. The

functionality provided by the Veritas Multi-Volume Support (MVS) feature replaces most of the functionality provided by QuickLog.

The following procedure describes how to convert from QuickLog to MVS. Unlike QuickLog, which allowed logging of up to 31 VxFS file systems to one device, MVS allows intent logging of only one file system per device. Therefore, the following procedure must be performed for each file system that is logged to a QuickLog device if Version 6 or Version 7 disk layout is used.

The QuickLog device did not need to be related to the file system. For MVS, the log volume and the file system volume must be in the same disk group.

### To convert Quicklog to MVS

- 1 Select a QuickLog-enabled file system to convert to MVS and unmount it.

```
# umount myfs
```

- 2 Detach one of the QuickLog volumes from the QuickLog device that the file system had been using. This volume will be used as the new intent log volume for the file system.

```
# qllogdetach -g diskgroup log_vol
```

- 3 Create the volume set.

```
# vxvset make myvset myfs_volume
```

- 4 Mount the volume set.

```
# mount -F vxfs /dev/vx/dsk/rootdg/myvset /mnt1
```

- 5 Upgrade the volume set's file system to Version 6 or Version 7 disk layout.

For example:

```
# vxupgrade -n 6 /mnt1
```

- 6 Add the log volume from step 2 to the volume set.

```
# vxvset addvol myvset log_vol
```

- 7 Add the log volume to the file system. The size of the volume must be specified.

```
# fsvoladm add /mnt1 log_vol 50m
```

- 8 Move the log to the new volume.

```
# fsadm -o logdev=log_vol,logsize=16m /mnt1
```

## Verifying the cluster

The installer performs the following checks on each node during the process:

- Verifies that SF Oracle RAC is installed.
- Verifies whether the Oracle process daemon (oproc) is running or not.
- If Oracle RAC 10g Release 2 or Oracle RAC 11g Release 1 is installed, verifies that LLT and Oracle Clusterware uses the same node IDs for each node in the cluster.
- Verifies that the Veritas libraries (VCS IPC, ODM, and VCSMM) are linked with the Oracle RAC libraries. The VCS IPC check is skipped if Oracle RAC 11g is installed.
- Verifies whether the tunables for LLT/LMX/VCSMM tunables are set correctly.
- Verifies that the CSSD resource is configured under the CVM service group.
- Verifies whether the specified Oracle user has read permissions on the /etc/llttab and /etc/llthosts files.
- Verifies that LLT and Oracle Clusterware/Grid Infrastructure have the same view of the number of nodes in the cluster.
- Verifies that the Oracle RAC database version is the same on all nodes in the cluster. The check fails if the version or patchset information varies across nodes.

For information on resolving issues that you may encounter during the checks, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*, Chapter: Troubleshooting SF Oracle RAC, Section: Troubleshooting installation and configuration check failures.

The installer skips the following checks when LLT is configured over UDP:

- LLT links' full duplex setting
- LLT link jumbo frame setting (MTU)
- LLT links' cross connection

- LLT links' speed and auto negotiation settings

Run the SF Oracle RAC installer to verify the functional health of the cluster.

#### To verify the cluster

- 1 Start the SF Oracle RAC installer:

```
# cd /opt/VRTS/install

# ./installsffrac -configure galaxy nebula
```

- 2 Select the option **SF Oracle RAC Installation and Configuration Checks**.

```
1) Configure SF Oracle RAC sub-components
2) SF Oracle RAC Installation and Configuration Checks
3) Prepare to Install Oracle
4) Install Oracle Clusterware/Grid Infrastructure and Database
5) Post Oracle Installation Tasks
6) Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 2
```

- 3 Enter **y** to confirm that Oracle RAC is installed on the nodes.
- 4 Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory.

---

**Note:** The installer tries to discover the location of the Oracle Clusterware/Grid Infrastructure home directory from the Oracle inventory. If the installer discovers the information, you are not prompted for this information.

---

The installer verifies the path information and re-prompts if the path is incorrect.

- 5 Enter the full path of the Oracle RAC database home directory.

---

**Note:** The installer tries to discover the location of the Oracle RAC database home directory from the Oracle inventory. If the installer discovers the information, you are not prompted for this information.

---

The installer verifies the path information.

- 6 Enter the Oracle user name.

The installer displays the location of the temporary log file and starts the verification process.

## Installation and upgrade of Oracle RAC

- [Chapter 19. Before installing Oracle RAC](#)
- [Chapter 20. Installing Oracle RAC](#)
- [Chapter 21. Performing Oracle RAC post-installation tasks](#)
- [Chapter 22. Upgrading Oracle RAC](#)



# Before installing Oracle RAC

This chapter includes the following topics:

- [About preparing to install Oracle RAC](#)
- [Preparing to install Oracle RAC using the SF Oracle RAC installer or manually](#)

## About preparing to install Oracle RAC

Use one of the following ways to perform the pre-installation tasks:

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC installer | <p>The SF Oracle RAC installer provides a menu-driven command line interface to step you through the pre-installation tasks.</p> <p><b>Note:</b> Some of the pre-installation steps are not supported by the SF Oracle RAC installer and must be done manually as described in the manual procedures.</p>                                                                                                                                                                                                                           |
| Manual                  | <p>You need to perform the pre-installation tasks manually as described in the manual procedures.</p>                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Response file           | <p>You can pre-configure the systems for Oracle RAC installation using an SF Oracle RAC response file. The SF Oracle RAC response file in tandem with the Oracle RAC response files simplify the process of automating and standardizing Oracle RAC installations.</p> <p><b>Note:</b> You can use the response file to automate only those pre-configuration tasks that are supported by the SF Oracle RAC installer.</p> <p>For instructions, see the chapter "Installing Oracle RAC using a response file" in this document.</p> |

The examples in this chapter assume a two-node cluster comprising the nodes galaxy and nebula.

Before installing Oracle RAC, review the Oracle installation manuals and the appropriate Oracle support Web sites.

---

**Note:** Some of the pre-installation tasks, wherever indicated in the document, must be done in accordance with the instructions in the Oracle installation manuals. The instructions for these tasks are not provided in this document.

---

## Preparing to install Oracle RAC using the SF Oracle RAC installer or manually

This section provides instructions for both manual as well as SF Oracle RAC installer-based procedures.

### To prepare to install Oracle RAC

- 1 Identify the public virtual IP addresses for use by Oracle.  
See [“Identifying the public virtual IP addresses for use by Oracle”](#) on page 329.
- 2 Set the kernel parameters.  
See [“Setting the kernel parameters”](#) on page 330.
- 3 Verify that the user nobody exists.  
See [“Verifying the user “nobody” exists”](#) on page 330.
- 4 Launch the SF Oracle RAC installer.  
See [“Launching the SF Oracle RAC installer”](#) on page 330.
- 5 Create Oracle user and groups.  
See [“Creating users and groups for Oracle RAC”](#) on page 331.
- 6 Create the storage for OCR and voting disk.  
See [“Creating storage for OCR and voting disk ”](#) on page 334.
- 7 Configure the private network for Oracle RAC.  
See [“Configuring private IP addresses for Oracle RAC”](#) on page 347.
- 8 Create the Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually.  
See [“Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually”](#) on page 360.
- 9 Set up Oracle user equivalence on all nodes.  
See [“Setting up user equivalence”](#) on page 376.



- 10 Verify whether the Veritas Membership library is linked to the Oracle library.  
See [“Verifying whether the Veritas Membership library is linked to Oracle libraries”](#) on page 376.
- 11 Verify the cluster configuration before the installation to ensure that there are no issues that may prevent a successful installation.  
See [“Verifying the systems for Oracle RAC installation”](#) on page 377.

## Identifying the public virtual IP addresses for use by Oracle

Identify separate public virtual IP addresses for each node in the cluster. Oracle requires one public virtual IP address for the Oracle listener process on each node. Public virtual IP addresses are used by client applications to connect to the Oracle database. Oracle Clusterware/Grid Infrastructure manages the virtual IP addresses.

The IP address and the corresponding host name should be registered in the domain name service (DNS). Alternatively, an entry for the virtual IP address and virtual public name can be placed in the `/etc/hosts` file as shown in the following example:

```
10.182.79.239 galaxy-vip
10.182.79.240 nebula-vip
```

The `/etc/hosts` file on each node of the cluster should have these entries.

Oracle recommends that the public node name for the virtual IP address be in the following format *hostname-vip*. For example, *galaxy-vip*.

---

**Note:** The public node name (in other words, the alias for the virtual IP address) for the nodes must be different from the host's current fully qualified domain name (FQDN).

---

For Oracle RAC 11g Release 2: Additionally, you need a Single Client Access Name (SCAN) registered in Enterprise DNS that resolves to three IP addresses (recommended) using a round robin algorithm or at least one IP address. The IP addresses must be on the same subnet as your public network in the cluster. SCAN provides a single name for clients to access an Oracle database running in a cluster.

---

**Note:** The virtual IP addresses that are used for SCAN IP resolution must be on the same subnet. Oracle RAC does not support their configuration on different subnets.

---

## Setting the kernel parameters

Set the kernel parameter values to meet Oracle RAC deployment requirements. The Oracle Universal Installer (OUI) verifies the settings at the time of installation to ensure that they satisfy the minimum requirements. The Oracle RAC installation fails if the kernel parameters are not configured properly. The settings can also be tuned to optimize system performance.

For instructions and guidelines, see the Oracle Metalink document: 169706.1

Restart the nodes for the new values to take effect.

## Verifying the user "nobody" exists

To verify the user "nobody" exists on each system in the cluster:

```
# id nobody
uid=60001(nobody) gid=60001(nobody)
```

If the user does not exist, create the user:

```
# groupadd -g 60001 nobody
# useradd -g 60001 -u 60001 nobody
```

---

**Note:** Make sure that the user ID and group ID are the same across the nodes in your cluster.

---

## Launching the SF Oracle RAC installer

The following procedure describes the steps to launch the SF Oracle RAC installer for performing the Oracle RAC pre-installation tasks.

### To launch the SF Oracle RAC installer

- 1 Log in as the root user on any one node and start the SF Oracle RAC installer:

```
# cd /opt/VRTS/install
# ./installsfrc -configure galaxy nebula
```

The program displays the Symantec copyright information as well as the location of the installation logs.

- 2 Review the installer instructions and press **Enter** to proceed. From the configuration program menu, select **Prepare to Install Oracle**.
- 3 Select the pre-installation task you want to perform.

- Enter **1** to select the option **Create Oracle User and Group**:  
See [“Creating users and groups for Oracle RAC”](#) on page 331.
- Enter **2** to select the option **Create Storage for OCR and voting disk**:  
See [“Creating storage for OCR and voting disk ”](#) on page 334.
- Enter **3** to select the option **Oracle Network Configuration**:  
See [“Configuring private IP addresses for Oracle RAC”](#) on page 347.

## Creating users and groups for Oracle RAC

Depending on the Oracle RAC version, create the required users and groups.

Create the following groups and users for Oracle RAC 10g Release 2/ Oracle RAC 11g Release 1:

- Oracle Inventory group
- dba group
- Oracle user

Create the following groups and users for Oracle RAC 11g Release 2:

- Oracle Inventory group
- dba group
- Oracle grid user
- Oracle user

Create additional users and groups as required by Oracle. Before creating Oracle users and groups, see the Oracle documentation.

You must assign Oracle Inventory as the primary group and dba as the secondary group. Oracle requires secondary groups for identifying operating system accounts that have database administrative (SYSDBA) privileges and for those accounts that have limited sets of database administrative (SYSOPER) privileges. Create the groups on all systems and assign the Oracle user to these groups.

Use one of the following ways to create the Oracle user and groups:

|                                   |                                                                                                      |
|-----------------------------------|------------------------------------------------------------------------------------------------------|
| Using the SF Oracle RAC installer | See <a href="#">“Creating Oracle user and groups using the SF Oracle RAC installer”</a> on page 332. |
| Manual                            | See <a href="#">“Creating the Oracle user and groups manually”</a> on page 333.                      |

## Creating Oracle user and groups using the SF Oracle RAC installer

This procedure provides instructions for creating Oracle user and groups using the SF Oracle RAC installer.

### To create Oracle user and groups on all nodes in the cluster

- 1 From the SF Oracle RAC installer menu, enter **1** to select the option **Create Oracle User and Group**.
- 2 Provide the following information for creating the Oracle user and groups: user name, group name, user ID, group ID, and the full path of the Oracle user home directory.

The user ID and group ID must not be in use on any node in the cluster. The installer suggests unused values, which you may use or change as required. The configuration program assigns the same values on all the nodes.

---

**Note:** If you are configuring GCO, then the user IDs and group IDs of all nodes on both the primary and secondary clusters must be the same. While configuring the user ID and group ID values on the secondary site, make sure that they are identical to the values used at the primary site.

---

```
Enter Oracle UNIX user name: [b] oracle
Enter Oracle UNIX group name: [b] oinstall
Enter Oracle user's ID (numerical): [b] (1165) 1000
Enter Oracle group's ID (numerical): [b] (1165) 1000
Enter absolute path of Oracle user's Home directory: [b] /home/oracle
```

If the specified user name does not exist on any node in the cluster, the following prompt is displayed:

```
User oracle does not exist on any node. Do you want to create it
with the information provided [y,n,q] (y)
```

Enter **y** to create the Oracle user with the information provided.

### 3 Enter the information for the secondary group.

```
Do you want to create a secondary group
for Oracle user? [y,n,q,b,?] (n) y
Enter Oracle UNIX group name: [b] (dba, oper etc.) dba
Enter Oracle group's ID (numerical): [b] (1996) 1001

Creating secondary group dba for
Oracle user oracle on galaxy ..... Done
Creating secondary group dba for
Oracle user oracle on nebula ..... Done
Do you want to create another secondary group
for Oracle user? [y,n,q,b,?] (n)
```

### 4 For Oracle RAC 11g Release 2: Repeat the above steps to create the grid user.

```
Enter Oracle UNIX user name: [b] grid
Enter Oracle UNIX group name: [b] (oinstall)
Enter Oracle user's ID (numerical): [b] (1168) 1001
Enter Oracle group's ID (numerical): [b] (1000)
Enter absolute path of Oracle user's Home directory: [b] /home/grid
```

## Creating the Oracle user and groups manually

Depending on the Oracle RAC version, create the necessary Oracle groups and users. Be sure to assign the same group ID, user ID, and home directory for the user on each system.

---

**Note:** When you create the user and group, make sure that you specify a user and group ID that is not in use.

---

To create the operating system Oracle user and group on each system

- 1 Create the 'oinstall' group on each system.

```
# groupadd -g 1000 oinstall  
  
# groupadd -g 1001 dba
```

- 2 Create the Oracle user and the user home directory on each system:

```
# useradd -g oinstall -u 1000 \  
-G dba -m -d /home/oracle oracle
```

- 3 For Oracle RAC 11g Release 2: Create the grid user and the user home directory on each system:

```
# useradd -g oinstall -u 1001 \  
-G dba -m -d /home/grid grid
```

## Creating storage for OCR and voting disk

Create appropriate storage for Oracle Cluster Registry (OCR) and the voting disk depending on the version of Oracle RAC.

|                                                       |                                                                                                       |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| Oracle RAC 10g Release 2/<br>Oracle RAC 11g Release 1 | The OCR and the voting disk can reside on CVM raw volumes or in directories in a cluster file system. |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------|

|                          |                                                                                           |
|--------------------------|-------------------------------------------------------------------------------------------|
| Oracle RAC 11g Release 2 | The OCR and the voting disk can reside on ASM or in directories in a cluster file system. |
|--------------------------|-------------------------------------------------------------------------------------------|

**Note:** You can use CVM raw volumes to create ASM disk groups.

Use one of the following ways to create the storage:

|                         |                                                                                                 |
|-------------------------|-------------------------------------------------------------------------------------------------|
| SF Oracle RAC installer | See <a href="#">“Creating storage for OCR and voting disk using the installer”</a> on page 335. |
|-------------------------|-------------------------------------------------------------------------------------------------|

|        |                                                                                      |
|--------|--------------------------------------------------------------------------------------|
| Manual | See <a href="#">“Creating storage for OCR and voting disk manually”</a> on page 340. |
|--------|--------------------------------------------------------------------------------------|

You need to create CVM volumes or a CFS mount point for database file storage later in the installation process:

## Creating storage for OCR and voting disk using the installer

The SF Oracle RAC installer enables you to create OCR and voting disk storage on CVM raw volumes or on a clustered file system. After creating the storage, the installer adds the storage configuration to VCS for high availability.

If you are creating the OCR and voting disk storage on CVM raw volumes, the installer performs the following tasks:

- Creates CVM volume mirrors for OCR and voting disk
- Creates the OCR and voting disk volumes and sets the ownership
- Starts the volumes
- Adds the CVMVolDg resource to the VCS configuration in the cvm group so that the volumes are brought online automatically when the node starts
- Brings the CVMVolDg resource online

If you are creating the OCR and voting disk storage on CFS, the installer performs the following tasks:

- Creates CVM volume mirrors for OCR and voting disk
- Creates the OCR and voting disk volumes and sets the ownership
- Starts the volumes
- Creates the mount point and mounts it on all the nodes
- Sets the ownership for the CFS mount point
- Adds the CFSSMount and CVMVolDg resources to the VCS configuration in the cvm group so that the resources are brought online automatically when the node starts
- Brings the CFSSMount and CVMVolDg resources online

**To create storage for OCR and voting disk using the SF Oracle RAC installer**

- 1 From the SF Oracle RAC installer menu, enter 2 to select the option **Create Storage for OCR and Voting disk**.**

The following menu displays:

```
1) Create Oracle User and Group
2) Create Storage for OCR and Voting disk
3) Oracle Network Configuration
4) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-4,b,q] (1) 2
```

- 2 Enter y to create the storage.**

```
Do you want the installer to assist you in creating disk groups,
volumes and file systems for Oracle? (Mirroring enabled
by default) [y,n,q] (n) y
```

If you want to create the storage manually, enter **n**. The installer displays instructions for creating the storage manually. You may skip the remaining steps.

- 3 Select an appropriate option for the disk group.**

```
1) Create a disk group
2) Use an existing disk group
b) Back to previous menu
Choose option: [1-2,b,q]
```

- If you choose to create a disk group, the installer displays the list of existing disks that do not belong to any disk group. Specify the disks (by entering the serial numbers displayed next to the disk names) that you want to use to create the disk group.

Enter the name of the disk group.

```
Enter the disk group name: [b] (ocrvotedg)
```

- If you choose to use an existing disk group, the installer displays the existing disk groups that contain at least two disks for mirroring.



#### 4 Review and confirm the configuration information displayed:

```
CVM Master node: galaxy
Selected disks (including mirroring):
  1. Disk_2
  2. Disk_3
Disk group name: ocrvotedg
Is this information correct? [y,n,q] (y)
```

The installer initializes the disk groups.

#### 5 Choose the type of storage.

- To create the storage on CVM raw volumes  
See [“Creating the OCR and voting disk storage on CVM raw volumes”](#) on page 337.
- To create the storage on CFS:  
See [“Creating the OCR and voting disk storage on CFS”](#) on page 338.

### Creating the OCR and voting disk storage on CVM raw volumes

Perform the steps in the following procedure to create the storage for OCR and voting disk on CVM raw volumes.

---

**Note:** For Oracle RAC 11g Release 2, you may create CVM raw volumes to create ASM disk groups that may be used to store the OCR and voting disk information.

---

#### To create the OCR and voting disk storage on CVM raw volumes

##### 1 Enter 1 to select the option **CVM raw volumes**.

```
1 CVM Raw Volume
2 Clustered File System
b Back to previous menu
Select the storage scheme to be used: [1-2,b,q] 1
```

##### 2 Enter the name and size of the volume on which you want to store OCR information.

```
Enter the volume name for OCR: [b] (ocrvol)
Enter the volume size for OCR (in MB): [b] (320)
```

- 3 Enter the name and size of the volume on which you want to store voting disk information.

```
Enter the volume name for Vote: [b] (votevol)
Enter the volume size for Vote (in MB): [b] (320)
```

- 4 Enter the Oracle UNIX user name.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
Enter Oracle UNIX user name: [b] oracle
```

For Oracle RAC 11g Release 2:

```
Enter Oracle UNIX user name: [b] grid
```

- 5 Enter the Oracle UNIX group name.

```
Enter Oracle UNIX group name: [b] (oinstall)
```

- 6 Press **Return** to continue.

- 7 Review and confirm the configuration information. The installer creates the volumes and brings the corresponding resources online.

Press **Return** to continue.

- 8 Verify that the resource is online on all nodes in the cluster.

---

**Note:** It takes a few minutes for the CVMVolDg resource to come online.

---

```
# hares -state ocrvote_voldg_ocrvotedg
```

## Creating the OCR and voting disk storage on CFS

Perform the steps in the following procedure to create the storage for OCR and voting disk on CFS.

## To create the OCR and voting disk storage on CFS

### 1 Enter **2** to select the option **CFS Mounted**.

```
1 CVM Raw Volume
2 Clustered File System
b Back to previous menu
Select the storage scheme to be used: [1-2,b,q] 2
```

### 2 Enter the name and size of the volume on which you want to store OCR and voting disk information.

```
Enter the volume name for OCR and Voting disk: [b] (ocrvotevol)
Enter the volume size for OCR and Voting disk (in MB): [b] (640)
```

### 3 Enter the Oracle UNIX user name.

The Oracle UNIX user name was previously set up when you performed the pre-installation task of creating Oracle user and group.

```
Enter Oracle UNIX user name: [b] oracle
```

### 4 Enter the Oracle UNIX group name.

The Oracle UNIX group name was previously set up when you performed the pre-installation task of creating Oracle user and group.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
Enter Oracle UNIX group name: [b] (oinstall)
```

For Oracle RAC 11g Release 2:

```
Enter Oracle UNIX group name: [b] (oinstall)
```

### 5 Press **Return** to continue.

### 6 Review and confirm the configuration information. The installer creates and starts the volumes on all nodes in the cluster.

- 7 Enter the CFS mount point for OCR and voting disk information.

Enter the mount point location for CFS (common for all the nodes) [b]  
(/ocrvote)

The installer creates the CFS mount points and sets the ownership. Press **Return** to continue.

- 8 Verify that the ocrvote\_mnt\_ocrvotedg and ocrvote\_voldg\_ocrvotedg resources are online on all nodes in the cluster:

---

**Note:** It takes a few minutes for the CVMVoldg resource to come online.

---

```
# hares -state ocrvote_mnt_ocrvotedg

# hares -state ocrvote_voldg_ocrvotedg
```

## Creating storage for OCR and voting disk manually

Use one of the following storage options to create the OCR and voting disk storage:

|                 |                                                                                          |
|-----------------|------------------------------------------------------------------------------------------|
| CVM raw volumes | See <a href="#">“To create OCR and voting disk volumes on raw volumes”</a> on page 340.  |
| CFS             | See <a href="#">“To create the storage for OCR and voting disks on CFS”</a> on page 341. |

---

**Note:** Whether you create volumes or file system directories, you can add them to the VCS configuration to make them highly available.

---

### To create OCR and voting disk volumes on raw volumes

- 1 Log in as the root user.
- 2 Determine the CVM master:

```
# vxdctl -c mode
```

- 3 On the master node, create a shared disk group:

```
# vxdg -s init ocrvotedg Disk_2 Disk_3
```

- 4 Create mirrored volumes in the shared group for OCR and voting disk:

```
# vxassist -g ocrvotedg make ocrvol 300M nmirrors=2

# vxassist -g ocrvotedg make votevol 300M nmirrors=2
```

- 5 Set the ownership for the volumes:

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# vxedit -g ocrvotedg set group=oinstall user=oracle mode=660
ocrvol
# vxedit -g ocrvotedg set group=oinstall user=oracle mode=660
votevol
```

For Oracle RAC 11g Release 2:

```
# vxedit -g ocrvotedg set group=oinstall user=grid mode=660
ocrvol
# vxedit -g ocrvotedg set group=oinstall user=grid mode=660
votevol
```

- 6 Add the storage resources to the VCS configuration to make them highly available.

See [“Adding the storage resources to the VCS configuration”](#) on page 342.

#### To create the storage for OCR and voting disks on CFS

- 1 Determine the CVM master:

```
# vxdctl -c mode
```

- 2 On the master node, create a shared VxVM disk group:

```
# vxdg -s init ocrvotedg Disk_2 Disk_3
```

- 3 From the CVM master, create a mirrored volume ocrvotevol:

```
# vxassist -g ocrvotedg make ocrvotevol 640M nmirrors=2
```

- 4 From the CVM master, create a file system with the volume ocrvotevol.

```
# mkfs -F vxfs /dev/vx/rdisk/ocrvotedg/ocrvotevol
```

- 5 On each system, create a directory /ocrvote:

```
# mkdir /ocrvote
```

- 6 On each system, mount the directory `/ocrvote`:

```
# mount -F vxfs -o cluster /dev/vx/dsk/ocrvotedg/ocrvotevol  
/ocrvote
```

- 7 From any system, change permissions on `/ocrvote`:

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# chown -R oracle:oinstall /ocrvote
```

For Oracle RAC 11g Release 2:

```
# chown -R grid:oinstall /ocrvote
```

- 8 Add the storage resources to the VCS configuration to make them highly available.

See [“Adding the storage resources to the VCS configuration”](#) on page 342.

### Adding the storage resources to the VCS configuration

The type of storage resource you add to the VCS configuration depends on whether you chose to create the OCR and voting disk storage on raw volumes or CFS. If you chose to create the storage on raw volumes, you need to add a CVMVolDg resource to the VCS configuration. If you chose to create the storage on CFS, you need to add the CVMVolDg and CFSSMount resources to the VCS configuration.

Depending on the type of storage, follow the steps in one of the following procedures to add the storage resources to the VCS configuration using the command line interface (CLI):

For storage resources on CVM raw volumes See [“To add the storage resources created on raw volumes to the VCS configuration”](#) on page 345.

For storage resources on CFS See [“To add the storage resources created on CFS to the VCS configuration”](#) on page 343.

---

**Note:** Set the attribute "Critical" to "0" for all the resources in the cvm service group. This ensures that critical CVM and CFS resources are always online.

---

## To add the storage resources created on CFS to the VCS configuration

### 1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

### 2 Configure the CVM volumes under VCS:

```
# hares -add ocrvote_voldg_ocrvotedg CVMVolDg cvm
# hares -modify ocrvote_voldg_ocrvotedg Critical 0
# hares -modify ocrvote_voldg_ocrvotedg CVMDiskGroup ocrvotedg
# hares -modify ocrvote_voldg_ocrvotedg CVMVolume -add ocrvotevol
# hares -modify ocrvote_voldg_ocrvotedg CVMActivation sw
```

### 3 Set up the file system under VCS:

```
# hares -add ocrvote_mnt_ocrvotedg CFSMount cvm
# hares -modify ocrvote_mnt_ocrvotedg Critical 0
# hares -modify ocrvote_mnt_ocrvotedg MountPoint "/ocrvote"
# hares -modify ocrvote_mnt_ocrvotedg BlockDevice \
"/dev/vx/dsk/ocrvotedg/ocrvotevol"
# hares -modify ocrvote_mnt_ocrvotedg
```

### 4 Link the parent and child resources:

```
# hares -link ocrvote_voldg_ocrvotedg cvm_clus
# hares -link ocrvote_mnt_ocrvotedg ocrvote_voldg_ocrvotedg
# hares -link ocrvote_mnt_ocrvotedg vxfsckd
```

### 5 Enable the resources:

```
# hares -modify ocrvote_voldg_ocrvotedg Enabled 1
# hares -modify ocrvote_mnt_ocrvotedg Enabled 1
# haconf -dump -makero
```

**6** Verify the configuration of the ocrvote\_mnt\_ocrvotedg and ocrvote\_voldg\_ocrvotedg resources in main.cf.

```
CFSMount ocrvote_mnt_ocrvotedg (  
    Critical = 0  
    MountPoint = "/ocrvote"  
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"  
    MountOpt= "mincache=direct"  
)  
CVMVoldg ocrvote_voldg_ocrvotedg (  
    Critical = 0  
    CVMDiskGroup = ocrvotedg  
    CVMVolume = { ocrvotevol }  
    CVMActivation = sw  
)  
ocrvote_mnt_ocrvotedg requires ocrvote_voldg_ocrvotedg  
ocrvote_mnt_ocrvotedg requires vxfsckd  
ocrvote_voldg_ocrvotedg requires cvm_clus
```

**7** Bring the ocrvote\_mnt\_ocrvotedg and ocrvote\_voldg\_ocrvotedg resources online on all systems in the cluster:

```
# hares -online ocrvote_voldg_ocrvotedg -sys sys_name  
  
# hares -online ocrvote_mnt_ocrvotedg -sys sys_name
```

Verify that the resources are online on all systems in the cluster:

```
# hares -state ocrvote_voldg_ocrvotedg  
  
# hares -state ocrvote_mnt_ocrvotedg
```



To add the storage resources created on raw volumes to the VCS configuration

- 1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

- 2 Configure the CVM volumes under VCS:

```
# hares -add ocrvote_voldg_ocrvotedg CVMVolDg cvm
# hares -modify ocrvote_voldg_ocrvotedg Critical 0
# hares -modify ocrvote_voldg_ocrvotedg CVMDiskGroup ocrvotedg
# hares -modify ocrvote_voldg_ocrvotedg CVMVolume -add ocrvol
# hares -modify ocrvote_voldg_ocrvotedg CVMVolume -add votevol
# hares -modify ocrvote_voldg_ocrvotedg CVMActivation sw
```

- 3 Link the parent and child resources:

```
# hares -link ocrvote_voldg_ocrvotedg cvm_clus
```

- 4 Enable the resources:

```
# hares -modify ocrvote_voldg_ocrvotedg Enabled 1
# haconf -dump -makero
```

- 5 Verify the configuration of the ocrvote\_voldg\_ocrvotedg resource in main.cf.

```
CVMVolDg ocrvote_voldg_ocrvotedg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvol, votevol }
    CVMActivation = sw
)
ocrvote_voldg_ocrvotedg requires cvm_clus
```

- 6 Bring the ocrvote\_voldg\_ocrvotedg resource online on all systems in the cluster:

```
# hares -online ocrvote_voldg_ocrvotedg -sys sys_name
```

Verify that the resource is online on all systems in the cluster:

```
# hares -state ocrvote_voldg_ocrvotedg
```

## Setting CVM disk policies for OCR, voting disk, and data disk groups

Table 19-1 describes the CVM disk policies you need to set for Oracle Cluster Registry (OCR), voting disk, and data disk groups.

Table 19-1 Disk group fail policy for OCR, voting disk, and data disk groups

| Disk group          | Policy setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OCR and voting disk | <p>leave</p> <p>If the disk groups containing the OCR and voting disk uses the default policy setting (<code>dgdisable</code>), and the CVM master node loses connectivity to the storage, then a cluster-wide panic results. The CVM master node disables the corresponding disk group after a storage loss and prevents access to volumes from any cluster node. This in turn prevents the slave nodes from performing I/O operations to shared volumes even though issues with storage may be confined to the master node. If Oracle's clusterware is unable to perform I/O operations to voting disks, the corresponding node is panicked and evicted from the cluster.</p> <p>Setting the policy to <code>leave</code> for the disk group ensures that only the node which loses connectivity to the storage is panicked.</p> |
| Data disk groups    | <p>leave</p> <p>If the disk groups containing the Oracle data uses the default policy setting (<code>dgdisable</code>), and if one of the nodes in the cluster loses connectivity to the storage, then the disk group is disabled across the cluster. This causes the database to freeze on all nodes in the cluster.</p> <p>Setting the policy to <code>leave</code> for the disk group ensures that only the node which loses connectivity to the storage is panicked.</p>                                                                                                                                                                                                                                                                                                                                                       |

**Note:** Symantec strongly recommends retaining the default setting (`global`) for the disk detach policy. For other disk detach policy options, see the *Veritas Volume Manager Administrator's Guide*.

To set the disk policies for OCR, voting disk, and data disk groups

- 1
- Log into the CVM master node as the root user.
- 2
- Set the disk group fail policy as follows:

```
# vxdbg -g dg_name set dgfailpolicy=leave
```

## Configuring private IP addresses for Oracle RAC

Private IP addresses are required by Oracle RAC to provide communication between the cluster nodes. Depending on your private network configuration, you may need one or more IP addresses. You can configure the private IP addresses for high availability using the PrivNIC or MultiPrivNIC agents.

**Note:** IPv6 addresses are not supported in this release.

[Table 19-2](#) lists the available options for configuring the private network for Oracle RAC. Use one of the following options to configure the private network.

**Table 19-2** Options for configuring the private network for Oracle RAC

| Option                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PrivNIC configuration | <div>Perform this configuration if you plan to:</div> <div><div>■</div>Use a common IP address for Oracle Clusterware/Grid Infrastructure heartbeat and Oracle RAC database cache fusion and you plan to configure the IP address for high availability using the PrivNIC agent.</div> <div><div>■</div>Or, in the case of Oracle RAC 10g Release 2, use VCSIPC/LMX/LLT for Oracle RAC database cache fusion (you need an IP address only for Oracle Clusterware heartbeat in this scenario) and you plan to configure the IP address for high availability using the PrivNIC agent.</div> |

For instructions:

See [“Configuring the private IP address and PrivNIC resource”](#) on page 348.

**Table 19-2** Options for configuring the private network for Oracle RAC  
(continued)

| Option                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MultiPrivNIC configuration | <p>Perform this configuration if you plan to:</p> <ul style="list-style-type: none"><li>■ Use an IP address on each node for Oracle Clusterware/Grid Infrastructure heartbeat</li><li>■ And, use one or more IP addresses on each node for Oracle database cache fusion</li><li>■ And, configure the IP addresses for high availability using the MultiPrivNIC agent</li></ul> <p>For instructions:</p> <p>See <a href="#">“Configuring the private IP address information and MultiPrivNIC resource”</a> on page 353.</p> |

Configuring the private IP address and PrivNIC resource

You need to configure the following information:

- An IP address on each node  
If you plan to install Oracle RAC 10g Release 2 or Oracle RAC 11g Release 1, the IP addresses configured must be added to the `/etc/hosts` file (using the SF Oracle RAC installer or manually) on all the nodes before installing Oracle Clusterware/Grid Infrastructure.
- The PrivNIC agent for failing over IP addresses in the event of link failures. The Oracle Clusterware/Grid Infrastructure interconnects need to be protected against NIC failures and link failures. The installer discovers the existing private NICs on which LLT is configured. For maximum failover options, all available LLT links are used for PrivNIC configuration.

Use one of the following ways to configure the PrivNIC and private IP address information:

|                         |                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC installer | See <a href="#">“Configuring the private IP address and PrivNIC using the SF Oracle RAC installer”</a> on page 348. |
| Manual                  | See <a href="#">“Configuring the private IP address and PrivNIC resource manually”</a> on page 352.                 |

Configuring the private IP address and PrivNIC using the SF Oracle RAC installer

The SF Oracle RAC installer performs the following tasks:

- Backs up the /etc/hosts file and adds the IP address information to the file (only if you specified that the installer update the file).
- Adds the PrivNIC resource in the CVM group.

Perform the steps in the following procedure to configure the PrivNIC and private IP address using the installer.

#### To configure the PrivNIC and private IP address information

#### 1 From the SF Oracle RAC menu, enter **3** to select the option **Oracle Network Configuration**.

```

1) Create Oracle User and Group
2) Create Storage for OCR and Voting disk
3) Oracle Network Configuration
4) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-4,b,q] (1) 3

```

#### 2 Enter **1** to select the option **Configure private IP addresses (PrivNIC Configuration)**.

```

1) Configure private IP addresses (PrivNIC Configuration)
2) Configure private IP addresses (MultiPrivNIC Configuration)
3) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-3,b,q] (1) 1

```

The installer discovers available LLT links and PrivNIC resources.

If PrivNIC resources exist, you can choose to delete and reconfigure the resources using the installer.

---

**Note:** The installer only removes the corresponding PrivNIC resources from the configuration file. You must manually disassociate the IP addresses from the corresponding network interfaces and remove the IP addresses from the /etc/hosts file.

---

#### 3 Enter the name for the PrivNIC resource.

```

Enter the PrivNIC resource name: [b] (ora_priv)

```

- 4 Enter **y** to modify the priority of the PrivNIC interfaces.

---

**Note:** The priority you set determines the interface that the PrivNIC agent chooses during failover.

---

```
Do you want to update the priority of the PrivNIC
interfaces? [y,n,q] (n) y
```

- 5 Set the interface priority in decreasing order. The PrivNIC agent will assign the highest priority to the first interface specified in the list.

```
Enter the interface name in the decreasing priority order,
separated by a space: [b] (bge2 bge3) bge3 bge2
```

- 6 ■ Enter **y** to add the IP addresses to the `/etc/hosts` file.

```
Do you want the Installer to add IP addresses in /etc/hosts
and /etc/inet/ipnodes files? [y,n,q] (y)
```

---

**Note:** The `/etc/inet/ipnodes` file is used only for IPv6 address updates. Ignore the reference to the `/etc/inet/ipnodes` file in the prompt as this release does not support IPv6 addresses.

---

- Enter **n** if you choose to add the IP addresses to the file manually. Go to step [8](#).

- 7 Perform this step only if you enabled the installer to add the IP address to the `/etc/hosts` file in the previous step.

Provide the private IP address and the private node name for the IP address that must be added to the file.

---

**Note:** All IP addresses must be in the same subnet, failing which Oracle Clusterware/Grid Infrastructure will not be able to communicate properly across the nodes.

If the private IP address entries are already present in the `/etc/hosts` file on one of nodes in the cluster, the installer does not update the file with the specified IP addresses on any of the nodes in the cluster.

---

```
Enter the private IP for galaxy: [b] 192.168.12.1
Enter Hostname alias for the above IP address: [b] galaxy-priv
Enter the private IP for nebula: [b] 192.168.12.2
Enter Hostname alias for the above IP address: [b] nebula-priv
```

Go to step 9.

- 8 Perform this step only if you have chosen to add the IP address to the `/etc/hosts` file manually.

Enter the private IP address information.

```
Enter the private IP for galaxy: [b] 192.168.12.1
Enter the private IP for nebula: [b] 192.168.12.2
```

- 9 Enter the netmask information for the private network:

```
Enter the Netmask for private network: [b] (255.255.255.0)
```

The SF Oracle RAC installer now displays the configuration information.

- 10 Enter **y** to review and confirm the configuration information. The installer adds the PrivNIC resources to the VCS configuration and updates the `/etc/hosts` file (if you chose an installer-based update).
- 11 If you chose to add the IP address information to the `/etc/hosts` file manually, proceed to update the file as described in the following procedure.  
See [“Adding private IP addresses to the /etc/hosts file manually”](#) on page 358.
- 12 Verify the PrivNIC configuration updates made by the program in the `main.cf` file.  
See [“Verifying the VCS configuration for PrivNIC and MultiPrivNIC”](#) on page 359.

## Configuring the private IP address and PrivNIC resource manually

Perform the steps in the following procedure to configure the private IP address and PrivNIC resource manually. Configure the PrivNIC resource in the VCS group where you have configured the OCR and voting disk resources.

The sample procedure creates the PrivNIC resource in the cvm group. The PrivNIC agent plumbs the IP address to the specified network interface.

### To configure the private IP address and PrivNIC resource manually

- 1 Log in as the root user on one of the nodes in the cluster.
- 2 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Create the PrivNIC resource and add the resource to the cvm group:

```
# hares -add ora_priv PrivNIC cvm
```

- 4 Modify the PrivNIC resource:

```
# hares -modify ora_priv Critical 0
# hares -local ora_priv Device
# hares -local ora_priv Address
# hares -modify ora_priv \
Device -add bge1 0 -sys galaxy
# hares -modify ora_priv \
Device -add bge2 1 -sys galaxy
# hares -modify ora_priv \
Address 192.168.12.1 -sys galaxy
# hares -modify ora_priv \
Device -add bge1 0 -sys nebula
# hares -modify ora_priv \
Device -add bge2 1 -sys nebula
# hares -modify ora_priv \
Address 192.168.12.2 -sys nebula
# hares -modify ora_priv \
NetMask 255.255.255.0
# hares -modify ora_priv Enabled 1
```

- 5 Change the cluster configuration to read-only mode:

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



## Configuring the private IP address information and MultiPrivNIC resource

You need to configure the following information:

- An IP address on each node for Oracle Clusterware/Grid Infrastructure heartbeats
- One or more IP addresses on each node for the Oracle database
- The MultiPrivNIC agent for failing over IP addresses in the event of link failures.  
The Oracle Clusterware/Grid Infrastructure interconnects need to be protected against NIC failures and link failures. The MultiPrivNIC agent protects the links against failures, if multiple links are available. The installer discovers the existing private NICs on which LLT is configured. For maximum failover options, all available LLT links are used for MultiPrivNIC configuration.

If you plan to use Oracle RAC 10g Release 2 or Oracle RAC 11g Release 1, the IP addresses configured must be added to the `/etc/hosts` file (using the SF Oracle RAC installer or manually) on all the nodes before installing Oracle Clusterware/Grid Infrastructure.

Use one of the following ways to configure the MultiPrivNIC and private IP address information:

|                         |                                                                                                                                      |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC installer | See <a href="#">“Configuring the MultiPrivNIC and private IP address information using the SF Oracle RAC installer”</a> on page 353. |
| Manual                  | See <a href="#">“Configuring MultiPrivNIC and private IP addresses manually”</a> on page 356.                                        |

### Configuring the MultiPrivNIC and private IP address information using the SF Oracle RAC installer

Perform the steps in the following procedure to configure the MultiPrivNIC and private IP address using the installer.

The installer performs the following tasks:

- Backs up the `/etc/hosts` file and adds the IP address information to the file (only if you specified that the installer update the file).
- Adds the MultiPrivNIC resource in the CVM group.

To configure the MultiPrivNIC and private IP address information using the installer

**1 From the SF Oracle RAC menu, enter 3 to select the option Oracle Network Configuration.**

- 1) Create Oracle User and Group
- 2) Create Storage for OCR and Voting disk
- 3) Oracle Network Configuration
- 4) Exit SF Oracle RAC Configuration

b) Back to previous menu

Choose option: [1-4,b,q] (1) **3**

**2 Enter 2 to select the option Configure private IP addresses (MultiPrivNIC Configuration).**

- 1) Configure private IP addresses (PrivNIC Configuration)
- 2) Configure private IP addresses (MultiPrivNIC Configuration)
- 3) Exit SF Oracle RAC Configuration

b) Back to previous menu

Choose option: [1-3,b,q] (1) **2**

The installer discovers available LLT links and MultiPrivNIC resources. If MultiPrivNIC resources exist, you can choose to delete and reconfigure the resources using the installer.

---

**Note:** The installer only removes the corresponding MultiPrivNIC resources from the configuration file. You must manually disassociate the IP addresses from the corresponding network interfaces and remove the IP addresses from the /etc/hosts file.

---

**3 Enter the name for the MultiPrivNIC resource.**

Enter the MultiPrivNIC resource name: [b] (ora\_mpriv) **multi\_priv**

**4 ■ Enter y to add the IP addresses to the /etc/hosts file.**

Do you want the Installer to add IP addresses in /etc/hosts and /etc/inet/ipnodes files? [y,n,q] (y)

---

**Note:** The /etc/inet/ipnodes file is used only for IPv6 address updates. Ignore the reference to the /etc/inet/ipnodes file in the prompt as this release does not support IPv6 addresses.

---

- Enter **n** if you choose to add the IP addresses to the file manually. Go to step 6.
- 5 Perform this step only if you enabled the installer to add the IP address to the /etc/hosts file in the previous step.

Provide the private IP address and the private node name for the IP addresses that must be added to the file. When you do not want to enter information at the prompts, enter **x**.

---

**Note:** The IP addresses used for a particular NIC on all nodes of a cluster must be in the same subnet. This subnet must be different from the subnets for the IP addresses on other NICs. Otherwise, Oracle Clusterware/Grid Infrastructure and UDP IPC will not be able to communicate properly across the nodes.

If the private IP address entries are already present in the /etc/hosts file on one of the nodes in the cluster, the installer does not update the file with the specified IP addresses on any of the nodes in the cluster.

---

```
Enter IP addresses for galaxy for bge1 separated by space: [b,q,?]
192.168.12.1
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] galaxy-priv
Enter IP addresses for galaxy for bge2
separated by space: [b,q,?] 192.168.2.1
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] galaxy-priv1
Enter IP addresses for nebula for bge1
separated by space: [b,q,?] 192.168.12.2
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] nebula-priv
Enter IP addresses for nebula for bge2
separated by space: [b,q,?] 192.168.2.2
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] nebula-priv1
```

Go to step 7.

- 6 Perform this step only if you have chosen to add the IP address to the /etc/hosts file manually.

Enter the private IP address information.

```
Enter IP addresses for galaxy for bge1
separated by space: [b,q,?] 192.168.12.1
Enter IP addresses for galaxy for bge2
separated by space: [b,q,?] 192.168.2.1
Enter IP addresses for nebula for bge1
separated by space: [b,q,?] 192.168.12.2
Enter IP addresses for nebula for bge2
separated by space: [b,q,?] 192.168.2.2
```

- 7 Enter the netmask information for the private network:

```
Enter the Netmask for private network: [b] (255.255.255.0)
```

The SF Oracle RAC installer displays the configured parameters.

- 8 Enter **y** to review and confirm the configuration information. The installer adds the MultiPrivNIC resources and updates the /etc/hosts file (if you chose installer-based update).
- 9 If you chose to add the IP address information to the /etc/hosts file manually, proceed to update the file as described in the following procedure.  
See [“Adding private IP addresses to the /etc/hosts file manually”](#) on page 358.
- 10 Verify the MultiPrivNIC configuration updates made by the program in the main.cf file.  
See [“Verifying the VCS configuration for PrivNIC and MultiPrivNIC”](#) on page 359.

### Configuring MultiPrivNIC and private IP addresses manually

Perform the steps in the following procedure to configure MultiPrivNIC and private IP addresses for Oracle Clusterware/Grid Infrastructure and UDP IPC manually. Configure the MultiPrivNIC resource in the VCS group where you have configured the OCR and voting disk resources.

Make sure that the number of links are the same for every node. For example, in the following configuration, if the bge1 and bge2 links fail, the MultiPrivNIC agent fails over the IP addresses to the bge3 link on galaxy; On nebula, the absence of a third link results in the agent failing over the link to bge1, which is already down. This results in loss of communication between the nodes and causes Oracle Clusterware/Grid Infrastructure to reboot the cluster.

```
MultiPrivNIC multi_priv (
    Critical = 0
    Device@galaxy = {bge1 = 0, bge2 = 1, bge3 = 2}
    Device@nebula = {bge1 = 0, bge2 = 1}
    Address@galaxy = {"192.168.12.1" = 0, "192.168.2.1" = 1}
    Address@nebula = {"192.168.12.2" = 0, "192.168.2.2" = 1}
    NetMask = "255.255.255.0"
)
```

---

**Note:** Avoid configurations where the number of links differ between the nodes.

---

The sample procedure creates the MultiPrivNIC resource in the cvm group. The MultiPrivNIC agent plumbs the IP addresses to the appropriate network interfaces.

#### To configure the private IP address and MultiPrivNIC manually

- 1 Log in as the root user on one of the nodes in the cluster.
- 2 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Create the MultiPrivNIC resource and add the resource to the cvm group:

```
# hares -add multi_priv MultiPrivNIC cvm
```

**4** Modify the MultiPrivNIC resource:

```
# hares -modify multi_priv Critical 0
# hares -local multi_priv Device
# hares -local multi_priv Address
# hares -modify multi_priv \
Device -add bge1 0 -sys galaxy
# hares -modify multi_priv \
Device -add bge2 1 -sys galaxy
# hares -modify multi_priv \
Address -add 192.168.12.1 0 -sys galaxy
# hares -modify multi_priv \
Address -add 192.168.2.1 1 -sys galaxy
# hares -modify multi_priv \
Device -add bge1 0 -sys nebula
# hares -modify multi_priv \
Device -add bge2 1 -sys nebula
# hares -modify multi_priv \
Address -add 192.168.12.2 0 -sys nebula
# hares -modify multi_priv \
Address -add 192.168.2.2 1 -sys nebula
# hares -modify multi_priv \
NetMask 255.255.255.0
# hares -modify multi_priv Enabled 1
```

**5** Change the cluster configuration to read-only mode:

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

## Adding private IP addresses to the /etc/hosts file manually

Perform the steps in the following procedure only if you plan to install Oracle RAC 10g Release 2 or Oracle RAC 11g Release 1 and you chose to add the private IP address information manually to the `/etc/hosts` file at the time of configuring them using the SF Oracle RAC installer.

### To add private IP addresses to the /etc/hosts file manually

- 1 Log in to each system as the root user.
- 2 For a configuration using the PrivNIC agent, add the following entries to the /etc/hosts file:

```
192.168.12.1 galaxy-priv
192.168.12.2 nebula-priv
```

For a configuration using the MultiPrivNIC agent, add the following entries to the /etc/hosts file:

```
192.168.12.1 galaxy-priv
192.168.2.1 galaxy-priv1
192.168.12.2 nebula-priv
192.168.2.2 nebula-priv1
```

## Verifying the VCS configuration for PrivNIC and MultiPrivNIC

After you complete the steps for configuring PrivNIC/MultiPrivNIC and the private IP addresses, verify the configuration in the VCS main.cf configuration file.

### To verify the VCS configuration for PrivNIC and MultiPrivNIC

- 1 View the main.cf file located in the directory /etc/VRTSvcs/conf/config:
 

```
# more /etc/VRTSvcs/conf/config/main.cf
```
- 2 For a configuration using the PrivNIC agent:
  - Verify that the PrivNIC resource ora\_priv displays in the file:

```
PrivNIC ora_priv (
    Critical = 0
    Device @galaxy = {bge1= 0, bge2= 1}
    Device @nebula = {bge1= 0, bge2= 1}
    Address @galaxy = "192.168.12.1"
    Address @nebula = "192.168.12.2"
    NetMask = "255.255.255.0"
)
```

- Verify that the ora\_priv resource is online on all systems in the cluster:

```
# hares -state ora_priv
Resource Attribute System Value
ora_priv State      galaxy  ONLINE
ora_priv State      nebula  ONLINE
```

- 3
- For a configuration using the MultiPrivNIC agent:
- Verify that the MultiPrivNIC resource `multi_priv` displays in the file:

```
MultiPrivNIC multi_priv (
    Critical = 0
    Device @galaxy = {bge1= 0, bge2 = 1}
    Device @nebula = {bge1= 0, bge2 = 1}
    Address @galaxy = {"192.168.12.1" =0, "192.168.2.1" =1}
    Address @nebula = {"192.168.12.2" =0, "192.168.2.2" =1}
    NetMask = "255.255.255.0"
)
```

- Verify that the `multi_priv` resource is online on all systems in the cluster:

```
# hares -state multi_priv
Resource      Attribute      System      Value
multi_priv    State         galaxy     ONLINE
multi_priv    State         nebula     ONLINE
```

- 4
- On each system, check the output of "ifconfig". Make sure that the specified private IP addresses and devices are displayed.
- # ifconfig -a
- 5
- From each system, verify that the private IP addresses are operational using the `ping` command.

# Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually

You can create the Oracle Clusterware/Grid Infrastructure and Oracle database home directories on the local file system, on a local Veritas file system, or on a Veritas cluster file system. Perform this step if you plan to place the directories on VxFS or CFS. When the installer prompts for the home directories at the time of installing Oracle Clusterware/Grid Infrastructure and Oracle database, it creates the directories locally on each node, if they do not exist.



**Note:** Symantec recommends that you install the Oracle RAC binaries local to each node in the cluster.

Table 19-3 lists the Oracle RAC directories you need to create:

**Table 19-3** List of directories

| Directory                                                                                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oracle Clusterware Home Directory (CRS_HOME)<br><br>(For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1) | <p>The path to the home directory that stores the Oracle Clusterware binaries. The Oracle Universal Installer (OUI) installs Oracle Clusterware into this directory, also referred to as CRS_HOME.</p> <p>The directory must be owned by the installation owner of Oracle Clusterware (oracle), with the permission set to 755.</p> <p>Follow Oracle Optimal Flexible Architecture (OFA) guidelines while choosing the path.</p>                                                                                                                           |
| Oracle Grid Infrastructure Home Directory (GRID_HOME)<br><br>(For Oracle RAC 11g Release 2)                 | <p>The path to the home directory that stores the Oracle Grid Infrastructure binaries. The Oracle Universal Installer (OUI) installs Oracle Grid Infrastructure and Oracle ASM into this directory, also referred to as GRID_HOME.</p> <p>The directory must be owned by the installation owner of Oracle Grid Infrastructure (oracle or grid), with the permission set to 755.</p> <p>The path to the grid home directory must be the same on all nodes.</p> <p>Follow Oracle Optimal Flexible Architecture (OFA) guidelines while choosing the path.</p> |
| Oracle base directory (ORACLE_BASE)                                                                         | <p>The base directory that contains all the Oracle installations. For Oracle RAC 11g Release 2, create separate Oracle base directories for the grid user and the Oracle user.</p> <p>It is recommended that installations of multiple databases maintain an Optimal Flexible Architecture (OFA) configuration.</p> <p>The path to the Oracle base directory must be the same on all nodes. The permission on the Oracle base directory must be at least 755.</p>                                                                                          |
| Oracle home directory (ORACLE_HOME)                                                                         | <p>The directory in which the Oracle database software is installed. The path to the Oracle home directory must be the same on all nodes. The permission on the Oracle home directory must be at least 755.</p>                                                                                                                                                                                                                                                                                                                                            |

Table 19-4 lists the sample directories for Oracle RAC 10g Release 2/Oracle RAC 11g Release 1 used in the procedures.

Table 19-4 Sample directories for Oracle RAC 10g Release 2/Oracle RAC 11g Release 1

| Directory   | Path                                                                                                                              |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------|
| ORACLE_BASE | /u01/app/oracle                                                                                                                   |
| CRS_HOME    | /u01/app/oracle/product/ <i>release</i> /crshome<br>where <i>release</i> is 10.2.0 or 11.1.0 depending on the Oracle RAC version  |
| ORACLE_HOME | /u01/app/oracle/product/ <i>release</i> /dbhome_1<br>where <i>release</i> is 10.2.0 or 11.1.0 depending on the Oracle RAC version |

Table 19-5 lists the sample directories for Oracle RAC 11g Release 2 used in the procedures.

Table 19-5 Sample directories for Oracle RAC 11g Release 2

| Directory                 | Path                                    |
|---------------------------|-----------------------------------------|
| ORACLE_BASE (Grid user)   | /u01/app/gridbase                       |
| GRID_HOME (Grid user)     | /u01/app/grid/product/11.2.0/gridhome   |
| ORACLE_BASE (Oracle user) | /u01/app/oracle                         |
| ORACLE_HOME (Oracle user) | /u01/app/oracle/product/11.2.0/dbhome_1 |

Use one of the following options to create the directories:

|                     |                                                                                                                                                        |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local file system   | See <a href="#">“To create the directories on the local file system”</a> on page 363.                                                                  |
| Veritas File System | See <a href="#">“To create the file system and directories on the Veritas File system (VxFS)”</a> on page 363.                                         |
| Cluster File System | See <a href="#">“To create the file system and directories on cluster file system for Oracle Grid Infrastructure and Oracle database”</a> on page 368. |

---

**Note:** All procedures use sample values. Replace them with the values that are appropriate for your installation.

---

### To create the directories on the local file system

- 1 Log in as the root user on each node.
- 2 Create the directories and set the ownership and permissions:

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/10.2.0/crshome
# mkdir -p /u01/app/oracle/product/10.2.0/dbhome_1
# chown -R oracle:oinstall /u01/app
# chmod -R 775 /u01/app
```

For Oracle RAC 11g Release 2:

```
# mkdir -p /u01/app/gridbase
# mkdir -p /u01/app/grid/product/11.2.0/gridhome
# chown -R grid:oinstall /u01/app
# chmod -R 775 /u01/app

# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

### To create the file system and directories on the Veritas File system (VxFS)

The sample commands in the procedure are for node galaxy. Repeat the steps on each node of the cluster.

- 1 As the root user, create a VxVM local diskgroup `bindg_hostname` on each node.

For Oracle RAC 11g Release 2, make sure that the size of the disk group is approximately 22 GB; if not, add another disk to the disk group.

For Oracle RAC 10g Release 2 and Oracle RAC 11g Release 1, make sure that the size of the disk group is approximately 12 GB; if not, add another disk to the disk group.

```
# vxdg init bindg_galaxy Disk_1
```

- 2 Create separate volumes for Oracle Clusterware/Oracle Grid Infrastructure (`crsbinvol`) binaries and Oracle binaries (`orabinvol`):

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# vxassist -g bindg_galaxy make crsbinvol 5G
# vxassist -g bindg_galaxy make orabinvol 7G
```

For Oracle RAC 11g Release 2:

```
# vxassist -g bindg_galaxy make crsbinvol 10G
# vxassist -g bindg_galaxy make orabinvol 12G
```

- 3 Create the directories for Oracle RAC.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/10.2.0/crshome
# mkdir -p /u01/app/oracle/product/10.2.0/dbhome_1
```

For Oracle RAC 11g Release 2:

```
# mkdir -p /u01/app/gridbase
# mkdir -p /u01/app/grid/product/11.2.0/gridhome
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
```

- 4 Create file systems with the volumes `crsbinvol` and `orabinvol`:

```
# mkfs -F vxfs /dev/vx/rdisk/bindg_galaxy/crsbinvol
# mkfs -F vxfs /dev/vx/rdisk/bindg_galaxy/orabinvol
```

## 5 Mount the file systems. Perform this step on each node.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mount -F vxfs /dev/vx/dsk/bindg_galaxy/crsbinvol \
/u01/app/oracle/product/10.2.0/crshome
# mount -F vxfs /dev/vx/dsk/bindg_galaxy/orabinvol \
/u01/app/oracle/product/10.2.0/dbhome_1
```

For Oracle RAC 11g Release 2:

```
# mount -F vxfs /dev/vx/dsk/bindg_galaxy/crsbinvol \
/u01/app/grid/product/11.2.0/gridhome
# mount -F vxfs /dev/vx/dsk/bindg_galaxy/orabinvol \
/u01/app/oracle/product/11.2.0/dbhome_1
```

## 6 Set appropriate ownership and permissions for the directories.

---

**Note:** For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1: The ownership and permissions must be changed on all nodes of the cluster because `/u01/app/oracle` must be owned by `oracle:oinstall`, otherwise `/u01/app/oracle/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

For Oracle RAC 11g Release 2: The ownership and permissions must be changed on all nodes of the cluster because `/u01/app` must be owned by `grid:oinstall`, otherwise `/u01/app/oraInventory` is not created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

---

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# chown -R oracle:oinstall /u01/app
# chmod -R 775 /u01/app
```

For Oracle RAC 11g Release 2:

```
# chown -R grid:oinstall /u01/app
# chmod -R 775 /u01/app
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

- 7 Add the resources to the VCS configuration.

See [“To add the storage resources created on VxFS to the VCS configuration”](#) on page 366.

- 8 Repeat all the steps on each node of the cluster.

#### To add the storage resources created on VxFS to the VCS configuration

- 1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

- 2 Configure the VxVM volumes under VCS:

```
# hares -add bin_dg DiskGroup cvm
# hares -modify bin_dg DiskGroup bindg_galaxy -sys galaxy
# hares -modify bin_dg DiskGroup bindg_nebula -sys nebula
# hares -modify bin_dg Enabled 1
```

### 3 Set up the file system under VCS:

```
# hares -add crsbin_mnt Mount cvm
```

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# hares -modify crsbin_mnt MountPoint \  
"/u01/app/oracle/product/10.2.0/crshome"
```

For Oracle RAC 11g Release 2:

```
# hares -modify crsbin_mnt MountPoint \  
"/u01/app/grid/product/11.2.0/gridhome"
```

The following commands are common to all Oracle RAC versions:

```
# hares -modify crsbin_mnt BlockDevice \  
"/dev/vx/dsk/bindg_galaxy/crsbinvol" -sys galaxy  
# hares -modify crsbin_mnt BlockDevice \  
"/dev/vx/dsk/bindg_nebula/crsbinvol" -sys nebula  
# hares -modify crsbin_mnt FSType vxfs  
# hares -modify crsbin_mnt FsckOpt "-n"  
# hares -modify crsbin_mnt Enabled 1  
# hares -add orabin_mnt Mount cvm
```

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# hares -modify orabin_mnt MountPoint \  
"/u01/app/oracle/product/10.2.0/dbhome_1"
```

For Oracle RAC 11g Release 2:

```
# hares -modify orabin_mnt MountPoint \  
"/u01/app/oracle/product/11.2.0/dbhome_1"
```

The following commands are common to all Oracle RAC versions:

```
# hares -modify orabin_mnt BlockDevice \  
"/dev/vx/dsk/bindg_galaxy/orabinvol" -sys galaxy  
# hares -modify orabin_mnt BlockDevice \  
"/dev/vx/dsk/bindg_nebula/orabinvol" -sys nebula  
# hares -modify orabin_mnt FSType vxfs  
# hares -modify orabin_mnt FsckOpt "-n"  
# hares -modify orabin_mnt Enabled 1
```

**4** Link the parent and child resources:

```
# hares -link crsbin_mnt bin_dg
# hares -link orabin_mnt bin_dg
```

**To create the file system and directories on cluster file system for Oracle Grid Infrastructure and Oracle database**

Perform the following steps on the CVM master node in the cluster.

**1** As the root user, create a VxVM shared disk group `bindg`:

```
# vxdg -s init bindg Disk_1
```

**2** Create separate volumes for Oracle Clusterware/Oracle Grid Infrastructure (`crsbinvol`) and Oracle database (`orabinvol`):

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# vxassist -g bindg make crsbinvol 5G
# vxassist -g bindg make orabinvol 7G
```

For Oracle RAC 11g Release 2:

```
# vxassist -g bindg make crsbinvol 10G
# vxassist -g bindg make orabinvol 12G
```



- 3 Create the following directories for Oracle. The file system and directories created on shared storage in this procedure are based on the following layout:

|                                                         |                                                                                                                                            |
|---------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| \$ORACLE_BASE                                           | <i>/u01/app/oracle</i>                                                                                                                     |
|                                                         | <i>/u01/app/oracle</i> is on local storage.                                                                                                |
|                                                         | For Oracle RAC 11g Release 2, create a separate Oracle base directory for the grid user, in addition to the directory for the Oracle user: |
|                                                         | <i>/u01/app/gridbase</i>                                                                                                                   |
|                                                         | <i>/u01/app/gridbase</i> is on local storage.                                                                                              |
| \$CRS_HOME                                              | <i>/u01/app/oracle/product/10.2.0/crshome</i>                                                                                              |
| (For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1) | <i>/u01/app/oracle/product/10.2.0</i> is on local storage.                                                                                 |
|                                                         | <i>/u01/app/oracle/product/10.2.0/crshome</i> is on shared storage.                                                                        |
| \$GRID_HOME                                             | <i>/u01/app/grid/product/11.2.0/gridhome</i>                                                                                               |
| (For Oracle RAC 11g Release 2)                          | <i>/u01/app/grid/product/11.2.0</i> is on local storage.                                                                                   |
|                                                         | <i>/u01/app/grid/product/11.2.0/gridhome</i> is on shared storage.                                                                         |
| \$ORACLE_HOME                                           | <i>/u01/app/oracle/product/release/dbhome_1</i>                                                                                            |
|                                                         | where <i>release</i> is 10.2.0, 11.1.0, or 11.2.0 depending on the Oracle RAC version                                                      |
|                                                         | <i>/u01/app/oracle/product/release</i> is on local storage.                                                                                |
|                                                         | <i>/u01/app/oracle/product/release/dbhome_1</i> is on shared storage.                                                                      |

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/10.2.0/crshome
# mkdir -p /u01/app/oracle/product/10.2.0/dbhome_1
```

For Oracle RAC 11g Release 2:

```
# mkdir -p /u01/app/gridbase
# mkdir -p /u01/app/grid/product/11.2.0/gridhome
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
```

- 4 Create file systems with the volumes `crsbinvol` and `orabinvol`:

```
# mkfs -F vxfs /dev/vx/rdisk/bindg/crsbinvol
# mkfs -F vxfs /dev/vx/rdisk/bindg/orabinvol
```

- 5 Mount the file systems. Perform this step on each node.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/crsbinvol \
/u01/app/oracle/product/10.2.0/crshome
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/orabinvol \
/u01/app/oracle/product/10.2.0/dbhome_1
```

For Oracle RAC 11g Release 2:

```
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/crsbinvol \
/u01/app/grid/product/11.2.0/gridhome
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/orabinvol \
/u01/app/oracle/product/11.2.0/dbhome_1
```

## 6 Change the ownership and permissions on all nodes of the cluster.

---

**Note:** For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1: The ownership and permissions must be changed on all nodes of the cluster because `/u01/app/oracle` must be owned by `oracle:oinstall`, otherwise `/u01/app/oracle/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

For Oracle RAC 11g Release 2: The ownership and permissions must be changed on all nodes of the cluster because `/u01/app` must be owned by `grid:install`, otherwise `/u01/app/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

---

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# chown -R oracle:oinstall /u01/app
# chmod -R 775 /u01/app
```

For Oracle RAC 11g Release 2:

```
# chown -R grid:install /u01/app
# chmod -R 775 /u01/app
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

## 7 Add the CVMVolDg and CFSSMount resources to the VCS configuration.

See [“To add the CFSSMount and CVMVolDg resources to the VCS configuration using CLI”](#) on page 372.

To add the CFSMount and CVMVolDg resources to the VCS configuration using CLI

- 1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

- 2 Configure the CVM volumes under VCS:

```
# hares -add crsorabin_voldg CVMVolDg cvm
```

```
# hares -modify crsorabin_voldg Critical 0
```

```
# hares -modify crsorabin_voldg CVMDiskGroup bindg
```

```
# hares -modify crsorabin_voldg CVMVolume -add crsbinvol
```

```
# hares -modify crsorabin_voldg CVMVolume -add orabinvol
```

```
# hares -modify crsorabin_voldg CVMActivation sw
```

### 3 Set up the file system under VCS:

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# hares -add crsbin_mnt CFSMount cvm
# hares -modify crsbin_mnt Critical 0
# hares -modify crsbin_mnt MountPoint \
"/u01/app/oracle/product/10.2.0/crshome"
# hares -modify crsbin_mnt BlockDevice \
"/dev/vx/dsk/bindg/crsbinvol"
# hares -add orabin_mnt CFSMount cvm
# hares -modify orabin_mnt Critical 0
# hares -modify orabin_mnt MountPoint \
"/u01/app/oracle/product/10.2.0/dbhome_1"
# hares -modify orabin_mnt BlockDevice \
"/dev/vx/dsk/bindg/orabinvol"
```

For Oracle RAC 11g Release 2:

```
# hares -add crsbin_mnt CFSMount cvm
# hares -modify crsbin_mnt Critical 0
# hares -modify crsbin_mnt MountPoint \
"/u01/app/grid/product/11.2.0/gridhome"
# hares -modify crsbin_mnt BlockDevice \
"/dev/vx/dsk/bindg/crsbinvol"
# hares -add orabin_mnt CFSMount cvm
# hares -modify orabin_mnt Critical 0
# hares -modify orabin_mnt MountPoint \
"/u01/app/oracle/product/11.2.0/dbhome_1"
# hares -modify orabin_mnt BlockDevice \
"/dev/vx/dsk/bindg/orabinvol"
```

### 4 Link the parent and child resources:

```
# hares -link crsorabin_voldg cvm_clus

# hares -link crsbin_mnt crsorabin_voldg

# hares -link crsbin_mnt vxfsckd

# hares -link orabin_mnt crsorabin_voldg

# hares -link orabin_mnt vxfsckd
```

**5** Enable the resources:

```
# hares -modify crsorabin_voldg Enabled 1

# hares -modify crsbin_mnt Enabled 1

# hares -modify orabin_mnt Enabled 1

# haconf -dump -makero
```

## 6 Verify the resource configuration in the main.cf file.

The following is a sample resource configuration for Oracle RAC 11g Release 2:

```
CFSMount crsbin_mnt (
    Critical = 0
    MountPoint = "/u01/app/grid/product/11.2.0/gridhome"
    BlockDevice = "/dev/vx/dsk/bindg/crsbinvol"
)

CFSMount orabin_mnt (
    Critical = 0
    MountPoint = "/u01/app/oracle/product/11.2.0/dbhome_1"
    BlockDevice = "/dev/vx/dsk/bindg/orabinvol"
)

CVMVolDg crsorabin_voldg (
    Critical = 0
    CVMDiskGroup = bindg
    CVMVolume = { crsbinvol, orabinvol }
    CVMActivation = sw
)

crsbin_mnt requires crsorabin_voldg
crsbin_mnt requires vxfsckd
orabin_mnt requires crsorabin_voldg
orabin_mnt requires vxfsckd
crsorabin_voldg requires cvm_clus
```

- 7 Verify that the resources are online on all systems in the cluster.

```
# hares -state crsorabin_voldg

# hares -state crsbin_mnt

# hares -state orabin_mnt
```

---

**Note:** At this point, the crsorabin\_voldg resource is reported offline, and the underlying volumes are online. Therefore, you need to manually bring the resource online on each node.

---

To bring the resource online manually:

```
# hares -online crsorabin_voldg -sys galaxy

# hares -online crsorabin_voldg -sys nebula
```

## Setting up user equivalence

You must establish grid user (Oracle RAC 11g Release 2) and Oracle user equivalence on all nodes to allow the Oracle Universal Installer to securely copy files and run programs on the nodes in the cluster without requiring password prompts.

Set up passwordless SSH communication between the cluster nodes for the Oracle user and the grid user.

For more information, see the Oracle documentation.

## Verifying whether the Veritas Membership library is linked to Oracle libraries

The Veritas Membership library (VCSMM) must be linked with Oracle libraries to enable coordinated exchange of cluster membership information and protection of data integrity. Oracle uses the linked skgxn library (libskgxn) to make ioctl calls to VCSMM, which in turn obtains membership information for clusters and instances.



To verify whether the Veritas Membership library is linked to Oracle libraries

- ◆ Verify that the library `/opt/ORCLcluster/lib/libskgxn2.so` is linked to the `/opt/VRTSvcs/rac/lib/libskgxn2_64.so` library:

```
# ls -l /opt/ORCLcluster/lib
libskgxn2.so -> /opt/VRTSvcs/rac/lib/libskgxn2_64.so
```

If the link does not exist, create a symbolic link to the Oracle library as follows:

```
# ln -s /opt/VRTSvcs/rac/lib/libskgxn2_64.so \
/opt/ORCLcluster/lib/libskgxn2.so
```

## Verifying the systems for Oracle RAC installation

Run the SF Oracle RAC installer as described in this section to verify that the cluster configuration settings are conducive for a successful Oracle RAC installation.

Table 19-6 lists the checks performed by the installer on each node. For information on resolving issues that you may encounter during the process, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*, Chapter: Troubleshooting SF Oracle RAC, Section: Troubleshooting installation and configuration check failures.

Table 19-6 List of checks

| Check                                   | Description                                                                                                                                                                                                                                         |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time synchronization                    | The time settings are synchronized across nodes. If the time difference exceeds five seconds, the check fails.                                                                                                                                      |
| System checks                           | All the nodes in the cluster have the same architecture.<br><br>For example, the nodes in the cluster must be either Solaris SPARC systems or Solaris x64 systems.<br><br>All the nodes in the cluster have the same number and type of processors. |
| Operating system and patch level checks | All the nodes in the cluster have the same operating system version and patch levels.                                                                                                                                                               |
| Kernel parameter                        | Verifies that the kernel parameter settings and tunables are in accordance with Oracle RAC recommendations.                                                                                                                                         |

Table 19-6            List of checks (continued)

| Check      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LLT checks | <p>The LLT checks are performed on only two links, the first and the last link specified in the <code>/etc/llttab</code> file.</p> <p>Verifies that the media speed setting for the links are set to full duplex. If the links are not set to full duplex, the check fails.</p> <p>Verifies that the media speed and auto-negotiation settings are the same across all nodes in the cluster.</p> <p>The check fails if the settings vary across the nodes and if auto-negotiation is enabled for the interface on any node in the cluster.</p> <p>Verifies that the payload (MTU) for the jumbo frames is between 1500 and 9000 bytes. Verifies that this setting is the same for all LLT links in the cluster. The check fails if the value is different for different LLT links in the cluster.</p> <p>Verifies that the links are not cross-connected. If the links are cross-connected, the check fails. Make sure that the links do not have a direct path of communication through switch interconnects.</p> <p>Verifies that the cluster has at least two high-priority links that do not belong to the public network.</p> <p>Verifies that the cluster configuration in the <code>/etc/llthosts</code> file is the same on all nodes in the cluster.</p> <p>Verifies that the cluster ID in the <code>/etc/llttab</code> file is the same on all the nodes.</p> <p><b>Note:</b> The following checks are skipped if LLT is configured over UDP:</p> <p>LLT links' full duplex setting</p> <p>LLT link jumbo frame setting (MTU)</p> <p>LLT links' cross connection</p> <p>LLT links' speed and auto negotiation settings</p> |

**Table 19-6** List of checks (*continued*)

| Check              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| I/O fencing checks | <p>Verifies that I/O fencing is enabled on each node in the cluster.</p> <p><b>Note:</b> The verification applies only for disk-based fencing.</p> <p>Verifies that all nodes in the cluster have registered keys with the coordinator disks. If registration keys are missing for any node in the cluster, the check fails for that node.</p> <p>The possible causes for check failure are:</p> <ul style="list-style-type: none"> <li>■ Fencing is not configured on the cluster.</li> <li>■ Fencing is configured in disabled mode.</li> <li>■ Fencing is configured using CPS.</li> </ul> |
| ODM checks         | <p>Verifies that port d is up and running on the nodes.The check fails if the port is not up.</p> <p>The possible causes for check failure are:</p> <ul style="list-style-type: none"> <li>■ ODM is not running in the cluster.</li> <li>■ ODM is running in standalone mode.</li> </ul>                                                                                                                                                                                                                                                                                                      |
| VCSMM checks       | Verifies that port o is up and running on the nodes.The check fails if port o is not up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| GAB checks         | Verifies that the ports a, b, d, f, h, o, v, w are up and running on the nodes. The check fails if any of the SF Oracle RAC components are not running.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| LMX checks         | Verifies that the LMX helper thread is disabled. The check fails if the helper thread is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**To verify the systems for Oracle RAC installation**

- 1** Start the SF Oracle RAC installer:

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

- 2** Enter **2** to select the option **SF Oracle RAC Installation and Configuration Checks**.

```
1) Configure SF Oracle RAC sub-components  
2) SF Oracle RAC Installation and Configuration Checks  
3) Prepare to Install Oracle  
4) Install Oracle Clusterware/Grid Infrastructure and Database  
5) Post Oracle Installation Tasks  
6) Exit SF Oracle RAC Configuration  
Choose option: [1-6,q] (1) 2
```

The installer displays the location of the temporary log file and starts the verification process.

- 3** Enter **n** to return to the menu.

```
Is Oracle installed on the system(s) [y,n,q] (y) n
```

# Installing Oracle RAC

This chapter includes the following topics:

- [About installing Oracle RAC](#)
- [Installing the Oracle Clusterware/Grid Infrastructure software](#)
- [Configuring LLT links in the GPnP profile](#)
- [Installing the Oracle RAC database software](#)
- [Verifying the Oracle Clusterware/Grid Infrastructure and database installation](#)

## About installing Oracle RAC

You can install Oracle RAC on shared storage or locally on each node.

Use one of the following ways to install Oracle RAC:

**SF Oracle RAC installer** The SF Oracle RAC installer starts the installation process and prompts for information that is required by the Oracle Universal Installer. The Oracle Universal Installer launches with these installation values pre-filled and installs Oracle RAC.

You need to invoke the SF Oracle RAC installer to start the installation.

**Oracle Universal Installer**

The Oracle Universal Installer installs Oracle RAC. The installation values must be manually entered at the time of installation.

You need to invoke the Oracle Universal Installer to install Oracle RAC.

Response files      The SF Oracle RAC installer supports silent installation of Oracle RAC using its own response files. You need to modify the SF Oracle RAC response files to include the path information of the Oracle RAC software binaries and the Oracle RAC response files.

For more information and instructions:  
See [“About response files”](#) on page 425.

For instructions, see the chapter "Installing Oracle RAC using a response file" in this document.

Before you start the installation:

- Keep the Oracle worksheets handy as you perform the installation tasks.  
See [“Required installation information for Oracle Clusterware/Grid Infrastructure”](#) on page 661.  
See [“Required installation information for Oracle database”](#) on page 666.
- Review your Oracle installation manuals and the appropriate Oracle support Web sites for additional information required during the installation.

## Installing the Oracle Clusterware/Grid Infrastructure software

This section provides instructions for installing Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC installer. The SF Oracle RAC installer prompts for information required to invoke the Oracle Universal Installer and launches it. The responses provided to the SF Oracle RAC installer are pre-filled in the Oracle Universal Installer wizard. When you step through the installation, review or change these installation values in the Oracle Universal Installer.

---

**Note:** Before you begin the installation, verify that the nodes in the cluster are connected with network links using similar network devices. For example, if you use bge0 as a public link on one node in the cluster, all other nodes in the cluster must also use bge0 as the public link. Similarly, if you use bge1 as a private link on one node in the cluster, all other nodes in the cluster must also use bge1 as the private link.

---

Oracle Clusterware/Grid Infrastructure software is installed on each node in the CRS\_HOME or GRID\_HOME location, depending on the version of Oracle RAC you install.

---

**Note:** If you want to install Oracle Clusterware/Grid Infrastructure on VxFS or CFS, make sure that you created the appropriate storage before proceeding with the installation.

See [“Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually”](#) on page 360.

---

Install the software using one of the following methods:

|                            |                                                                                                                           |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC installer    | See <a href="#">“Installing Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC installer”</a> on page 383.    |
| Oracle Universal Installer | See <a href="#">“Installing Oracle Clusterware/Grid Infrastructure using the Oracle Universal Installer”</a> on page 386. |

## Installing Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC installer

The SF Oracle RAC installer performs the following tasks:

- Invokes the Oracle Universal Installer to install Oracle Clusterware/Grid Infrastructure
- Verifies the Oracle Oracle Clusterware/Grid Infrastructure installation

**To install Oracle Clusterware/Grid Infrastructure using the SF Oracle RAC installer**

- 1** Make sure that you completed the required pre-installation steps.
- 2** Return to the following SF Oracle RAC installer menu and select the option **Install Oracle Clusterware/Grid Infrastructure and Database**.

```
1) Configure SF Oracle RAC sub-components
2) SF Oracle RAC Installation and Configuration Checks
3) Prepare to Install Oracle
4) Install Oracle Clusterware/Grid Infrastructure and Database
5) Post Oracle Installation Tasks
6) Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 4
```

**3 Select the option `Install Oracle Clusterware/Grid Infrastructure`.**

```
1)  Install Oracle Clusterware/Grid Infrastructure
2)  Install Oracle Database
3)  Exit SF Oracle RAC Configuration
b)  Back to previous menu
Choose option: [1-3,b,q] (1) 1
```

**4 Verify that the nodes in the cluster are connected with network links using similar network devices. Review the related information on screen and press `Enter` to confirm.**

```
Do you want to continue? [y,n,q] (y)
```

**5 Set the `DISPLAY` environment variable that is required for the Oracle Universal Installer:**

```
Enter DISPLAY environment variable: [b] 10.20.12.150:0.0
```

where 10.20.12.150 is the IP address of X client where you want to export the display for the installer.

**6 Enter the Oracle UNIX user name. The Oracle UNIX user name was previously set up during the pre-installation process.**

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
Enter Oracle UNIX user name:  [b] (oracle)
```

For Oracle RAC 11g Release 2:

```
Enter Oracle UNIX user name:  [b] (grid)
```

**7 Enter Oracle UNIX group name. The Oracle UNIX group name was previously set up during the pre-installation process.**

```
Enter Oracle UNIX group name:  [b] (oinstall)
```

**8 Enter the full path of the Oracle base directory.**

---

**Note:** The `ORACLE_BASE` directory must be a local directory.

---



- 9 Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory.  
  
If the Oracle Clusterware/Grid Infrastructure home directory you specified does not exist, the installer creates the directory locally on each node and sets appropriate permissions for the Oracle user.
- 10 Enter the full path of the Oracle Clusterware/Grid Infrastructure installation image. Press **Return** to proceed.  
  
The installer detects the version of the Oracle Clusterware/Grid Infrastructure software.
- 11 Enter **y** to continue with the installation.
- 12 Review and confirm the configuration information. The installer invokes the Oracle Universal Installer:
- 13 Enter the required information when prompted by the Oracle Universal Installer.  
  
See [“Required installation information for Oracle Clusterware/Grid Infrastructure”](#) on page 661.
- 14 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle Clusterware/Grid Infrastructure installation.
- 15 At the end of the Oracle Clusterware/Grid Infrastructure installation, run the following configuration scripts as the root user from each node of the cluster, in the listed order.
  - `orainstRoot.sh` (located in the `oraInventory` directory)  
Make sure the script exists on each node before proceeding.
  - `root.sh` (located in the `CRS_HOME` or `GRID_HOME` directory, depending on your Oracle RAC version)

Do not run the scripts simultaneously on the cluster nodes.

---

**Note:** For Oracle RAC 10g Release 2: If `vipca` fails to run silently, run the script manually on one of the nodes as the root user.

```
# export DISPLAY=10.20.12.150:0.0
# cd $CRS_HOME/bin
# ./vipca
```

---

---

**Note:** For Oracle RAC 11g Release 2: The Oracle Cluster Verification Utility fails during the installation of the Oracle Grid Infrastructure software. You may ignore this error. This is a known issue with Oracle.

---

- 16 Return to the Oracle Universal Installer window and click "OK" to continue. The Oracle Universal Installer informs you that the Oracle Clusterware/Grid Infrastructure installation was successful.
- 17 Return to the SF Oracle RAC installer and press Return to proceed. The installer verifies whether Oracle Clusterware/Grid Infrastructure is installed properly.  
This completes the Oracle Clusterware/Grid Infrastructure installation.

## Installing Oracle Clusterware/Grid Infrastructure using the Oracle Universal Installer

This section provides instructions for installing the Oracle Clusterware/Grid Infrastructure software using the Oracle Universal Installer. The software is installed on each node in the Oracle Clusterware/Grid Infrastructure home directory.

### To install Oracle Clusterware/Grid Infrastructure using the Oracle Universal Installer

- 1 Log in as the Oracle grid user (Oracle RAC 11g Release 2) or the Oracle user. On the first node, set the DISPLAY variable.
  - For Bourne Shell (bash), type:  

```
$ DISPLAY=10.20.12.150:0.0;export DISPLAY
```

where 10.20.12.150 is the IP address of X client where you want to export the display for the installer.
  - For C Shell (csh or tcsh), type:  

```
$ setenv DISPLAY 10.20.12.150:0.0
```

where 10.20.12.150 is the IP address of X client where you want to export the display for the installer.
- 2 Start the Oracle Universal Installer on the first node.  

```
$ cd /dvd_mount
```

```
$ ./runInstaller
```

- 3 Enter the required information when prompted by the Oracle Universal Installer.  
See [“Required installation information for Oracle Clusterware/Grid Infrastructure”](#) on page 661.
- 4 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle Clusterware/Grid Infrastructure installation.

---

**Note:** For Oracle RAC 11g Release 2: If you want to save the Oracle Grid Infrastructure installation configuration into a response file for future installations, click the **Save Response File** option on the Summary page of the Oracle Universal Installer.

---

- 5 Run the `orainstRoot.sh` script as prompted by the Oracle Universal Installer.
- 6 Run the `root.sh` script on each node as prompted by the Oracle Universal Installer:

The Oracle Clusterware daemons are started on the node.

---

**Note:** The Oracle Cluster Verification Utility fails during the installation of the Oracle Clusterware/Grid Infrastructure software. You may ignore this error. This is a known issue with Oracle.

---

## Configuring LLT links in the GPnP profile

Perform this step only for Oracle RAC 11g Release 2 installations.

Update the GPnP profile to include the remaining LLT links that were not added to the profile during the Oracle Grid Infrastructure installation.

### To configure the LLT links in the GPnP profile

- 1 View the currently configured interfaces:

```
# $GRID_HOME/bin/oifcfg getif
bge0 10.2.156.0          global      public
bge1 192.168.12.0       global      cluster_interconnect
```

The interfaces that are currently stored in the GPnP profile, their subnets, and their role (public or cluster\_interconnect) are displayed.

- 2 Add the remaining LLT links to the GPnP profile:

```
# $GRID_HOME/bin/oifcfg setif -global \
bge2/192.168.12.0:cluster_interconnect
```

If you are using multiple IP addresses on different subnet for cluster interconnect (for load balancing), add the remaining interface subnets to the GPnP profile.

```
# $GRID_HOME/bin/oifcfg setif -global \
bge2/192.168.2.0:cluster_interconnect
# $GRID_HOME/bin/oifcfg setif -global \
bge1/192.168.2.0:cluster_interconnect
```

- 3 Verify that the correct interface subnet is in use:

```
# $GRID_HOME/bin/oifcfg getif
bge0 10.2.156.0          global      public
bge1 192.168.12.0       global      cluster_interconnect
bge2 192.168.12.0       global      cluster_interconnect
bge1 192.168.2.0        global      cluster_interconnect
bge2 192.168.2.0        global      cluster_interconnect
```

Make sure all the LLT links are configured and listed in the GPnP profile.

## Installing the Oracle RAC database software

Before you start the installation of Oracle database, make sure that Oracle Clusterware/Grid Infrastructure is up and running. Symantec recommends you to install the Oracle database locally on each node.

---

**Note:** If you want to install Oracle database on VxFS or CFS, make sure that you created the appropriate storage before proceeding with the installation.

See [“Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories manually”](#) on page 360.

---

Install the software using one of the following methods:

|                            |                                                                                                            |
|----------------------------|------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC installer    | See <a href="#">“Installing the Oracle RAC database using the SF Oracle RAC installer”</a> on page 389.    |
| Oracle Universal Installer | See <a href="#">“Installing the Oracle RAC database using the Oracle Universal Installer”</a> on page 391. |

## Installing the Oracle RAC database using the SF Oracle RAC installer

The SF Oracle RAC installer performs the following tasks:

- Verifies the status of Oracle Clusterware/Grid Infrastructure on all nodes
- Invokes the Oracle Universal Installer to install the Oracle database
- Verifies the Oracle RAC database installation

### To install the Oracle RAC database

- 1** Return to the SF Oracle RAC installer and type **4** to select the option **Install Oracle Clusterware/Grid Infrastructure and Database**.

```
1)  Configure SF Oracle RAC sub-components
2)  SF Oracle RAC Installation and Configuration Checks
3)  Prepare to Install Oracle
4)  Install Oracle Clusterware/Grid Infrastructure and Database
5)  Post Oracle Installation Tasks
6)  Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 4
```

- 2** Select the option **Install Oracle Database**.

```
1)  Install Oracle Clusterware/Grid Infrastructure
2)  Install Oracle Database
3)  Exit SF Oracle RAC Configuration
b)  Back to previous menu
Choose option: [1-3,b,q] (1) 2
```

- 3** Set the DISPLAY environment variable that is required for the Oracle Universal Installer.

Enter the DISPLAY environment variable: [b] 10.20.12.150:0.0

- 4** Enter Oracle UNIX user name. The Oracle UNIX user name was previously set up during the pre-installation process.

Enter Oracle UNIX user name: [b] (oracle)

- 5** Enter Oracle UNIX group name. The Oracle UNIX group name was previously set up during the pre-installation process.

Enter Oracle UNIX group name: [b] (oinstall)

- 6** Enter the full path of the Oracle base directory.

Enter absolute path of Oracle Base directory: [b]

- 7** Enter the full path of the Oracle Clusterware/Grid Infrastructure home directory.

Enter absolute path of Oracle Clusterware/Grid Infrastructure Home directory: [b]

- 8** Enter the full path of the Oracle database home directory.

Enter absolute path of Oracle Database Home directory: [b]

If the Oracle RAC database home directory you specified does not exist, the installer creates the directory locally on each node and sets appropriate permissions for the Oracle user.

- 9** Enter the full path of the database installation image.

Enter absolute path of Oracle Database install image: [b]

The installer determines the version of the Oracle software from the binaries.

- 10** Enter **y** to proceed with the installation.

- 11** Review and confirm the configuration information.

The installer verifies that Oracle Clusterware/Grid Infrastructure is running and invokes the Oracle Universal Installer:

- 12 Enter the required information when prompted by the Oracle Universal Installer.  
See [“Required installation information for Oracle database”](#) on page 666.
- 13 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle database installation.
- 14 Run the root.sh script as the root user on the cluster nodes:  
Return to the Oracle Universal Installer window and click "OK" to continue.
- 15 Return to the SF Oracle RAC installer and press **Return** to continue. The installer verifies the Oracle database installation.

---

**Note:** Symantec recommends you to relink the SF Oracle RAC libraries only after completing all the required patch additions.

---

## Installing the Oracle RAC database using the Oracle Universal Installer

The following procedure describes how to install the Oracle RAC database using the Oracle Universal Installer. Symantec recommends that you install the Oracle RAC database locally on each node.

### To install Oracle RAC database using the Oracle Universal Installer

- 1 Log in as the Oracle user. On the first node, set the DISPLAY variable.
  - For Bourne Shell (bash), type:  

```
$ DISPLAY=10.20.12.150:0.0;export DISPLAY
```
  - For C Shell (csh or tcsh), type:  

```
$ setenv DISPLAY 10.20.12.150:0.0
```
- 2 Start the Oracle Universal Installer.  

```
$ cd /dvd_mount
```

```
$ ./runInstaller
```
- 3 Enter the required information when prompted by the Oracle Universal Installer.  
See [“Required installation information for Oracle database”](#) on page 666.

- 4
- Review the configuration summary presented by the Oracle Universal Installer.

**Note:** For Oracle RAC 11g Release 2: If you want to save the Oracle RAC database installation configuration into a response file for future installations, click the **Save Response File** option on the Summary page of the Oracle Universal Installer.

The Oracle Universal Installer begins the Oracle database installation.

- 5
- Run the root.sh script as prompted by the Oracle Universal Installer.

```
# cd $ORACLE_HOME

# ./root.sh
```

Verifying the Oracle Clusterware/Grid Infrastructure and database installation

The following procedure verifies the Oracle Clusterware/Grid Infrastructure and Oracle RAC database installation by verifying that the Oracle processes are running on all nodes.

To verify the installation, run the following command from any node in the cluster. Verify in the command output that the Oracle Clusterware/Grid Infrastructure processes are online on the nodes.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# $CRS_HOME/bin/crs_stat
```

| Name           | Type        | Target | State  | Host   |
|----------------|-------------|--------|--------|--------|
| ora.galaxy.vip | application | ONLINE | ONLINE | galaxy |
| ora.galaxy.gsd | application | ONLINE | ONLINE | galaxy |
| ora.galaxy.ons | application | ONLINE | ONLINE | galaxy |
| ora.nebula.vip | application | ONLINE | ONLINE | nebula |
| ora.nebula.gsd | application | ONLINE | ONLINE | nebula |
| ora.nebula.ons | application | ONLINE | ONLINE | nebula |

For Oracle RAC 11g Release 2:

```
# $GRID_HOME/bin/crsctl stat res -t
```



| NAME                    | TARGET  | STATE   | SERVER | STATE_DETAILS |
|-------------------------|---------|---------|--------|---------------|
| Local Resources         |         |         |        |               |
| ora.LISTENER.lsnr       |         |         |        |               |
|                         | ONLINE  | ONLINE  | galaxy |               |
|                         | ONLINE  | ONLINE  | nebula |               |
| ora.asm                 |         |         |        |               |
|                         | OFFLINE | OFFLINE | galaxy |               |
|                         | OFFLINE | OFFLINE | nebula |               |
| .                       |         |         |        |               |
| .                       |         |         |        |               |
| .                       |         |         |        |               |
| Cluster Resources       |         |         |        |               |
| ora.LISTENER_SCAN1.lsnr |         |         |        |               |
| 1                       | ONLINE  | ONLINE  | nebula |               |
| ora.LISTENER_SCAN2.lsnr |         |         |        |               |
| 1                       | ONLINE  | ONLINE  | galaxy |               |
| ora.LISTENER_SCAN3.lsnr |         |         |        |               |
| 1                       | ONLINE  | ONLINE  | galaxy |               |
| .                       |         |         |        |               |
| .                       |         |         |        |               |
| .                       |         |         |        |               |

To verify the Oracle RAC database installation, check the oraInventory logs located in the ORACLE\_BASE directory.

## Node numbering discrepancies in Oracle RAC 11g Release 2

For Oracle RAC 11g Release 2, you will observe that the output of the `lsnodes` command (used to obtain membership information from SF Oracle RAC) and the `olsnodes` command (used to obtain membership information from Oracle Grid Infrastructure) differ. Oracle RAC starts the count at 1 while SF Oracle RAC starts the count at 0.

For example, in a two-node cluster:

The `lsnodes` command produces the following output:

```
# lsnodes -n
Node1      0
Node2      1
```

The `olsnodes` command produces the following output:

```
# olsnodes -n
Node1      1
Node2      2
```

In previous Oracle RAC releases (starting with Oracle RAC 10g Release 1), the commands produced the same output unless there were configuration issues in the cluster. Whenever node numbering discrepancies were observed, further issues were seen in the cluster.

However, the discrepancy in node numbering in Oracle RAC 11g Release 2 is conformant with the design change in Oracle Grid Infrastructure. It does not indicate any configuration issues in the cluster.

# Performing Oracle RAC post-installation tasks

This chapter includes the following topics:

- Adding Oracle RAC patches or patchsets
- Configuring the CSSD resource
- Preventing automatic startup of Oracle Clusterware/Grid Infrastructure
- Relinking the SF Oracle RAC libraries with Oracle RAC
- Creating the Oracle RAC database
- Adding Oracle UDP IPC private IP addresses to the Oracle initialization parameter file
- Preventing automatic database startup
- Configuring VCS service groups for Oracle RAC
- Removing permissions for communication
- Configuring the SFDB repository database after installation

## Adding Oracle RAC patches or patchsets

To install the required patches or patchsets, review the notes that accompany the patch or patchset.

Before installing any Oracle RAC patch or patchset software:

- Review the latest information on supported Oracle RAC patches and patchsets:  
<http://www.symantec.com/docs/TECH44807>

- You must have installed the base version of the Oracle RAC software.

## Configuring the CSSD resource

You must configure the CSSD resource to ensure that the CSSD dependencies on the resources that manage OCR and voting disk and the private IP address are satisfied before Oracle Clusterware/Grid Infrastructure starts.

---

**Note:** It is mandatory to use CSSD agent in SF Oracle RAC installations. Using the CSSD agent ensures adequate handling of inter-dependencies, thus preventing the premature startup of Oracle Clusterware/Grid Infrastructure, which causes cluster failures.

---

Use one of the following ways to configure the CSSD resource:

|                         |                                                                                                    |
|-------------------------|----------------------------------------------------------------------------------------------------|
| SF Oracle RAC installer | See <a href="#">“Configuring the CSSD resource using the SF Oracle RAC installer”</a> on page 396. |
| Manual                  | See <a href="#">“Configuring the CSSD resource manually”</a> on page 398.                          |

## Configuring the CSSD resource using the SF Oracle RAC installer

Configure the CSSD resource using the SF Oracle RAC installer if the OCR and voting disk storage is configured on CVM raw volumes or CFS.

---

**Note:** If the OCR and voting disk storage is configured on ASM disk groups, configure the CSSD resource manually.

---

Before you configure the CSSD resource, make sure that the following requirements are satisfied:

1. Oracle Clusterware/Grid Infrastructure is up and running.
2. OCR and voting disk is configured on CVM raw volumes or CFS and managed by VCS.
3. The private IP address for Oracle Clusterware/Grid Infrastructure is configured under the PrivNIC or MultiPrivNIC resource in the same VCS group as that of OCR and voting disk.

**Note:** The SF Oracle RAC installer configures the OCR, voting disk, and PrivNIC/MultiPrivNIC resources in the cvm group. If you configured one of these resources manually, make sure that these resources are placed in the cvm group. If the resources are not in the same group, configure the the CSSD resource manually.

See [“Configuring the CSSD resource manually”](#) on page 398.

The installer performs the following configuration tasks:


- Adds the CSSD resource to the VCS configuration in the cvm group.

**Note:** If the CSSD resource already exists, the installer enables reconfiguration of the resource by deleting the existing resource.

- Sets the dependency of the CSSD resource on the PrivNIC or MultiPrivNIC resource that manages the private IP address for Oracle Clusterware/Grid Infrastructure.
- Sets the dependency of the CSSD resource on the CFSMount or CVMVolDg resources that manage OCR and voting disk.
- Enables the CSSD resource and saves the new configuration.

[Figure 21-1](#) illustrates the configuration performed by the SF Oracle RAC installer. In the figure, the CSSD resource is configured under the CVM group.

**Figure 21-1** CSSD configuration by SF Oracle RAC installer



### To configure the CSSD resource

- 1 Start the SF Oracle RAC installer, if it is not already running. Select the option **Post Oracle Installation Tasks**.

```
1)  Configure SF Oracle RAC sub-components
2)  SF Oracle RAC Installation and Configuration Checks
3)  Prepare to Install Oracle
4)  Install Oracle Clusterware/Grid Infrastructure and Database
5)  Post Oracle Installation Tasks
6)  Exit SF Oracle RAC Configuration
Choose option: [1-6,q] (1) 5
```

- 2 Select the option **Configure CSSD agent**.

The installer verifies that Oracle Clusterware/Grid Infrastructure is running on all the nodes.

- 3 Press **Return** to continue with the configuration. The installer reads the resource and group mappings for the CSSD agent from the VCS configuration file and displays the information.
- 4 Enter **y** to continue with the configuration. Review the messages as the installer configures the CSSD agent and sets the appropriate dependencies.
- 5 Press **Return** to return to the installer menu.
- 6 If the Oracle Clusterware/Grid Infrastructure and the Oracle database binaries are on CFS, set the dependencies between the cssd resource and the CFMount resources for the binaries manually:

```
# hares -link cssd crsbin_mnt
# hares -link cssd orabin_mnt
```

## Configuring the CSSD resource manually

Add the `cssd` resource to the VCS configuration and set CSSD dependencies on the resources that manage OCR and voting disk and the private IP addresses for Oracle Clusterware/Grid Infrastructure.

---

**Note:** It is recommended that the OCR, voting disk, and PrivNIC/MultiPrivNIC resources be configured in the same VCS group as that of the `cssd` resource. If the resources are not in the same group, set the appropriate dependencies between the service groups.

---

### To configure the CSSD resource

- 1 Change the permission on the VCS configuration file to read-write mode:

```
# haconf -makerw
```

- 2 Add the CSSD resource to the cvm group:

```
# hares -add cssd Application cvm
```

- 3 Modify the CSSD resource attributes:

```
# hares -modify cssd StartProgram /opt/VRTSvcs/rac/bin/cssd-online
# hares -modify cssd StopProgram /opt/VRTSvcs/rac/bin/cssd-offline
# hares -modify cssd MonitorProgram /opt/VRTSvcs/rac/bin/cssd-monitor
# hares -modify cssd CleanProgram /opt/VRTSvcs/rac/bin/cssd-clean
# hares -modify cssd Critical 0
# hares -override cssd OnlineWaitLimit
# hares -modify cssd OnlineWaitLimit 5
```

- 4 Enable the CSSD resource:

```
# hares -modify cssd Enabled 1
```

- 5 Set the dependency of the CSSD resource on the CFSMount or CVMVolDg resources that manage OCR and voting disk.

If you configured OCR and voting disk on CVM raw volumes:

```
# hares -link cssd ocr_voldg_ocrvotedg
# hares -link cssd vote_voldg_ocrvotedg
```

If you configured OCR and voting disk on CFS:

```
# hares -link cssd ocrvote_mnt_ocrvotedg
```

- 6 Set the dependency of the CSSD resource on the PrivNIC or MultiPrivNIC resources that manage the private IP address for Oracle Clusterware/Grid Infrastructure.

If you configured the PrivNIC or MultiPrivNIC resource:

```
# hares -link cssd resource_name
```

where *resource\_name* is the name of the PrivNIC or MultiPrivNIC resource.

- 7 If the Oracle Clusterware/Grid Infrastructure and the Oracle database binaries are on CFS, set the dependencies between the CSSD resource and the CFSMount resources for the binaries manually:

```
# hares -link cssd crsbin_mnt
# hares -link cssd orabin_mnt
```

- 8 Change the permission on the VCS configuration file to read-only mode:

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

## Preventing automatic startup of Oracle Clusterware/Grid Infrastructure

The use of the CSSD agent is mandatory to ensure adequate handling of service group inter-dependencies and thereby prevent the premature startup of Oracle Clusterware/Grid Infrastructure. Therefore, disable automatic startup of Oracle Clusterware/Grid Infrastructure when the system starts.

### To prevent automatic startup of Oracle Clusterware/Grid Infrastructure

- 1 Log in as the root user on each node in the cluster.
- 2 Disable automatic startup of Oracle Clusterware/Grid Infrastructure.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# $CRS_HOME/bin/crsctl disable crs
```

For Oracle RAC 11g Release 2:

```
# $GRID_HOME/bin/crsctl disable crs
```

## Relinking the SF Oracle RAC libraries with Oracle RAC

If you added or upgraded the Oracle patches, you must relink the SF Oracle RAC libraries to Oracle. If you are using Oracle RAC 10g, you need to link Oracle RAC with VCSIPC, VCSMM, and ODM libraries. If you are using Oracle RAC 11g, you need to link Oracle RAC with VCSMM and ODM libraries. Relinking the libraries enables coordinated exchange of cluster membership information and protection of data.



---

**Note:** Symantec recommends that you relink the SF Oracle RAC libraries only after completing all the required patch additions.

---

Use one of the following ways to relink the libraries:

SF Oracle RAC installer See [“Relinking the SF Oracle RAC libraries with Oracle RAC using the SF Oracle RAC installer”](#) on page 401.

Manual See [“Relinking SF Oracle RAC libraries with Oracle RAC manually”](#) on page 402.

## Relinking the SF Oracle RAC libraries with Oracle RAC using the SF Oracle RAC installer

Perform the steps in the following procedure to relink the libraries using the SF Oracle RAC installer.

### To relink the SF Oracle RAC libraries with Oracle RAC

**1** Return to the SF Oracle RAC installer menu, and select the option **Post Oracle Installation**.

- 1) Configure SF Oracle RAC sub-components
  - 2) SF Oracle RAC Installation and Configuration Checks
  - 3) Prepare to Install Oracle
  - 4) Install Oracle Clusterware/Grid Infrastructure and Database
  - 5) Post Oracle Installation Tasks
  - 6) Exit SF Oracle RAC Configuration
- Choose option: [1-6,q] (1) **5**

**2** Select the option **Relink Oracle Database Binary**.

- 3 Provide the Oracle environment information—user name, group name, CRS\_HOME (or GRID\_HOME), ORACLE\_HOME. Based on this information, the installer detects the version of Oracle installed.

---

**Note:** You need to provide the Oracle environment information only if you quit the installer after installing Oracle Clusterware/Grid Infrastructure and the Oracle database.

---

```
Enter Oracle UNIX user name: [b] oracle
Enter Oracle UNIX group name: [b] (oinstall)
Enter absolute path of Oracle Clusterware/Grid Infrastructure
Home directory: [b]
Enter absolute path of Oracle Database Home directory: [b]
.
.
Do you want to continue? [y,n,q] (y)
```

- 4 Review and confirm the Oracle database information.  
The installer starts relinking the libraries.
- 5 After the relinking completes, confirm that the correct ODM library is used.

```
$ ldd $ORACLE_HOME/bin/oracle | grep odm
```

If the library is not linked correctly, no output is displayed. Sometimes, the Oracle ODM library may not be correctly linked with the Veritas ODM library because of the presence of a static ODM library at \$ORACLE\_HOME/rdbms/lib/libodm.a. This is a known issue with Oracle.

To resolve this issue, run the following commands:

```
$ cd $ORACLE_HOME/rdbms/lib
$ mv libodm10.a libodm10.a.backup
$ /usr/ccs/bin/make -f ins_rdbms.mk ioracle
```

Verify that the correct ODM library is used:

```
$ ldd $ORACLE_HOME/bin/oracle | grep odm
libodm10.so => /app/oracle/orahome/lib/libodm10.so
```

For more information, see the Oracle metalink document: 725903.1

## Relinking SF Oracle RAC libraries with Oracle RAC manually

The relinking process involves the following tasks:

- Copying VCS IPC libraries for an Oracle RAC 10g installation
- Linking the ODM library

The Oracle RAC and Veritas library locations for linking the Oracle libraries with SF Oracle RAC are listed in the following tables. You may use these tables as reference when you relink the libraries as described in this section.

**Note:** VCS IPC is not applicable to Oracle RAC 11g.

Table 21-1 lists the Oracle RAC and Veritas library locations for Oracle RAC 10g Release 2.

**Table 21-1** Oracle RAC and Veritas library locations for VCS IPC and ODM - Oracle RAC 10g Release 2

| Libraries | Oracle RAC library                                                                                                                             | Veritas library                                                                                                                        |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| VCSMM     | <code>\$CRS_HOME/lib/libskgxn2.so</code>                                                                                                       | <code>/opt/ORCLcluster/lib/libskgxn2.so</code>                                                                                         |
| VCS IPC   | <code>\$ORACLE_HOME/lib/libskgxp10.so</code>                                                                                                   | <code>/opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.so</code>                                                                               |
| ODM       | For Solaris SPARC:<br><br><code>\$ORACLE_HOME/lib/libodm10.so</code><br><br>For Solaris x86:<br><br><code>\$ORACLE_HOME/lib/libodm10.so</code> | For Solaris SPARC:<br><br><code>/usr/lib/sparcv9/libodm.so</code><br><br>For Solaris x86:<br><br><code>/usr/lib/amd64/libodm.so</code> |

Table 21-2 lists the Oracle RAC and Veritas library locations for Oracle RAC 11g Release 1.

**Table 21-2** Oracle RAC and Veritas library locations for VCSMM and ODM - Oracle RAC 11g Release 1

| Libraries | Oracle RAC library                       | Veritas library                                |
|-----------|------------------------------------------|------------------------------------------------|
| VCSMM     | <code>\$CRS_HOME/lib/libskgxn2.so</code> | <code>/opt/ORCLcluster/lib/libskgxn2.so</code> |

**Table 21-2** Oracle RAC and Veritas library locations for VCSMM and ODM - Oracle RAC 11g Release 1 *(continued)*

| Libraries | Oracle RAC library                                                                                                   | Veritas library                                                                                                      |
|-----------|----------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| ODM       | For Solaris SPARC:<br><br>\$ORACLE_HOME/lib/libodm11.so<br><br>For Solaris x86:<br><br>\$ORACLE_HOME/lib/libodm11.so | For Solaris SPARC:<br><br>/usr/lib/sparcv9/<br>libodm.so<br><br>For Solaris x86:<br><br>/usr/lib/amd64/<br>libodm.so |

Table 21-3 lists the Oracle RAC and Veritas library locations for Oracle RAC 11g Release 2.

**Table 21-3** Oracle RAC and Veritas library locations for VCSMM and ODM - Oracle RAC 11g Release 2

| Libraries | Oracle RAC library                                                                                                   | Veritas library                                                                                                      |
|-----------|----------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| VCSMM     | \$GRID_HOME/lib/libskgxn2.so                                                                                         | /opt/ORCLcluster/lib/<br>libskgxn2.so                                                                                |
| ODM       | For Solaris SPARC:<br><br>\$ORACLE_HOME/lib/libodm11.so<br><br>For Solaris x86:<br><br>\$ORACLE_HOME/lib/libodm11.so | For Solaris SPARC:<br><br>/usr/lib/sparcv9/<br>libodm.so<br><br>For Solaris x86:<br><br>/usr/lib/amd64/<br>libodm.so |

**Copying VCS IPC libraries for an Oracle RAC 10g installation**

Perform this step only for Oracle RAC 10g installations and if you are using VCS IPC.

Perform the steps in the procedure on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, copy the libraries on one node only. Use the mount command to check that the file system containing the Oracle libraries are mounted.

### To copy VCS IPC libraries

- 1 Log in as the Oracle user.
- 2 Change to the `$ORACLE_HOME/lib` directory:

```
$ cd $ORACLE_HOME/lib
```

- 3 Back up Oracle's `libskgxp10` library:

```
$ mv libskgxp10.so libskgxp10.so.oracle.`date +%m_%d_%Y-%H_%M_%S`
```

- 4 Replace the file `$ORACLE_HOME/lib/libskgxp10.so` with the Veritas VCS IPC library:

```
$ cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.so libskgxp10.so
```

## Linking the ODM library

Perform the steps in the procedure on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, copy the libraries on one node only. Use the mount command to check that the file system containing the Oracle libraries are mounted.

### To link the Veritas ODM library

- 1 Log in as the Oracle user.
- 2 Change to the `$ORACLE_HOME/lib` directory:

```
$ cd $ORACLE_HOME/lib
```

- 3 Back up Oracle's ODM library:

For Oracle RAC 10g:

```
$ mv libodm10.so libodm10.so.oracle-`date +%m_%d_%Y-%H_%M_%S`
```

For Oracle RAC 11g:

```
$ mv libodm11.so libodm11.so.oracle-`date +%m_%d_%Y-%H_%M_%S`
```

**4** Link the Veritas ODM library with Oracle's `libodm` library:

For Oracle RAC 10g:

For Solaris SPARC:

```
$ cp /usr/lib/sparcv9/libodm.so libodm10.so
```

For Solaris x64:

```
$ cp /usr/lib/amd64/libodm.so libodm10.so
```

For Oracle RAC 11g:

For Solaris SPARC:

```
$ cp /usr/lib/sparcv9/libodm.so libodm11.so
```

For Solaris x64:

```
$ cp /usr/lib/amd64/libodm.so libodm11.so
```

**5** After the relinking completes, confirm that the correct ODM library is used.

```
$ ldd $ORACLE_HOME/bin/oracle | grep odm
```

If the library is not linked correctly, no output is displayed. Sometimes, the Oracle ODM library may not be correctly linked with the Veritas ODM library because of the presence of a static ODM library at

`$ORACLE_HOME/rdbms/lib/libodm.a`. This is a known issue with Oracle.

To resolve this issue, run the following commands:

```
$ cd $ORACLE_HOME/rdbms/lib
$ mv libodm10.a libodm10.a.backup
$ /usr/ccs/bin/make -f ins_rdbms.mk ioracle
```

Verify that the correct ODM library is used:

```
$ ldd $ORACLE_HOME/bin/oracle | grep odm
libodm10.so => /app/oracle/orahome/lib/libodm10.so
```

For more information, see the Oracle metalink document: 725903.1

## Creating the Oracle RAC database

Create the Oracle RAC database on shared raw VxVM volumes or shared VxFS file systems. Use your own tools or scripts or review the guidelines on using the Oracle DBCA (Database Creation Assistant) tool to create the database.

For instructions on creating an Oracle RAC database:

See [“About creating a test database”](#) on page 719.

For more information, see the Oracle RAC documentation.

---

**Note:** If you plan to configure global clusters, then set up the Oracle RAC database only on the primary site. On the secondary site, the database will be replicated.

---

---

**Note:** For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1: Make sure that you configure the database for availability if the database cache fusion traffic is configured to use Oracle UDP IPC private IP addresses and if these addresses are configured as a PrivNIC or MultiPrivNIC resource.

See [“Adding Oracle UDP IPC private IP addresses to the Oracle initialization parameter file”](#) on page 407.

---

## Adding Oracle UDP IPC private IP addresses to the Oracle initialization parameter file

Perform this step only for Oracle RAC 10 Release 2/Oracle RAC 11 Release 1 installations.

Add the Oracle UDP IPC private IP addresses to the Oracle initialization parameter file (for example, `pfile` or `spfile`) if the database cache fusion traffic is configured to use Oracle UDP IPC private IP addresses and if these addresses are configured as a PrivNIC or MultiPrivNIC resource for high availability. You need to configure the `cluster_interconnects` parameter in the Oracle initialization parameter file to use the appropriate private IP addresses.

**To add Oracle UDP IPC private IP addresses to the initialization parameter file**

- 1 Log in as the Oracle user on one of the nodes in the cluster.
- 2 Set the `cluster_interconnects` parameter in the Oracle initialization parameter file to use the Oracle UDP IPC private IP addresses that are used for database cache fusion and are managed by PrivNIC or MultiPrivNIC.

---

**Note:** Configure multiple IP addresses for each database instance in the `cluster_interconnects` parameter to distribute the database cache fusion traffic across multiple links.

---

```
$ sqlplus '/ as sysdba'
```

```
SQL> alter system set cluster_interconnects='192.168.2.1:192.168.3.1' \
scope=spfile sid='vrts1';
SQL> alter system set cluster_interconnects='192.168.2.2:192.168.3.2' \
scope=spfile sid='vrts2';
SQL> exit;
```

- 3 Stop the Oracle database on all nodes:

If the database is configured under VCS:

```
# hares -offline oracle_resource -sys system_name
```

If the database is not configured under VCS:

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

- 4 For Oracle RAC 10g: Unlink the Veritas VCS IPC library.

Perform this step on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, perform the step on one node only.

```
$ cd $ORACLE_HOME/lib
```

```
$ cp libskgxp10.so libskgxp10.so
```



- 5 As the Oracle user on one of the nodes in the cluster, restart the Oracle database:

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

- 6 As the root user on each node in the cluster, verify that the private IP addresses are running.

```
# ifconfig bge1
```

## Preventing automatic database startup

Configure the Oracle RAC database for manual startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent. Before you configure the VCS service groups for Oracle, you need to prevent the Oracle database from starting automatically. The Oracle Clusterware/Grid Infrastructure and Oracle agent may attempt to start the database instance at the same time if the database mount is available. To prevent the Oracle database from starting automatically, you must change the management policy for the database from automatic to manual using the Oracle `SRVCTL` command. The command changes the `AUTO_START` attribute of the Oracle database and instance resources.

### To prevent automatic database startup

- 1 Register the database, if not already registered:

```
$ srvctl add database -d db-name -o oracle_home \  
-p location-of-parameterfile -y manual
```

- 2 Once the database is registered, change the management policy for the database to manual:

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

- 3 Start the database:

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

# Configuring VCS service groups for Oracle RAC

You can set up the Oracle database to be managed by VCS or by Oracle Clusterware/Grid Infrastructure. Symantec recommends that the Oracle database be configured under VCS.


**Note:** The VCS Oracle agent is supported only for administrator-managed databases. For policy-managed databases, use the service group configuration without the VCS Oracle agent.

When the database is configured under VCS:

- You can choose to configure the service group in a way that insulates all the databases from failure in any of the databases in the group.
- VCS manages the start and stop sequence of the applications and the database.

Figure 21-2 illustrates a service group configuration with the VCS Oracle agent.

**Figure 21-2** Service group configuration with the VCS Oracle agent



If your configuration does not use the VCS Oracle agent, you need to modify the existing configuration as illustrated.

**Figure 21-3** illustrates a sample configuration in the absence of the VCS Oracle agent for single database configurations.

**Figure 21-3** Service group configuration without VCS Oracle agent (single database)





**Figure 21-4** illustrates a sample service group configuration in the absence of the VCS Oracle agent for multiple database configurations.

**Figure 21-4** Service group configuration without the VCS Oracle agent (multiple databases)



In a service group configuration without the VCS Oracle agent, Oracle Clusterware/Grid Infrastructure controls the database. An online local firm dependency exists between the Oracle Clusterware/Grid Infrastructure group and the CVM group. When the system starts, the CVM group brings up the volume and mount points for the databases. The Oracle Clusterware/Grid Infrastructure group brings up the OCR and voting disk, configures the private IP address for Oracle Clusterware/Grid Infrastructure, and starts Oracle Clusterware/Grid Infrastructure. Oracle Clusterware/Grid Infrastructure starts the database and the application is brought online.

**Note:** In a service group configuration without the VCS Oracle agent, when the system starts, all volumes and mount points MUST be online for the dependent service groups to be online.

To configure the Oracle database under VCS, create Oracle service groups after installing Oracle RAC and creating a database.

See [“Using the CLI to configure VCS service groups for Oracle RAC”](#) on page 413.

## Using the CLI to configure VCS service groups for Oracle RAC

This section describes how to configure the Oracle service group using the CLI.

The following procedure assumes that you have created the database.

### To configure the Oracle service group using the CLI

- 1 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the service group to the VCS configuration:

```
# hagrps -add oradb1_grp
```

- 3 Modify the attributes of the service group:

```
# hagrps -modify oradb1_grp Parallel 1
```

```
# hagrps -modify oradb1_grp SystemList galaxy 0 nebula 1
```

```
# hagrps -modify oradb1_grp AutoStartList galaxy nebula
```

- 4 Add the CVMVolDg resource for the service group:

```
# hares -add oradata_voldg CVMVolDg oradb1_grp
```

- 5 Modify the attributes of the CVMVolDg resource for the service group:

```
# hares -modify oradata_voldg CVMDiskGroup oradatadg
```

```
# hares -modify oradata_voldg CVMActivation sw
```

```
# hares -modify oradata_voldg CVMVolume oradatavol
```

- 6 Add the CFSSMount resource for the service group:

```
# hares -add oradata_mnt CFSSMount oradb1_grp
```

- 7 Modify the attributes of the CFSSMount resource for the service group:

```
# hares -modify oradata_mnt MountPoint "/oradata"
```

```
# hares -modify oradata_mnt BlockDevice \  
"/dev/vx/dsk/oradatadg/oradatavol"
```

- 8 Add the Oracle RAC database instance to the service group:

```
# hares -add oral Oracle oradb1_grp
```

**9** Modify the attributes of the Oracle resource for the service group:

```
# hares -modify ora1 Owner oracle
# hares -modify ora1 Home "/app/oracle/orahome"
# hares -modify ora1 StartUpOpt SRVCTLSTART
# hares -modify ora1 ShutDownOpt SRVCTLSTOP
```

**10** Localize the Sid attribute for the Oracle resource:

```
# hares -local ora1 Sid
```

**11** Set the Sid attributes for the Oracle resource on each system:

```
# hares -modify ora1 Sid vrts1 -sys galaxy
# hares -modify ora1 Sid vrts2 -sys nebula
```

**12** Set the dependencies between the CFSSMount resource and the CVMVolDg resource for the Oracle service group:

```
# hares -link oradata_mnt oradata_voldg
```

**13** Set the dependencies between the Oracle resource and the CFSSMount resource for the Oracle service group:

```
# hares -link ora1 oradata_mnt
```

**14** Create an online local firm dependency between the oradb1\_grp service group and the cvm service group:

```
# hagrps -link oradb1_grp cvm online local firm
```

**15** Enable the Oracle service group:

```
# hagrps -enableresources oradb1_grp
```

**16** Change the cluster configuration to the read-only mode:

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

**17** Bring the Oracle service group online on all the nodes:

```
# hagrps -online oradb1_grp -any
```

For more information and instructions on configuring the service groups using the CLI:

See the *Veritas Cluster Server Administrator's Guide*.

## Location of VCS log files

You may want to review the log files at `/var/VRTSvcs/log/engine_A.log` for errors or status messages. When large amounts of data are written, multiple log files may be written, such as `engine_B.log`, `engine_C.log`, and so on. The `engine_A.log` contains the most recent data.

## Removing permissions for communication

Make sure you completed the installation of SF Oracle RAC and the verification of disk support for I/O fencing. If you used `rsh`, remove the temporary `rsh` access permissions that you set for the nodes and restore the connections to the public network.

If the nodes use `ssh` for secure communications, and you temporarily removed the connections to the public network, restore the connections.

## Configuring the SFDB repository database after installation

If you want to use the Storage Foundation Database (SFDB) tools, you must set up the SFDB repository after installing and configuring SF Oracle RAC and Oracle. For SFDB repository set up procedures:

See *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases*





# Upgrading Oracle RAC

This chapter includes the following topics:

- [Supported upgrade paths](#)
- [Preparing to upgrade Oracle RAC](#)
- [Upgrading Oracle RAC binaries](#)
- [Migrating the Oracle RAC database](#)

## Supported upgrade paths

[Table 22-1](#) lists the upgrade paths for Oracle RAC.

**Table 22-1** Supported upgrade paths for Oracle RAC

| From current version     | Upgrade to                                                                       |
|--------------------------|----------------------------------------------------------------------------------|
| Oracle RAC 9i Release 2  | Oracle RAC 10g Release 2<br>Oracle RAC 11g Release 1                             |
| Oracle RAC 10g Release 1 | Oracle RAC 10g Release 2<br>Oracle RAC 11g Release 1<br>Oracle RAC 11g Release 2 |
| Oracle RAC 10g Release 2 | Oracle RAC 11g Release 1<br>Oracle RAC 11g Release 2                             |
| Oracle RAC 11g Release 1 | Oracle RAC 11g Release 2                                                         |

---

**Note:** When you upgrade to a different version of Oracle RAC, make sure that the full path of the Oracle Clusterware/Grid Infrastructure home directory and the Oracle database home directory is different from the path where the existing version of Oracle RAC resides.

---

The upgrade procedure assumes that the beginning configuration includes the following components, and that these components are running on the cluster nodes:

- SF Oracle RAC 5.1 SP1
- A supported version of the operating system

## Preparing to upgrade Oracle RAC

Complete the following preparatory tasks before you upgrade Oracle RAC:

1. Depending on the version you are upgrading from, perform the steps in one of the following sections:
  - See [“Preparing to upgrade from Oracle RAC 10g or Oracle RAC 11g ”](#) on page 418.
  - See [“Preparing to upgrade from Oracle RAC 9i”](#) on page 419.
2. Verify that the cluster configuration is compatible for upgrading Oracle RAC.
  - See [“Verifying the systems for Oracle RAC installation”](#) on page 377.

## Preparing to upgrade from Oracle RAC 10g or Oracle RAC 11g

Perform the following tasks before upgrading Oracle RAC.

### To prepare for upgrade from Oracle RAC 10g or Oracle RAC 11g

- 1 Take a hot or cold backup of the existing database.
- 2 Back up the existing Oracle home and central inventory.

- 3 If the Oracle RAC database is under VCS control, freeze the Oracle service groups to prevent VCS from reporting the resource as faulted when Oracle RAC stops and starts the database during the upgrade:

```
# haconf -makerw  
  
# hagrps -freeze oracle_group -persistent
```

- 4 Freeze the cvm service group to prevent VCS from reporting the resource as faulted when Oracle Clusterware is stopped and started during the upgrade:

```
# hagrps -freeze cvm_group -persistent  
  
# haconf -dump -makero
```

## Preparing to upgrade from Oracle RAC 9i

Perform the following pre-upgrade tasks before upgrading from Oracle RAC 9i.

### To prepare for upgrade from Oracle RAC 9i

- 1 Take a hot or cold backup of the existing Oracle RAC 9i Release 2 database.
- 2 Back up the existing Oracle Home and Central Inventory.
- 3 If the Oracle RAC database is under VCS control, freeze the Oracle service groups to prevent VCS from reporting the resource as faulted when Oracle RAC stops and starts the database during the upgrade:

```
# haconf -makerw  
  
# hagrps -freeze oracle_group -persistent
```

- 4 Freeze the CVM service group to prevent VCS from reporting the resource as faulted when Oracle Clusterware is stopped and started during the upgrade:

```
# hagrps -freeze cvm_group -persistent  
  
# haconf -dump -makero
```

- 5 Perform the following pre-installation steps:
  - Identify the public virtual addresses for use by Oracle
  - Set the kernel tunables and other parameters
  - Configure private IP addresses for Oracle RAC
  - Create voting disk volumes and mount points

- Set up Oracle user equivalence

For instructions on performing the tasks, see the chapter *Before installing Oracle RAC* in this guide.

## Upgrading Oracle RAC binaries

Review your Oracle installation manuals and the appropriate Oracle support Web sites before upgrading Oracle RAC.

---

**Note:** If you upgrade to Oracle RAC 11g Release 2, make sure that you upgrade with the same user credentials as that of the existing installation.

---

### To upgrade Oracle RAC binaries

- 1 Upgrade Oracle Clusterware.

---

**Note:** If you upgrade to Oracle RAC 11g Release 2, upgrade Oracle Clusterware to a new directory called the Oracle grid infrastructure home directory (GRID\_HOME).

Starting with Oracle RAC 11g Release 2, ASM must reside in the Oracle Grid Infrastructure home directory. If you plan to upgrade ASM to Release 11g R2, make sure that you upgrade it to run in the Oracle Grid Infrastructure home directory.

---

For instructions, see the Oracle RAC documentation.

- 2 Make sure that Oracle Clusterware is running.
- 3 Install the Oracle RAC database binaries.

For instructions, see the Oracle RAC documentation.

- 4 Complete the following post-installation tasks:

- Add Oracle RAC patches or patchsets.  
See [“Adding Oracle RAC patches or patchsets”](#) on page 395.
- Relink the SF Oracle RAC libraries with Oracle RAC.  
See [“Relinking the SF Oracle RAC libraries with Oracle RAC”](#) on page 400.
- For upgrades from Oracle RAC 9i: Add the CSSD resource to the VCS configuration.

# Migrating the Oracle RAC database

For instructions on migrating the existing Oracle RAC database, see the Oracle metalink documentation.

After migrating the database, complete the post-upgrade tasks:

See [“Performing post-upgrade tasks”](#) on page 421.

## Performing post-upgrade tasks

Perform the steps in the following procedure to complete the upgrade.

### To perform post-upgrade tasks

- 1 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 For upgrades from Oracle RAC 9i: Modify the resources in the Oracle service groups:

```
# hares -modify oracle_resource StartUpOpt SRVCTLSTART  
# hares -modify oracle_resource ShutDownOpt SRVCTLSTOP  
# hares -modify oracle_resource pfile "" -sys galaxy  
# hares -modify oracle_resource pfile "" -sys nebula
```

- 3 Add the Oracle UDP IPC private IP addresses to the Oracle `init.ora` file if the database cache fusion traffic is configured to use Oracle UDP IPC private IP addresses and if these addresses are configured as a PrivNIC or MultiPrivNIC resource for high availability.

See [“Adding Oracle UDP IPC private IP addresses to the Oracle initialization parameter file”](#) on page 407.

- 4 Modify the Oracle RAC configuration to prevent automatic startup of Oracle Clusterware.

See [“Preventing automatic startup of Oracle Clusterware/Grid Infrastructure”](#) on page 400.

- 5 Modify the Oracle RAC database configuration to prevent automatic database startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent.

See [“Preventing automatic database startup”](#) on page 409.

**6** Unfreeze the VCS service groups that were frozen earlier.

As root user, enter:

```
# hagrps -unfreeze oracle_group -persistent  
# hagrps -unfreeze cvm_group -persistent  
# haconf -dump -makero
```

**7** Migrate OCR and voting disk to CFS or ASM.

For migrating OCR and voting disk to CFS:

- Create the required storage on CFS.  
For instructions:  
See [“Creating the OCR and voting disk storage on CFS”](#) on page 338.
- Migrate OCR and voting disk to CFS.  
For instructions, see the Oracle documentation.

For instructions on migrating OCR and voting disk to ASM, see the Oracle documentation.

# Installation of SF Oracle RAC and Oracle RAC using a response file

- [Chapter 23. About response files](#)
- [Chapter 24. Installing and configuring SF Oracle RAC using a response file](#)
- [Chapter 25. Configuring I/O fencing using a response file](#)
- [Chapter 26. Installing Oracle RAC using a response file](#)
- [Chapter 27. Installing SF Oracle RAC and Oracle RAC using a response file](#)
- [Chapter 28. Response file variable definitions](#)





# About response files

This chapter includes the following topics:

- [About response files](#)
- [Response file syntax](#)
- [Guidelines for creating the SF Oracle RAC response file](#)
- [Installation scenarios for response files](#)

## About response files

Use response files to standardize and automate installations on multiple clusters. You can perform the following installation activities using a response file:

- Installing and configuring SF Oracle RAC
- Preparing the nodes for installation of Oracle RAC
- Installing Oracle RAC
- Relinking the SF Oracle RAC libraries with Oracle RAC
- Configuring the CSSD agent
- Upgrading SF Oracle RAC
- Uninstalling SF Oracle RAC

You can perform end-to-end installations or modular deployments.

For more information on modular deployments:

See “[Modular deployments using response files](#)” on page 427.

[Table 23-1](#) lists the various options available for creating or obtaining a response file.

Table 23-1 Options for obtaining a response file

| Option                                                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create a response file                                             | <p>Create a response file based on the response file template provided with SF Oracle RAC.</p> <p>The file is located at <code>/opt/VRTSvcs/rac/install.</code></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Reuse or customize the response files generated by an installation | <p>The Veritas installation programs generate a response file during the installation, configuration, upgrade, or uninstallation of SF Oracle RAC and for installer-based pre-installation tasks for Oracle RAC.</p> <p>The response file generated by the installer is located in the following directory:</p> <pre>/opt/VRTS/install/logs/installsfrac-installernumber\<br/>/installsfrac-installernumber.response file</pre> <p>You can reuse or customize the response files to perform a complete end-to-end installation or to perform a specialized activity, such as setting up the PrivNIC or MultiPrivNIC configuration on your clusters. All you need to do is update the response file variable definitions to enable or disable the options, depending on the task you want to perform.</p> <p><b>Note:</b> Response files are not created if the tasks terminated abruptly or if you entered <b>q</b> to quit the installation. To generate the response file when you plan to discontinue a task, use the <b>Exit SF Oracle RAC configuration</b> option.</p> |
| Use the installation simulator                                     | <p>Create a response file by specifying the <code>-makeresponsefile</code> option with the SF Oracle RAC installer.</p> <p>Mount the product disc and navigate to the folder that contains the installation program. Start the installation program.</p> <pre># ./installsfrac -makeresponsefile</pre> <p>The option creates a response file by simulating an installation. The response file is created in the directory <code>/opt/VRTS/install/logs/.</code></p> <p><b>Note:</b> You can use the <code>-makeresponsefile</code> option to create response files only for installing, configuring, and uninstalling SF Oracle RAC.</p> <p>For more information:</p> <p>See <a href="#">“About the installation simulator”</a> on page 428.</p>                                                                                                                                                                                                                                                                                                                             |

At the end of the SF Oracle RAC installation, the following files are created:

- A log file that contains executed system commands and output.
- A summary file that contains the output of the installation scripts.
- Response files to be used with the `-responsefile` option of the installer.

---

**Note:** The SF Oracle RAC response files also contain VCS variables used for the installation and configuration of VCS.

For the VCS variable definitions, see the *Veritas Cluster Server Installation Guide*.

---

## Modular deployments using response files

Modular deployments offer flexibility in planning your installation and configuration activities. You can choose to perform any installation or configuration task independent of other related activities. For example, you can disable Oracle user and group creation across all nodes in a cluster while installing Oracle RAC.

Modular deployments are supported for the following installation and configuration activities:

- Installing SF Oracle RAC
- Configuring SF Oracle RAC
- Creating Oracle user and group
- Creating storage for OCR and voting disk
- Configuring private IP addresses for Oracle RAC
- Installing Oracle Clusterware
- Installing Oracle database
- Oracle post-installation tasks:
  - Configuring CSSD agent
  - Relinking Oracle RAC libraries

The variables required for each of the above tasks are described in the chapter "Response file variable definitions".

All modular deployment activities require the following variable definitions:

```
$CFG{prod}  
$CFG{systems}
```

## About the installation simulator

The SF Oracle RAC installer includes the option (`-makeresponsefile`) to simulate the installation, configuration, or uninstallation of the product. The simulation option steps through the installation script, including all of the preinstallation checks on the systems. However, the simulation does not actually install the packages, uninstall previously installed packages, or start or stop any processes.

Use the installation simulator in the following situations:

- To understand the information that is required when you install, configure, or uninstall SF Oracle RAC  
As the simulator steps through the same code that is used by the installer, you can use the simulation process to preview the installation or configuration steps without disrupting your existing installation. You can skip the preinstallation checks, if not required.
- To create a response file  
The simulation process enables you to create a response file that can be used as a template for installing, configuring, or uninstalling SF Oracle RAC. You can customize the response file, as required.

To simulate an installation or configuration, specify the `-makeresponsefile` option with the installer or product installation script at the command line.

To simulate an uninstallation, specify the `-makeresponsefile` option with the installer or the product uninstall script at the command line.

## Response file syntax

The Perl statement syntax that is included in the response file varies, depending on whether “Scalar” or “List” values are required by the variables.

For example,

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=["value", "value", "value"];
```

# Guidelines for creating the SF Oracle RAC response file

This section provides guidelines for creating the SF Oracle RAC response file.

1. Create a response file using one of the available options.  
 For various options on creating or obtaining an SF Oracle RAC response file:  
 See [“About response files”](#) on page 425.
2. Set the following master values to 1 to enable SF Oracle RAC installation and configuration.

---

**Note:** The master settings must be set to 1 to enable the installer to read dependent variable definitions. For example, if the value `$CFG{opt}{install}` is not set to 1, the other dependent installation values in the response file will be disregarded. This is true for any master setting.

---

The following is the list of master values that must be set for installing and configuring SF Oracle RAC.

Installing SF Oracle RAC    `$CFG{opt}{install}=1;`

Configuring SF Oracle RAC    `$CFG{opt}{configure}=1;`  
                                      `$CFG{config_sfrac_subcomponents} = 1;`

3. Now, set the appropriate value in the dependent variable definitions for installing and configuring SF Oracle RAC.

The set of minimum definitions for a successful installation and configuration is as follows:

```
$CFG{accepteula}=1;
$CFG{opt}{install}=1;
$CFG{opt}{configure}=1;
$CFG{config_sfrac_subcomponents} = 1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{vxkeyless}=1;
$CFG{prod}="SFRAC51";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=101;
```

```
$CFG{vcs_clustername}="rac_cluster101";
$CFG{vcs_lltlink1}{galaxy}="bge1";
$CFG{vcs_lltlink1}{nebula}="bge1";
$CFG{vcs_lltlink2}{galaxy}="bge2";
$CFG{vcs_lltlink2}{nebula}="bge2";
$CFG{vcs_userenpw}=[ qw(gpqIpkPmqLqqOyqKpn) ];
$CFG{vcs_username}=[ qw(admin) ];
$CFG{vcs_userpriv}=[ qw(Administrators) ];
```

You can add more variable definitions, as required.

4.
- Set the following master values to 1 to enable the Oracle pre-installation tasks.

Note: This step assumes that you have already set the master value `$CFG{opt}{configure}` to 1 in the previous step.

The following is the list of additional master values that must be set for completing Oracle pre-installation tasks:

|                                                     |                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Private IP address configuration                    | <div>If you plan to use the PrivNIC agent, use the following master settings:</div> <div><pre>\$CFG{config_privnic}=1; \$CFG{config_multiprivnic}=0;</pre></div> <div>If you plan to use the MultiPrivNIC agent, use the following master settings:</div> <div><pre>\$CFG{config_multiprivnic}=1; \$CFG{config_privnic}=0;</pre></div> |
| OCR and voting disk storage                         | <pre>\$CFG{create_ocr_vote_storage}=1;</pre>                                                                                                                                                                                                                                                                                           |
| Oracle Clusterware/Grid Infrastructure installation | <pre>\$CFG{install_oracle_clusterware}=1;</pre>                                                                                                                                                                                                                                                                                        |
| Oracle database installation                        | <pre>\$CFG{install_oracle_database}=1;</pre>                                                                                                                                                                                                                                                                                           |
| Post Oracle tasks                                   | <div><pre>\$CFG{config_cssd_agent}=1; \$CFG{relink_oracle_database}=1;</pre></div>                                                                                                                                                                                                                                                     |

Set the master value for the `$CFG{create_oracle_user_group}` variable to 0 as follows:

```
$CFG{create_oracle_user_group}=0;
```

---

**Note:** If you are performing an end-to-end installation, the creation of Oracle user and group must be completed before starting the installation. This is because user equivalence must be manually established on all nodes to allow the Oracle Universal Installer to securely copy files and run programs on the nodes in the cluster without requiring password prompts.

---

5. Set the appropriate value for the dependent variable definitions for the Oracle pre-installation tasks.

---

**Note:** The following table discusses variable definitions with corresponding sample values. Replace the sample values with those that are appropriate for your installation requirements.

---

Private IP address  
configuration

If you plan to use the PrivNIC agent, provide the following definitions:

```
$CFG{privnic_resname} = "ora_priv";
$CFG{privnic_interface_priority}="bge2 bge3";
$CFG{galaxy}{privnicip} = "192.168.12.1";
$CFG{galaxy}->{hostname_for_ip} = "galaxy-priv";
$CFG{nebula}{privnicip} = "192.168.12.2";
$CFG{nebula}->{hostname_for_ip} = "nebula-priv";
$CFG{nic_netmask} = "255.255.255.0";
$CFG{nic_add_ip_to_files} = 1;
```

If you plan to use the MultiPrivNIC agent, use the following master settings:

```
$CFG{multiprivnic_resname} = "ora_priv";
$CFG{nic_add_ip_to_files} = 1;
$CFG{galaxy}->{bge1}->{multiprivnicip}
= "192.168.12.1";
$CFG{galaxy}->{bge1}->{hostname_for_ip}
= "galaxy-priv";
$CFG{galaxy}->{bge2}->{multiprivnicip}
= "192.168.2.1";
$CFG{galaxy}->{bge2}->{hostname_for_ip}
= "galaxy-priv1";
$CFG{nebula}->{bge1}->{multiprivnicip}
= "192.168.12.2";
$CFG{nebula}->{bge1}->{hostname_for_ip}
= "nebula-priv";
$CFG{nebula}->{bge2}->{multiprivnicip}
= "192.168.2.2";
$CFG{nebula}->{bge2}->{hostname_for_ip}
= "nebula-priv1";
$CFG{nic_netmask} = "255.255.255.0";
```



**OCR and voting disk storage**

If you choose to create the OCR and voting disk storage:

```
$CFG{ocrvotedgoption} = 0;
$CFG{ocrvotescheme} = 1;
$CFG{ocrvotedisks} = [ qw(Disk_1 Disk_2) ];
$CFG{ocrvotedgname} = "ocrvotedg";
$CFG{ocrvotevolname} = "ocrvotevol";
$CFG{ocrvotevolsize} = 640;
$CFG{ocrvotemount} = "/ocrvote";
$CFG{oracle_user} = "oracle";
$CFG{oracle_group} = "oinstall";
```

If you choose to use an existing storage for the OCR and voting disk storage:

```
$CFG{ocrvotedgoption} = 1;
$CFG{ocrvotescheme} = 0;
$CFG{ocrvotedgname} = "ocrvotedg";
$CFG{ocrvotevolname} = "ocrvotevol";
$CFG{ocrvotevolsize} = 640;
$CFG{ocrvotemount} = "/ocrvote";
$CFG{oracle_user} = "oracle";
$CFG{oracle_group} = "oinstall";
```

**Oracle Clusterware/Grid Infrastructure installation**

```
$CFG{oracle_user} = "oracle";
$CFG{oracle_group} = "oinstall";
$CFG{oracle_base} = "/u01/app/oracle";
$CFG{crs_home} = "/u01/app/oracle/product/11.2.0/crshome";
$CFG{crs_installpath} = "/cdrom/oracle/clusterware";
$CFG{oracle_version} = "11.2.0.1";
$CFG{crs_responsefile} = "/oracle/crs.rsp";
```

**Oracle database installation**

```
$CFG{oracle_user} = "oracle";
$CFG{oracle_group} = "oinstall";
$CFG{oracle_base} = "/u01/app/oracle";
$CFG{crs_home} = "/u01/app/oracle/product/11.2.0/crshome";
$CFG{db_home} = "/u01/app/oracle/product/11.2.0/dbhome";
$CFG{db_installpath} = "/cdrom/oracle/database";
$CFG{oracle_version} = "11.2.0.1";
$CFG{db_responsefile} = "/oracle/db.rsp";
```

Post Oracle tasks

The master value (set in the previous step) suffices for CSSD resource configuration.

For relinking libraries, the dependent variable settings are as follows:

```
$CFG{oracle_user} = "oracle";  
$CFG{oracle_group} = "oinstall";  
$CFG{crs_home} = "/u01/app/oracle/product/11.2.0/crshome";  
$CFG{db_home} = "/u01/app/oracle/product/11.2.0/dbhome_1";  
$CFG{oracle_version} = "11.2.0.1";
```

## Installation scenarios for response files

The chapters in this section cover the following installation scenarios using response files:

- Installing and configuring SF Oracle RAC  
See [“Installing and configuring SF Oracle RAC”](#) on page 435.
- Installing Oracle RAC  
See [“About installing Oracle RAC using response files”](#) on page 443.
- Installing both SF Oracle RAC and Oracle RAC  
See [“Installing SF Oracle RAC and Oracle RAC”](#) on page 452.

# Installing and configuring SF Oracle RAC using a response file

This chapter includes the following topics:

- [Installing and configuring SF Oracle RAC](#)
- [Sample response file for installing and configuring SF Oracle RAC](#)

## Installing and configuring SF Oracle RAC

You can create a single response file or separate response files for installing and configuring SF Oracle RAC.

The installer performs the following tasks:

- Installs SF Oracle RAC.
- Configures SF Oracle RAC.

---

**Note:** Fencing is configured in disabled mode. You must configure I/O fencing after the installation.

---

The following sample procedure uses a single response file for installing and configuring SF Oracle RAC.

### To install and configure SF Oracle RAC using response files

- 1 Make sure that the systems meet the installation requirements.

- 2 Complete the preparatory steps before starting the installation.  
For instructions, see the chapter "Preparing to install SF Oracle RAC" in this document.
- 3 Create a response file using one of the available options.  
For information on various options available for creating a response file:  
See [“About response files”](#) on page 425.

---

**Note:** You must replace the host names in the response file with that of the new systems in the cluster.

---

For guidelines on creating a response file:

See [“Guidelines for creating the SF Oracle RAC response file”](#) on page 429.

For a sample response file:

See [“Sample response file for installing and configuring SF Oracle RAC”](#) on page 437.

- 4 Mount the product disc and navigate to the product directory that contains the installation program.
- 5 Start the installation and configuration:

```
# ./installsffrac -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the full path name of the response file.

---

**Note:** Fencing is configured in disabled mode. You need to configure fencing after the configuration.

---

- 6 Configure I/O fencing.

---

**Note:** Before you configure I/O fencing, make sure that you complete the required pre-configuration tasks. For instructions, see the chapter *Preparing to configure SF Oracle RAC* in this document.

---

For instructions on configuring I/O fencing using a response file, see the chapter *Configuring I/O fencing using a response file* in this document.

- 7 Complete the SF Oracle RAC post-installation tasks.  
For instructions, see the chapter *Performing post-installation and configuration tasks* in this document.

# Sample response file for installing and configuring SF Oracle RAC

The following sample response file installs and configures SF Oracle RAC on two nodes, galaxy and nebula.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{opt}{vxkeyless}=1;
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{prod}="SFRAC51";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{opt}{configure}=1;
$CFG{config_sfrac_subcomponents} = 1;
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=101;
$CFG{vcs_clustername}="rac_cluster101";
$CFG{vcs_lltlink1}{galaxy}="bge1";
$CFG{vcs_lltlink1}{nebula}="bge1";
$CFG{vcs_lltlink2}{galaxy}="bge2";
$CFG{vcs_lltlink2}{nebula}="bge2";
$CFG{vcs_username}=[ qw(admin) ];
$CFG{vcs_userpriv}=[ qw(Administrators) ];
$CFG{vcs_userenpw}=[ qw(gpqIpkPmqLqqOyqKpn) ];
$CFG{uploadlogs}=0;

1;
```



# Configuring I/O fencing using a response file

This chapter includes the following topics:

- [Configuring I/O fencing using response files](#)
- [Sample response file for configuring disk-based I/O fencing](#)
- [Sample response file for configuring server-based I/O fencing](#)

## Configuring I/O fencing using response files

Typically, you can use the response file that the installer generates after you perform I/O fencing configuration to configure I/O fencing for SF Oracle RAC.

### To configure I/O fencing using response files

- 1 Make sure that SF Oracle RAC is configured.
- 2 Based on whether you want to configure disk-based or server-based I/O fencing, make sure you have completed the preparatory tasks.  
[See “About planning to configure I/O fencing” on page 67.](#)
- 3 Copy the response file to one of the cluster systems where you want to configure I/O fencing.  
[See “Sample response file for configuring disk-based I/O fencing” on page 440.](#)  
[See “Sample response file for configuring server-based I/O fencing” on page 440.](#)

- 4 Edit the values of the response file variables as necessary.  
See [“Response file variables to configure disk-based I/O fencing”](#) on page 469.  
See [“Response file variables to configure server-based I/O fencing”](#) on page 471.
- 5 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsfrac -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

## Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions.

See [“Response file variables to configure disk-based I/O fencing”](#) on page 469.

```
#
# Configuration Values:
#
our %CFG;

$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;

$CFG{prod}="SFRAC51";

$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="rac_cluster101";
$CFG{vxfen_config_fencing_dg}="fendg";
$CFG{vxfen_config_fencing_mechanism}="dmp";
$CFG{vxfen_config_fencing_newdg_disks}=
  [ qw(c1t1d0s2 c2t1d0s2 c3t1d0s2) ];
$CFG{vxfen_config_fencing_option}=2;
```

## Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing :



```
$CFG{fencing_cpc_config_cpagent}=0;
$CFG{fencing_cpc_cps}=[ qw(10.200.117.145) ];
$CFG{fencing_cpc_dgname}="vxfencoorddg";
$CFG{fencing_cpc_diffab}=0;
$CFG{fencing_cpc_disks}=[ qw(emc_clariion0_37 emc_clariion0_13) ];
$CFG{fencing_cpc_mechanism}="raw";
$CFG{fencing_cpc_ncps}=3;
$CFG{fencing_cpc_ndisks}=2;
$CFG{fencing_cpc_ports}{"10.200.117.145"}=14250;
$CFG{fencing_cpc_reusedg}=1;
$CFG{fencing_cpc_security}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFRAC51";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="rac_cluster101";
$CFG{vxfen_config_fencing_option}=1;
```



# Installing Oracle RAC using a response file

This chapter includes the following topics:

- [About installing Oracle RAC using response files](#)
- [Before you install](#)
- [Installing Oracle RAC](#)
- [Sample response file for installing Oracle RAC](#)

## About installing Oracle RAC using response files

You can perform an unattended pre-configuration of your systems by running the SF Oracle RAC installer with the `-responsefile` option. This capability in tandem with the Oracle response files for Oracle Clusterware and Oracle database installation enables you to standardize and automate Oracle RAC deployments in your cluster.

## Before you install

Make sure that you complete the tasks in the following procedure before starting the silent installation of Oracle RAC.

### To prepare the systems for installing Oracle RAC using response files

- 1 Complete the following preparatory tasks on the nodes manually or by using the SF Oracle RAC installer:
  - Identify the public virtual IP addresses for use by Oracle
  - Plan for the storage and the private network configuration for Oracle RAC

- Set the kernel parameters
- Verify the user "nobody" exists
- Create Oracle user and groups
- Set up Oracle user equivalence
- Edit the Oracle user profile

For instructions:

See [“About preparing to install Oracle RAC”](#) on page 327.

- 2 Comment the ORCLudlm pre-requisite check in the `crs_prereq.xml` file to suppress the check during the installation of Oracle Clusterware.

---

**Note:** The silent installation of Oracle Clusterware fails if the ORCLudlm check is not suppressed.

---

Change to the directory where you have mounted the Oracle Clusterware software. Navigate to the file `stage/prereq/crs/crs_prereq.xml` and comment the following line, if present in the file (the line may or may not be present on Solaris x64 systems):

```
<PREREQUISITEREF NAME="ORCLudlm_check" SEVERITY="Error"/>
```

Comment the line as follows:

```
<!--PREREQUISITEREF NAME="ORCLudlm_check" SEVERITY="Error"/-->
```

- 3 Create response files for Oracle Clusterware/Grid Infrastructure and Oracle database installation using the response file template provided by Oracle RAC.

For instructions on using the response file template, see the Oracle RAC documentation.

---

**Note:** Keep at hand the full path of the directory where these response files will be saved. The full path of the Oracle Clusterware/Grid Infrastructure response file must be set in the SF Oracle RAC response file variable `$CFG{crs_responsefile};`. The full path of the Oracle database response file must be set in the SF Oracle RAC response file variable `$CFG{db_responsefile};`.

---

**4** Create an SF Oracle RAC response file.

For information on various options available for creating a response file:

See [“About response files”](#) on page 425.

Make sure that you provide the full path information of the Oracle Clusterware/Grid Infrastructure and Oracle database response files.

For guidelines on creating a response file:

See [“Guidelines for creating the SF Oracle RAC response file”](#) on page 429.

For information on the list of required and optional variables:

See [“Response file variable definitions for Oracle RAC”](#) on page 473.

For a sample response file:

See [“Sample response file for installing Oracle RAC”](#) on page 446.

- 5** Make sure that the Oracle user has appropriate read and write permissions on the response files.
- 6** Make sure that passwordless communication between Oracle users is set up on the nodes of the cluster.

## Installing Oracle RAC

To perform a silent installation of Oracle RAC, run the SF Oracle RAC installer.

The installer supports completion of the following tasks using the response file:

- Creates storage for OCR and voting disk
- Configures the private network for Oracle RAC
- Installs Oracle Clusterware/Grid Infrastructure by leveraging the corresponding Oracle response file
- Installs Oracle database by leveraging the Oracle database response file
- Relinks the Oracle RAC libraries with SF Oracle RAC libraries
- Configures the CSSD agent

The sample procedure assumes that Oracle user equivalence is established on all the nodes in the cluster.

**To perform Oracle pre-installation and installation on the nodes**

- 1 Navigate to the directory containing the SF Oracle RAC installer:

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

- 2 Start the installation:

```
# ./installsfrc -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the full path name of the SF Oracle RAC response file.

- 3 Run the `root.sh` script as the root user on the cluster nodes.

---

**Note:** Do not run the script simultaneously on your cluster nodes.

---

- 4 Complete the following Oracle post-installation tasks:

- Add any patches or patchsets required by Oracle RAC .
- Create the Oracle RAC database.
- Add the Oracle UDP IPC private IP addresses to the Oracle `init.ora` file if the database cache fusion traffic is configured to use Oracle UDP IPC private IP addresses and if these addresses are configured as a PrivNIC or MultiPrivNIC resource for high availability.
- Configure the Oracle RAC database for manual startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent.
- Configure the VCS service groups for Oracle RAC.
- Verify the cluster.
- Remove the temporary communication permissions.

For instructions:

See the chapter *Performing Oracle RAC post-installation tasks*.

## Sample response file for installing Oracle RAC

The following sample response file installs Oracle RAC 11g Release 2 and performs the following Oracle RAC pre-installation and installation tasks on two nodes, galaxy and nebula, in the cluster:

- Creates a disk group for OCR and voting disk storage

- Creates OCR and voting disk storage on raw volumes
- Configures the PrivNIC agent for high availability on both nodes
- Installs Oracle Clusterware
- Installs Oracle database
- Relinks the Oracle RAC libraries with SF Oracle RAC libraries
- Configures the CSSD agent

```
#
# Configuration Values:
#
our %CFG;

$CFG{prod}="SFRAC51";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{opt}{configure}=1;

$CFG{config_privnic}=1;
$CFG{privnic_resname}="ora_priv";
$CFG{privnic_interface_priority}="bge2 bge3";
$CFG{galaxy}{privnicip}=" 192.168.12.1";
$CFG{galaxy}->{hostname_for_ip}="galaxy-priv";
$CFG{nebula}{privnicip}=" 192.168.12.2";
$CFG{nebula}->{hostname_for_ip}="nebula-priv";
$CFG{nic_netmask}="255.255.255.0";
$CFG{nic_add_ip_to_files}=1;
$CFG{nic_reuseip}=1;
$CFG{nic_reusealias}=1;

$CFG{create_ocr_vote_storage}=1;
$CFG{ocrvotedgoption}=0;
$CFG{oracle_group}="dba";
$CFG{grid_user}="grid";
$CFG{oracle_user}="oracle";
$CFG{ocrvolname}="ocrvotevol";
$CFG{ocrvolsize}=640;
$CFG{ocrvotedgname}="ocrvotedg";
$CFG{ocrvotedisks}=[qw(Disk_1 Disk_2)];
$CFG{ocrvotescheme}=0;
$CFG{votevolname}="votevol";
$CFG{votevolsize}=500;
```

```
$CFG{install_oracle_clusterware}=1;
$CFG{install_oracle_database}=1;
$CFG{oracle_base} = "/u01/app/oracle";
$CFG{crs_home}="/u01/app/oracle/product/11.2.0/crshome";
$CFG{db_home}="/u01/app/oracle/product/11.2.0/dbhome_1";
$CFG{crs_installpath}="/cdrom/oracle/clusterware";
$CFG{db_installpath}="/cdrom/oracle/database";
$CFG{crs_responsefile}="/oracle/crs.rsp";
$CFG{db_responsefile} = "/oracle/db.rsp";

$CFG{config_cssd_agent}=1;
$CFG{relink_oracle_database}=1;
$CFG{uploadlogs}=0;

1;
```



# Installing SF Oracle RAC and Oracle RAC using a response file

This chapter includes the following topics:

- [About installing SF Oracle RAC and Oracle RAC using response files](#)
- [Before you install](#)
- [Installing SF Oracle RAC and Oracle RAC](#)
- [Sample response file for installing SF Oracle RAC and Oracle RAC](#)

## About installing SF Oracle RAC and Oracle RAC using response files

SF Oracle RAC integrates with the silent installation capabilities of the Oracle Universal Installer (OUI) to perform an end-to-end silent installation of SF Oracle RAC and Oracle RAC. The installation process uses the SF Oracle RAC response file to perform the installation, configuration, and Oracle pre-installation tasks in tandem with the Oracle RAC response files to install Oracle Clusterware and Oracle database. The SF Oracle RAC response file must be supplied with the full path of the location where the Oracle Clusterware and Oracle database response files reside.

You must have the following response files to perform a complete end-to-end installation:

|                                                               |                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| An SF Oracle RAC installation and configuration response file | <p>Contains the information listed in the following section:</p> <p>See <a href="#">“Information required in the SF Oracle RAC response file”</a> on page 450.</p> <p>For various options on creating or obtaining an SF Oracle RAC response file:</p> <p>See <a href="#">“About response files”</a> on page 425.</p> |
| An Oracle Clusterware installation response file              | <p>Contains Oracle Clusterware installation information</p> <p>The response file may be created as described in the Oracle RAC documentation.</p>                                                                                                                                                                     |
| An Oracle database installation response file                 | <p>Contains Oracle database installation information</p> <p>The response file may be created as described in the Oracle RAC documentation.</p>                                                                                                                                                                        |

To perform the installation, you need to invoke the SF Oracle RAC installer with the `-responsefile` option.

---

**Note:** Some pre-installation and post-installation tasks must be manually completed before and after the installation, respectively. Complete these tasks as described in the procedures in this chapter.

---

## Information required in the SF Oracle RAC response file

The SF Oracle RAC response file must contain the following information for an end-to-end installation:

- SF Oracle RAC installation information
- SF Oracle RAC configuration information
- Oracle RAC pre-installation information, such as follows:
  - Grid user information
  - Oracle user and group information
  - OCR and voting disk storage information
  - Private IP address configuration information
  - Full path to the Oracle Clusterware and Oracle database installation binaries
  - Full path of the Oracle Clusterware and Oracle database home directories
  - Full path of the Oracle Clusterware response file

- Full path of the Oracle database response file
- Oracle RAC post-installation information, such as follows:
  - CSSD resource configuration
  - Relinking Oracle RAC libraries with SF Oracle RAC

## Before you install

SF Oracle RAC integrates with the silent installation capabilities of the Oracle Universal Installer (OUI) to provide an end-to-end installation of SF Oracle RAC and Oracle RAC. Complete the steps in the following procedure before you install SF Oracle RAC and Oracle RAC using response files.

### To prepare to install SF Oracle RAC and Oracle RAC using response files

- 1 Make sure that the systems meet the installation requirements.
- 2 Complete the preparatory tasks for installing SF Oracle RAC.  
See [“Preparing to install Oracle RAC using the SF Oracle RAC installer or manually”](#) on page 328.
- 3 Set the kernel parameters for Oracle RAC.  
For instructions, see the Oracle RAC documentation.
- 4 Create the users and groups required by Oracle.  
See [“Creating users and groups for Oracle RAC”](#) on page 331.
- 5 Set up Oracle user equivalence on all nodes.  
See [“Setting up user equivalence”](#) on page 376.
- 6 Comment the ORCLudlm pre-requisite check in the `crs_prereq.xml` file to suppress the check during the installation of Oracle Clusterware.

---

**Note:** The silent installation of Oracle Clusterware fails if the ORCLudlm check is not suppressed.

---

Change to the directory where you have mounted the Oracle Clusterware software. Navigate to the file `stage/prereq/crs/crs_prereq.xml` and comment the following line, if present in the file (the line may or may not be present on Solaris x64 systems):

```
<PREREQUISITEREF NAME="ORCLudlm_check" SEVERITY="Error"/>
```

- 7
- Create response files for Oracle Clusterware and Oracle database installation using the response file template provided by Oracle RAC.

For instructions on creating the response files for Oracle Clusterware and Oracle database, see the Oracle RAC documentation.

Make sure that the Oracle user has read/write permissions on the Oracle Clusterware and Oracle database response files.

**Note:** Keep at hand the full path of the directory where these response files are located. The full path of the Oracle Clusterware response file must be set in the SF Oracle RAC response file variable `$CFG{crs_responsefile};`. The full path of the Oracle database response file must be set in the SF Oracle RAC response file variable `$CFG{db_responsefile};`.

- 8
- Create an SF Oracle RAC response file. Make sure that you provide the full path information of the Oracle Clusterware and Oracle database response files.

For guidelines on creating an SF Oracle RAC response file:  
See [“Guidelines for creating the SF Oracle RAC response file”](#) on page 429.  
For the list of variable definitions:

|                                        |                                                                                                                                                          |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| For SF Oracle RAC variable definitions | For installing SF Oracle RAC<br>See <a href="#">“Response file variables for installing SF Oracle RAC”</a> on page 457.<br>For configuring SF Oracle RAC |
| For Oracle RAC variable definitions    | See <a href="#">“Response file variable definitions for Oracle RAC”</a> on page 473.                                                                     |

For a sample response file:  
See [“Sample response file for installing SF Oracle RAC and Oracle RAC”](#) on page 454.

# Installing SF Oracle RAC and Oracle RAC

The installer performs the following tasks using the response file:

- Installs and configures SF Oracle RAC.

---

**Note:** Fencing is configured in disabled mode. Configure fencing after the installation of SF Oracle RAC and Oracle RAC is complete.

---

- Installs Oracle Clusterware/Grid Infrastructure and Oracle database.

---

**Note:** The SF Oracle RAC installer pauses the installation if VIPCA fails to run silently during the installation of Oracle RAC 10g Release 2. Run the script manually on one of the nodes as the root user:

```
# export DISPLAY=10.20.12.150:0.0
# cd $CRS_HOME/bin
# ./vipca
```

Resume the installation.

---

- Configures the CSSD agent.
- Relinks the Oracle RAC libraries with SF Oracle RAC.

If any of the installation tasks fail, the installer quits the installation. To resume the installation:

- Review the log file to identify the cause of failure.
- Resolve the issue.
- Review the messages on the console to identify the tasks that have completed successfully. Disable these tasks in the response file.
- Run the installer again using the modified response file.

**To install and configure SF Oracle RAC and Oracle RAC using a response file**

- 1 Mount the product disc and navigate to the product directory that contains the installation program.
- 2 Start the installation and configuration:

```
# ./installsfrc -responsefile /tmp/response_file
```

Where /tmp/response\_file is the full path name of the response file.

---

**Note:** Fencing is configured in disabled mode. You need to configure fencing after the configuration.

---

- 3 Run the `root.sh` script as the root user on the cluster nodes.

---

**Note:** Do not run the script simultaneously on your cluster nodes.

---

- 4 After the installation completes, initialize disks as VxVM disks and configure fencing.

See [“Initializing disks as VxVM disks”](#) on page 169.

For instructions on configuring I/O fencing using a response file, see the chapter *Configuring I/O fencing using a response file* in this document.

- 5 Complete the SF Oracle RAC post-installation tasks.

For instructions, see the chapter *Performing post-installation and configuration tasks* in this document.

- 6 Complete the following Oracle post-installation tasks:

- Add any patches or patchsets required by Oracle RAC .
- Create the Oracle RAC database.
- Add the Oracle UDP IPC private IP addresses to the Oracle `init.ora` file if the database cache fusion traffic is configured to use Oracle UDP IPC private IP addresses and if these addresses are configured as a PrivNIC or MultiPrivNIC resource for high availability.
- Configure the Oracle RAC database for manual startup if you want the Oracle RAC database to be managed by VCS using the Oracle agent.
- Configure the VCS service groups for Oracle RAC.
- Verify the cluster.
- Remove the temporary communication permissions.

For instructions:

See the chapter *Performing Oracle RAC post-installation tasks* in this document.

## Sample response file for installing SF Oracle RAC and Oracle RAC

The sample response file installs SF Oracle RAC 5.1 SP1 and Oracle RAC 11g Release 2 and performs the following installation and configuration tasks on two nodes, galaxy and nebula, in the cluster:

- Installs all SF Oracle RAC packages
- Configures SF Oracle RAC with two private interconnects on each node
- Creates a disk group for OCR and voting disk storage
- Creates OCR and voting disk storage on CFS
- Configures the PrivNIC agent for high availability on both nodes
- Installs Oracle Clusterware
- Installs Oracle database
- Relinks the Oracle RAC libraries with SF Oracle RAC libraries
- Configures the CSSD agent

```
#
# Configuration Values:
#
our %CFG;

$CFG{accepteula}=1;
$CFG{prod}="SFRAC51";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{vxkeyless}=1;

$CFG{opt}{configure}=1;
$CFG{config_sfrac_subcomponents} = 1;
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=101;
$CFG{vcs_clustername}="rac_cluster101";
$CFG{vcs_lltlink1}{galaxy}="bge1";
$CFG{vcs_lltlink1}{nebula}="bge1";
$CFG{vcs_lltlink2}{galaxy}="bge2";
$CFG{vcs_lltlink2}{nebula}="bge2";

$CFG{create_oracle_user_group}=0;

$CFG{config_privnic}=1;
$CFG{privnic_resname}="ora_priv";
$CFG{privnic_interface_priority}="bge2 bge3";
$CFG{galaxy}{hostname_for_ip}="galaxy-priv";
$CFG{galaxy}{privnicip}="192.168.12.1";
$CFG{nebula}{hostname_for_ip}="nebula-priv";
```

```
$CFG{nebula}{privnicip}="192.168.12.2";
$CFG{nic_netmask}="255.255.255.0";
$CFG{nic_add_ip_to_files}=1;

$CFG{create_ocr_vote_storage}=1;
$CFG{ocrvotedgoption}=0;
$CFG{oracle_group}="dba";
$CFG{grid_user}="grid";
$CFG{oracle_user}="oracle";
$CFG{ocrvotedgname}="ocrvotedg";
$CFG{ocrvotedisks}=[ qw(Disk_1 Disk_2) ];
$CFG{ocrvotemount}="/ocrvote";
$CFG{ocrvotescheme}=1;
$CFG{ocrvotevolname}="ocrvotevol";
$CFG{ocrvotevolsize}=640;

$CFG{install_oracle_clusterware}=1;
$CFG{install_oracle_database}=1;
$CFG{oracle_base}="/u01/app/oracle";
$CFG{crs_home}="/u01/app/oracle/product/11.2.0/crshome";
$CFG{db_home}="/u01/app/oracle/product/11.2.0/dbhome_1";
$CFG{crs_installpath}="/cdrom/oracle/clusterware";
$CFG{db_installpath}="/cdrom/oracle/database";
$CFG{crs_responsefile} = "/oracle/crs.rsp";
$CFG{db_responsefile} = "/oracle/db.rsp";
$CFG{oracle_display}="10.20.12.150:0.0";

$CFG{config_cssd_agent}=1;
$CFG{relink_oracle_database}=1;

$CFG{uploadlogs}=0;

1;
```



# Response file variable definitions

This chapter includes the following topics:

- [Response file variables for installing SF Oracle RAC](#)
- [Response file variables to configure SF Oracle RAC](#)
- [Response file variables to configure disk-based I/O fencing](#)
- [Response file variables to configure server-based I/O fencing](#)
- [Response file variable definitions for Oracle RAC](#)

## Response file variables for installing SF Oracle RAC

[Table 28-1](#) lists the response file variables that you can define to install SF Oracle RAC.

**Table 28-1** Response file variables for installing SF Oracle RAC

| Variable          | List or Scalar | Description                                                          |
|-------------------|----------------|----------------------------------------------------------------------|
| CFG{opt}{install} | Scalar         | Required<br>Installs SF Oracle RAC packages.                         |
| CFG{opt}{systems} | List           | Required<br>List of systems on which the product is to be installed. |
| CFG{opt}{upi}     | Scalar         | Required for the installer<br>Defines the product to be installed.   |

**Table 28-1** Response file variables for installing SF Oracle RAC (*continued*)

| Variable                                                                                     | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------------------------------------------------------------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{installallpkgs}<br>or<br>CFG{opt}{installrecpkgs}<br>or<br>CFG{opt}{installminpkgs} | Scalar         | <p>Required</p> <p>Instructs the installer to install SF Oracle RAC packages based on the variable that has the value set to 1:</p> <ul style="list-style-type: none"> <li>■ installallpkgs: Installs all packages</li> <li>■ installrecpkgs: Installs recommended packages</li> <li>■ installminpkgs: Installs minimum packages</li> </ul> <p><b>Note:</b> Set only one of these variable values to 1. In addition to setting the value of one of these variables, you must set the variable <i>\$CFG{opt}{install}</i> to 1.</p> |
| CFG{opt}{rsh}                                                                                | Scalar         | <p>Optional</p> <p>Defines that remote shell must be used instead of ssh as the communication method between systems.</p>                                                                                                                                                                                                                                                                                                                                                                                                          |
| CFG{opt}{gco}                                                                                | Scalar         | <p>Optional</p> <p>Defines that the installer must enable the global cluster option. You must set this variable value to 1 if you want to configure global clusters.</p>                                                                                                                                                                                                                                                                                                                                                           |
| CFG{opt}{keyfile}                                                                            | Scalar         | <p>Optional</p> <p>Defines the location of an ssh keyfile that is used to communicate with all remote systems.</p>                                                                                                                                                                                                                                                                                                                                                                                                                 |
| CFG{opt}{patchpath}                                                                          | Scalar         | <p>Optional</p> <p>Defines a location, typically an NFS mount, from which all remote systems can install product patches. The location must be accessible from all target systems.</p>                                                                                                                                                                                                                                                                                                                                             |

**Table 28-1** Response file variables for installing SF Oracle RAC (*continued*)

| Variable                            | List or Scalar | Description                                                                                                                                                                                                                                                     |
|-------------------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{pkgpath}                   | Scalar         | Optional<br><br>Defines a location, typically an NFS mount, from which all remote systems can install product packages. The location must be accessible from all target systems.                                                                                |
| CFG{opt}{tmppath}                   | Scalar         | Optional<br><br>Defines the location where a working directory is created to store temporary files and the depots that are needed during the install. The default location is /var/tmp.                                                                         |
| CFG{opt}{logpath}                   | Scalar         | Optional<br><br>Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.<br><br><b>Note:</b> The installer copies the response files and summary files also to the specified <i>logpath</i> location.        |
| CFG{opt}{donotinstall}<br>{package} | List           | Optional<br><br>Instructs the installation to not install the optional packages in the list.                                                                                                                                                                    |
| \$CFG{opt}{vxkeyless}               | Scalar         | Optional<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates keyless installation.<br><br>The value 0 indicates that a license key is required for the product. You must set the variable <i>\$CFG{keys}{hostname}</i> with appropriate values. |

**Table 28-1** Response file variables for installing SF Oracle RAC *(continued)*

| Variable                | List or Scalar | Description                                                                                                                                                                                                                                       |
|-------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{keys}<br>{hostname} | Scalar         | Optional<br><br>List of keys to be registered on the system if the variable <i>\$CFG{opt}{vxkeyless}</i> is set to 0.                                                                                                                             |
| \$CFG{uploadlogs}       | Scalar         | Optional<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the installation logs are uploaded to the Symantec Web site.<br><br>The value 0 indicates that the installation logs are not uploaded to the Symantec Web site. |

# Response file variables to configure SF Oracle RAC

Table 28-2 lists the response file variables that you can define to configure SF Oracle RAC.

**Table 28-2** Response file variables specific to configuring SF Oracle RAC

| Variable                          | List or Scalar | Description                                                                                                                                                         |
|-----------------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{configure}               | Scalar         | Performs the configuration if the packages are already installed.<br><br>(Required)                                                                                 |
| \$CFG{config_sfrac_subcomponents} | Scalar         | Set the variable to 1 to configure the SF Oracle RAC components.<br><br>(Required)<br><br><b>Note:</b> You must set the <i>\$CFG{opt}{configure}</i> variable to 1. |
| CFG{accepteula}                   | Scalar         | Specifies whether you agree with <i>EULA.pdf</i> on the media.<br><br>(Required)                                                                                    |

**Table 28-2** Response file variables specific to configuring SF Oracle RAC  
*(continued)*

| Variable          | List or Scalar | Description                                                                                                                                                                                                                                                      |
|-------------------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{systems}      | List           | List of systems on which the product is to be configured.<br>(Required)                                                                                                                                                                                          |
| CFG{prod}         | Scalar         | Defines the product to be configured.<br>(Required)                                                                                                                                                                                                              |
| CFG{opt}{keyfile} | Scalar         | Defines the location of an ssh keyfile that is used to communicate with all remote systems.<br>(Optional)                                                                                                                                                        |
| CFG{opt}{rsh}     | Scalar         | Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems.<br>(Optional)                                                                                                                                            |
| CFG{opt}{logpath} | Scalar         | Mentions the location where the log files are to be copied. The default location is <code>/opt/VRTS/install/logs</code> .<br><b>Note:</b> The installer copies the response files and summary files also to the specified <i>logpath</i> location.<br>(Optional) |
| \$CFG{uploadlogs} | Scalar         | Defines Boolean value 0 or 1.<br><br>The value 1 indicates that the installation logs are uploaded to the Symantec Web site.<br><br>The value 0 indicates that the installation logs are not uploaded to the Symantec Web site.<br>(Optional)                    |

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group (the `csgnic`, `csgvip`, and `csgnetmask` variables) must be defined if any are defined. The same is true for the SMTP notification (the `smtpserver`, `smtprecp`, and `smtpsrsev` variables), the SNMP trap notification (the `snmppport`, `snmpcons`, and `snmpcsev` variables), and the Global Cluster Option (the `gconic`, `gcovip`, and `gconetmask` variables).

[Table 28-3](#) lists the response file variables that specify the required information to configure a basic SF Oracle RAC cluster.

**Table 28-3**            Response file variables specific to configuring a basic SF Oracle RAC cluster

| Variable             | List or Scalar | Description                                                                                                                                        |
|----------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_clusterid}   | Scalar         | An integer between 0 and 65535 that uniquely identifies the cluster.<br>(Required)                                                                 |
| CFG{vcs_clustername} | Scalar         | Defines the name of the cluster.<br>(Required)                                                                                                     |
| CFG{vcs_allowcomms}  | Scalar         | Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start).<br>(Required) |

[Table 28-4](#) lists the response file variables that specify the required information to configure LLT over Ethernet.

**Table 28-4** Response file variables specific to configuring private LLT over Ethernet

| Variable                              | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_lltlink#}<br>{"system"}       | Scalar         | <p>Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links.</p> <p>You must enclose the system name within double quotes.</p> <p>(Required)</p>                                                                                                                                                                                         |
| CFG{vcs_lltlinklowpri#}<br>{"system"} | Scalar         | <p>Defines a low-priority heartbeat link. Typically, lltlinklowpri is used on a public network link to provide an additional layer of communication.</p> <p>If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, lltlinklowpri1, lltlinklowpri2, and so on.</p> <p>You must enclose the system name within double quotes.</p> <p>(Optional)</p> |

[Table 28-5](#) lists the response file variables that specify the required information to configure LLT over UDP.

**Table 28-5** Response file variables specific to configuring LLT over UDP

| Variable          | List or Scalar | Description                                                                                |
|-------------------|----------------|--------------------------------------------------------------------------------------------|
| CFG{lltoverudp}=1 | Scalar         | <p>Indicates whether to configure heartbeat link using LLT over UDP.</p> <p>(Required)</p> |

**Table 28-5** Response file variables specific to configuring LLT over UDP  
*(continued)*

| Variable                                             | List or Scalar | Description                                                                                                                                                                                                                                                                            |
|------------------------------------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_udplink<n>_address}<br>{<system1>}           | Scalar         | Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1.<br><br>You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links.<br><br>(Required)                                              |
| CFG<br>{vcs_udplinklowpri<n>_address}<br>{<system1>} | Scalar         | Stores the IP address (IPv4 or IPv6) that the low-priority heartbeat link uses on node1.<br><br>You can have four low-priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low-priority heartbeat links.<br><br>(Required)       |
| CFG{vcs_udplink<n>_port}<br>{<system1>}              | Scalar         | Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1.<br><br>You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links.<br><br>(Required)                                        |
| CFG{vcs_udplinklowpri<n>_port}<br>{<system1>}        | Scalar         | Stores the UDP port (16-bit integer value) that the low-priority heartbeat link uses on node1.<br><br>You can have four low-priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low-priority heartbeat links.<br><br>(Required) |
| CFG{vcs_udplink<n>_netmask}<br>{<system1>}           | Scalar         | Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1.<br><br>You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links.<br><br>(Required)                                              |



**Table 28-5** Response file variables specific to configuring LLT over UDP  
*(continued)*

| Variable                                         | List or Scalar | Description                                                                                                                                                                                                                                                                      |
|--------------------------------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_udplinklowpri<n>_netmask}<br>{<system1>} | Scalar         | Stores the netmask (prefix for IPv6) that the low-priority heartbeat link uses on node1.<br><br>You can have four low-priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low-priority heartbeat links.<br><br>(Required) |

[Table 28-6](#) lists the response file variables that specify the required information to configure virtual IP for SF Oracle RAC cluster.

**Table 28-6** Response file variables specific to configuring virtual IP for SF Oracle RAC cluster

| Variable                    | List or Scalar | Description                                                                                                                                |
|-----------------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_csgnic}<br>{system} | Scalar         | Defines the NIC device to use on a system. You can enter 'all' as a system value if the same NIC is used on all systems.<br><br>(Optional) |
| CFG{vcs_csgvip}             | Scalar         | Defines the virtual IP address for the cluster.<br><br>(Optional)                                                                          |
| CFG{vcs_csgnetmask}         | Scalar         | Defines the Netmask of the virtual IP address for the cluster.<br><br>(Optional)                                                           |

[Table 28-7](#) lists the response file variables that specify the required information to configure the SF Oracle RAC cluster in secure mode.

**Table 28-7** Response file variables specific to configuring SF Oracle RAC cluster in secure mode

| Variable                         | List or Scalar | Description                                                                                                                                         |
|----------------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{at_rootdomain}               | Scalar         | Defines the name of the system where the root broker is installed.<br>(Optional)                                                                    |
| CFG{at_rootbroker}               | Scalar         | Defines the root broker's name.                                                                                                                     |
| CFG{vcs_securitymenuopt}         | Scalar         | Specifies the menu option to choose to configure the cluster in secure mode.<br><br>■ 1—Automatic<br>■ 2—Semi-automatic<br>■ 3—Manual<br>(Optional) |
| CFG{vcs_vssdefport}              | Scalar         | Specifies the default port address of the root broker.<br>(Optional)                                                                                |
| CFG{vcs_roothashpath}            | Scalar         | Specifies the path of the root hash file.<br>(Optional)                                                                                             |
| CFG{vcs_ab_prplname}<br>{system} | Scalar         | Specifies the authentication broker's principal name on system.<br>(Optional)                                                                       |
| CFG{vcs_ab_password}<br>{system} | Scalar         | Specifies the authentication broker's password on system.<br>(Optional)                                                                             |
| CFG{vcs_blobpath}<br>{system}    | Scalar         | Specifies the path of the encrypted BLOB file for system.<br>(Optional)                                                                             |

**Table 28-8** lists the response file variables that specify the required information to configure VCS users.

**Table 28-8** Response file variables specific to configuring VCS users

| Variable          | List or Scalar | Description                                                                                                                                                                                                                                                                 |
|-------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_userenpw} | List           | <p>List of encoded passwords for VCS users.</p> <p>The value in the list can be "Administrators Operators Guests."</p> <p><b>Note:</b> The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list.</p> <p>(Optional)</p> |
| CFG{vcs_username} | List           | <p>List of names of VCS users.</p> <p>(Optional)</p>                                                                                                                                                                                                                        |
| CFG{vcs_userpriv} | List           | <p>List of privileges for VCS users.</p> <p><b>Note:</b> The order of the values for the vcs_userpriv list must match the order of the values in the vcs_username list.</p> <p>(Optional)</p>                                                                               |

[Table 28-9](#) lists the response file variables that specify the required information to configure VCS notifications using SMTP.

**Table 28-9** Response file variables specific to configuring VCS notifications using SMTP

| Variable            | List or Scalar | Description                                                                                                                                        |
|---------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_smtpserver} | Scalar         | <p>Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification.</p> <p>(Optional)</p> |
| CFG{vcs_smtprecpl}  | List           | <p>List of full email addresses (example: user@symantecexample.com) of SMTP recipients.</p> <p>(Optional)</p>                                      |

**Table 28-9** Response file variables specific to configuring VCS notifications using SMTP (*continued*)

| Variable         | List or Scalar | Description                                                                                                                                                                                                                                            |
|------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_smtpsev} | List           | Defines the minimum severity level of messages (Information, Warning, Error, and SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients.<br>(Optional) |

[Table 28-10](#) lists the response file variables that specify the required information to configure VCS notifications using SNMP.

**Table 28-10** Response file variables specific to configuring VCS notifications using SNMP

| Variable          | List or Scalar | Description                                                                                                                                                                                                                                       |
|-------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_snmpport} | Scalar         | Defines the SNMP trap daemon port (default=162).<br>(Optional)                                                                                                                                                                                    |
| CFG{vcs_snmpcons} | List           | List of SNMP console system names.<br>(Optional)                                                                                                                                                                                                  |
| CFG{vcs_snmpcsev} | List           | Defines the minimum severity level of messages (Information, Warning, Error, and SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names.<br>(Optional) |

[Table 28-11](#) lists the response file variables that specify the required information to configure SF Oracle RAC global clusters.

**Table 28-11** Response file variables specific to configuring SF Oracle RAC global clusters

| Variable                    | List or Scalar | Description                                                                                                                                                         |
|-----------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_gconic}<br>{system} | Scalar         | Defines the NIC for the Virtual IP that the Global Cluster Option uses. You can enter 'all' as a system value if the same NIC is used on all systems.<br>(Optional) |
| CFG{vcs_gcovip}             | Scalar         | Defines the virtual IP address to that the Global Cluster Option uses.<br>(Optional)                                                                                |
| CFG{vcs_gconetmask}         | Scalar         | Defines the Netmask of the virtual IP address that the Global Cluster Option uses.<br>(Optional)                                                                    |

## Response file variables to configure disk-based I/O fencing

[Table 28-12](#) lists the response file variables that specify the required information to configure disk-based I/O fencing for SF Oracle RAC.

**Table 28-12** Response file variables specific to configuring disk-based I/O fencing

| Variable                         | List or Scalar | Description                                                                                                                                                                                                                           |
|----------------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{fencing}                | Scalar         | Performs the I/O fencing configuration.<br>(Required)                                                                                                                                                                                 |
| CFG{vxfen_config_fencing_option} | Scalar         | Specifies the I/O fencing configuration mode. <ul style="list-style-type: none"> <li>■ 1—Coordination Point Server-based I/O fencing</li> <li>■ 2—Coordinator disk-based I/O fencing</li> <li>■ 3—Disabled mode</li> </ul> (Required) |

Table 28-12

Response file variables specific to configuring disk-based I/O fencing

(continued)

| Variable                              | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG {vxfen_config_fencing_mechanism}  | Scalar         | <p>Specifies the I/O fencing mechanism.</p> <p>This variable is not required if you had configured fencing in disabled mode. For disk-based fencing, you must configure the vxfen_config_fencing_mechanism variable and either the vxfen_config_fencing_dg variable or the vxfen_config_fencing_newdg_disks variable.</p> <p>(Optional)</p>               |
| CFG{vxfen_config_fencing_dg}          | Scalar         | <p>Specifies the disk group for I/O fencing.</p> <p>(Optional)</p> <p><b>Note:</b> You must define the vxfen_config_fencing_dg variable to use an existing disk group. If you want to create a new disk group, you must use both the vxfen_config_fencing_dg variable and the vxfen_config_fencing_newdg_disks variable.</p>                              |
| CFG{vxfen_config_fencing_newdg_disks} | List           | <p>Specifies the disks to use to create a new disk group for I/O fencing.</p> <p>(Optional)</p> <p><b>Note:</b> You must define the vxfen_config_fencing_dg variable to use an existing disk group. If you want to create a new disk group, you must use both the vxfen_config_fencing_dg variable and the vxfen_config_fencing_newdg_disks variable.</p> |

# Response file variables to configure server-based I/O fencing

You can use a CP server response file to configure server-based customized I/O fencing. The installer uses the CP server response file for the following types of I/O fencing configurations:

- Client cluster fencing (server-based I/O fencing configuration itself)  
The installer configures server-based customized I/O fencing on the SF Oracle RAC cluster without prompting for user input.
- Disk-based fencing with the disk group already created  
The installer configures fencing in disk-based mode on the SF Oracle RAC cluster without prompting for user input.  
Disk-based fencing configuration is one in which SCSI-3 disks are used as the only coordination points.  
Disk-based fencing with the disk group already created means that the disk group consisting of the coordinating disks already exists on the SF Oracle RAC cluster nodes.
- Disk-based fencing with the disk group to be created  
The installer creates the disk group and configures fencing properly on all the nodes in the SF Oracle RAC cluster without user intervention.  
Disk-based fencing with the disk group to be created means that the disk group does not exist yet, but will be created with the disks mentioned as coordination point.

[Table 28-13](#) lists the fields in the response file that are relevant for server-based customized I/O fencing.

**Table 28-13** CP server response file definitions

| Response file field        | Definition                                                                                                                                                                                                                                                                                                                    |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fencing_cpc_config_cpagent | <p>Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not.</p> <p>Enter "0" if you do not want to configure the Coordination Point agent using the installer.</p> <p>Enter "1" if you want to use the installer to configure the Coordination Point agent.</p> |

**Table 28-13** CP server response file definitions (*continued*)

| Response file field    | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fencing_cpc_cpagentgrp | Name of the service group which will have the Coordination Point agent resource as part of it.<br><br><b>Note:</b> This field is obsolete if the <code>fencing_cpc_config_cpagent</code> field is given a value of '0'.                                                                                                                                                                                                                                                                                                                    |
| fencing_cpc_cps        | Virtual IP address or Virtual hostname of the CP servers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| fencing_cpc_reusedg    | This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks).<br><br>Enter either a "1" or "0".<br><br>Entering a "1" indicates reuse, and entering a "0" indicates do not reuse.<br><br>When reusing an existing DG name for the mixed mode fencing configuration, you need to manually add a line of text, such as<br>"\$CFG{fencing_cpc_reusedg}=0" or<br>"\$CFG{fencing_cpc_reusedg}=1" before proceeding with a silent installation. |
| fencing_cpc_dgname     | The name of the disk group to be used in the customized fencing, where at least one disk is being used.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| fencing_cpc_diffab     | This response field indicates whether the CP servers and SF Oracle RAC clusters use different root brokers.<br><br>Entering a "1" indicates that they are using different root brokers.<br><br>Entering a "0" indicates that they are not using different root brokers.                                                                                                                                                                                                                                                                    |
| fencing_cpc_disks      | The disks being used as coordination points if any.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| fencing_cpc_ncps       | Total number of coordination points being used, including both CP servers and disks.                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| fencing_cpc_ndisks     | The number of disks being used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| fencing_cpc_ports      | The port of the CP server that is denoted by <i>cps</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



**Table 28-13** CP server response file definitions (*continued*)

| Response file field   | Definition                                                                                                                                                                    |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fencing_cpc_ccab      | The name of the authentication broker (AB) for any one of the SF Oracle RAC cluster nodes.                                                                                    |
| fencing_cpc_cpsabport | The port at which the authentication broker (AB) mentioned above listens for authentication..                                                                                 |
| fencing_cpc_ccabport  | The port at which the authentication broker (AB) mentioned above listens for authentication.                                                                                  |
| fencing_cpc_mechanism | The disk mechanism that is used by customized fencing.<br>The value for this field is either "raw" or "dmp"                                                                   |
| fencing_cpc_cpsab     | The name of the authentication broker (AB) for any one of the CP servers.                                                                                                     |
| fencing_cpc_security  | This field indicates whether security is enabled or not<br>Entering a "1" indicates that security is enabled.<br>Entering a "0" indicates that security has not been enabled. |

## Response file variable definitions for Oracle RAC

The variable definitions for Oracle RAC are grouped in tabular format for the following Oracle tasks:

|                                                                        |                                              |
|------------------------------------------------------------------------|----------------------------------------------|
| Creating Oracle user and group                                         | See <a href="#">Table 28-14</a> on page 474. |
| Creating storage for OCR and voting disk                               | See <a href="#">Table 28-15</a> on page 475. |
| Configuring the private IP address and PrivNIC resource under VCS      | See <a href="#">Table 28-16</a> on page 478. |
| Configuring the private IP address and MultiPrivNIC resource under VCS | See <a href="#">Table 28-17</a> on page 482. |
| Installing Oracle Clusterware                                          | See <a href="#">Table 28-18</a> on page 485. |
| Installing Oracle database                                             | See <a href="#">Table 28-19</a> on page 486. |
| Configuring CSSD resource                                              | See <a href="#">Table 28-20</a> on page 487. |
| Relinking Oracle RAC libraries                                         | See <a href="#">Table 28-21</a> on page 488. |

**Note:** Some of the variable definitions may occur in multiple sections, for example \$CFG{oracle\_user}. These variables need not be repeated if all the tasks are performed as a single installation activity. However, if you perform these tasks independently, make sure that all the required variables, as indicated in the table for each task, are supplied in the response file.

Table 28-14 lists the variables that are used to create the Oracle user and group.

**Table 28-14** Variables for creating Oracle user and group

| Variable                        | List or Scalar | Description                                                                                                                                                                                           |
|---------------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{create_oracle_user_group} | Scalar         | Required<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that Oracle user and group will be created.<br><br>The value 0 indicates that Oracle user and group will not be created. |
| \$CFG{grid_user}                | Scalar         | Required<br><br>Defines the name of the grid user.                                                                                                                                                    |
| \$CFG{oracle_user}              | Scalar         | Required<br><br>Defines the name of the Oracle user.                                                                                                                                                  |
| \$CFG{oracle_uid}               | Scalar         | Required<br><br>Defines the user ID of the Oracle user.                                                                                                                                               |
| \$CFG{oracle_group}             | Scalar         | Required<br><br>Defines the primary group of the Oracle user.                                                                                                                                         |
| \$CFG{oracle_gid}               | Scalar         | Required<br><br>Defines the group ID of the Oracle user.                                                                                                                                              |
| \$CFG{oracle_user_home}         | Scalar         | Required<br><br>Defines the full path of the Oracle user's home directory.                                                                                                                            |

**Table 28-14** Variables for creating Oracle user and group (*continued*)

| Variable                      | List or Scalar | Description                                                                                                                                                                                             |
|-------------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{oracle_secondary_group} | List           | Optional<br>Defines the list of secondary groups for the Oracle user.                                                                                                                                   |
| \$CFG{oracle_secondary_gid}   | List           | Optional<br>Defines the list of secondary group IDs for the Oracle user. The elements of this variable must be in the same order as that of the elements in the variable \$CFG{oracle_secondary_group}. |

[Table 28-15](#) lists the variables that are used to create the storage for OCR and voting disk.

**Table 28-15** Variables for creating storage for OCR and voting disk

| Variable                       | List or Scalar | Description                                                                                                                                                                                                                   |
|--------------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{create_ocr_vote_storage} | Scalar         | Required<br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the storage for OCR and voting disk will be created.<br><br>The value 0 indicates that the storage for OCR and voting disk will not be created. |

**Table 28-15** Variables for creating storage for OCR and voting disk (*continued*)

| Variable               | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{ocrvotedgoption} | Scalar         | <p>Required</p> <p>Defines a Boolean value 0 or 1.</p> <p>The value 1 indicates that an existing disk group will be used to create the storage for OCR and voting disk.</p> <p><b>Note:</b> If you choose to use an existing disk group, use the \$CFG{ocrvotedgname} variable to specify the name of an existing disk group that has a minimum of two disks (for mirroring).</p> <p>The value 0 indicates that a new disk group will be created for OCR and voting disk storage.</p> <p><b>Note:</b> If you choose to create a disk group, you must set the following variables:<br/> \$CFG{ocrvotedisks},<br/> \$CFG{ocrvotedgname}</p> |
| \$CFG{ocrvotescheme}   | Scalar         | <p>Required</p> <p>Defines the storage scheme to be used for OCR and voting disk.</p> <p>The value 1 indicates Clustered File System.</p> <p>The value 0 indicates CVM raw volumes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| \$CFG{ocrvotedisks}    | List           | <p>Required</p> <p>Defines the list of shared disks to be used for OCR and voting disk.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| \$CFG{ocrvotedgname}   | Scalar         | <p>Required</p> <p>Defines the name of the disk group to be used for OCR and voting disk.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Table 28-15** Variables for creating storage for OCR and voting disk (*continued*)

| Variable              | List or Scalar | Description                                                                                                                                                          |
|-----------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{ocrvotevolname} | Scalar         | Required<br><br>Defines the volume name for OCR and voting disk. This variable must be used only if you have set the storage scheme to 1 (Clustered File System).    |
| \$CFG{ocrvotevolsize} | Scalar         | Required<br><br>Defines the size of the OCR and voting disk volume. This variable must be used only if you have set the storage scheme to 1 (Clustered File System). |
| \$CFG{ocrvotemount}   | Scalar         | Required<br><br>Defines the full path to the CFS mount point. This variable must be used only if you have set the storage scheme to 1 (Clustered File System).       |
| \$CFG{ocrvolname}     | Scalar         | Required<br><br>Defines the volume name for OCR. This variable must be used only if you have set the storage scheme to 0 (CVM Raw Volumes).                          |
| \$CFG{ocrvolsize}     | Scalar         | Required<br><br>Defines the size of the OCR volume. This variable must be used only if you have set the storage scheme to 0 (CVM Raw Volumes).                       |
| \$CFG{votevolname}    | Scalar         | Required<br><br>Defines the volume name for voting disk. This variable must be used only if you have set the storage scheme to 0 (CVM Raw Volumes).                  |

**Table 28-15** Variables for creating storage for OCR and voting disk (*continued*)

| Variable            | List or Scalar | Description                                                                                                                                            |
|---------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{votevolsize}  | Scalar         | Required<br><br>Defines the size of the voting disk volume. This variable must be used only if you have set the storage scheme to 0 (CVM Raw Volumes). |
| \$CFG{oracle_user}  | Scalar         | Required<br><br>Defines the name of the Oracle user.                                                                                                   |
| \$CFG{oracle_group} | Scalar         | Required<br><br>Defines the primary group of the Oracle user.                                                                                          |

[Table 28-16](#) lists the variables that are used to configure the private IP address and PrivNIC resource under VCS.

**Table 28-16** Variables for configuring the private IP address and PrivNIC resource under VCS

| Variable               | List or Scalar | Description                                                                                                                                                                                                                                                                                                 |
|------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{config_privnic}  | Scalar         | Required<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the PrivNIC and private IP address information will be configured for Oracle Clusterware.<br><br>The value 0 indicates that the PrivNIC and private IP address information will not be configured for Oracle Clusterware. |
| \$CFG{privnic_resname} | Scalar         | Required<br><br>Defines the PrivNIC resource name in the <code>main.cf</code> file.                                                                                                                                                                                                                         |

**Table 28-16** Variables for configuring the private IP address and PrivNIC resource under VCS (*continued*)

| Variable                          | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{privnic_interface_priority} | String         | <p>Required</p> <p>Defines the priority that determines which NIC will be used in the event of a failover. Set the priority in decreasing order.</p> <p>For example, the following priority setting indicates that bge2 will be given priority in the event of a failover:</p> <p><b>\$CFG{privnic_interface_priority}="bge2 bge3";</b></p>                                       |
| \$CFG{host1}{privnicip}           | Scalar         | <p>Required</p> <p>Defines the IP address to be configured for the PrivNIC resource on the node.</p> <p>Repeat this variable for each node in the cluster. For example, if you have two nodes in the cluster, you must provide this variable for each node.</p> <p>For example:</p> <pre>\$CFG{galaxy}{privnicip} ="192.168.12.1"  \$CFG{nebula}{privnicip} ="192.168.12.2"</pre> |
| \$CFG{nic_reuseip}                | Scalar         | <p>Required</p> <p>Defines a Boolean value 0 or 1.</p> <p>The value 1 indicates that the existing IP addresses in the <code>/etc/hosts</code> or <code>/etc/inet/ipnodes</code> files will be used.</p> <p>The value 0 indicates that the IP addresses will not be reused.</p>                                                                                                    |

**Table 28-16** Variables for configuring the private IP address and PrivNIC resource under VCS (*continued*)

| Variable                                  | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>\$CFG{host}{hostname_for_ip}</code> | Scalar         | <p>Required</p> <p>Defines the private node name of the IP address (<i>hostname_for_ip</i>) for the PrivNIC resource and the node (<i>system</i>) for which the resource is configured.</p> <p>Repeat this variable for each node in the cluster. For example, if you have two nodes in the cluster, you must provide this variable for each node.</p> <p>For example:</p> <pre>\$CFG{galaxy}{hostname_for_ip} ="galaxy-priv"  \$CFG{nebula}{hostname_for_ip} ="galaxy-priv"</pre>                                                                    |
| <code>\$CFG{nic_netmask}</code>           | Scalar         | <p>Required</p> <p>Defines the netmask for the private network.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <code>\$CFG{nic_add_ip_to_files}</code>   | Scalar         | <p>Required</p> <p>Defines a Boolean value 0 or 1.</p> <p>The value 1 indicates that the IP addresses are added to the <code>/etc/hosts</code> file or the <code>/etc/inet/ipnodes</code> file.</p> <p><b>Note:</b> Make sure that the IP addresses for the NIC resource are not already present in the files or set the <code>\$CFG{nic_reuseip}</code> and <code>\$CFG{nic_reusealias}</code> variables, otherwise the network configuration step fails.</p> <p>The value 0 indicates that the IP addresses may already be present in the file.</p> |



**Table 28-16** Variables for configuring the private IP address and PrivNIC resource under VCS (*continued*)

| Variable                                 | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{nic_reconfigure_existing_resource} | Scalar         | <p>Optional</p> <p>Defines a Boolean value 0 or 1.</p> <p>The value 1 indicates that the existing PrivNIC resource in the <code>main.cf</code> file will be deleted and reconfigured.</p> <p>The value 0 indicates that the existing PrivNIC resource in the <code>main.cf</code> file will be reused.</p>                                                                                                                                                                                                                                                                                                                               |
| \$CFG{nic_reusealias}                    | Scalar         | <p>Required</p> <p>Defines a Boolean value 0 or 1.</p> <p>The value 1 indicates that the installer will not check the <code>/etc/hosts</code> file to determine whether the host name alias for the private IP addresses exist or not. The installer assumes that the host names alias information is present in the file. Make sure that the alias information is present in the file.</p> <p>The value 0 indicates that the installer checks whether the host name alias information is present in the <code>/etc/hosts</code> file. Make sure that the alias information is present in the file otherwise the installation fails.</p> |

**Table 28-17** lists the variables that are used to configure the private IP addresses and the MultiPrivNIC resource under VCS.

**Table 28-17** Variables for configuring the private IP addresses and the MultiPrivNIC resource under VCS

| Variable                    | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{config_multiprivnic}  | Scalar         | <p>Required</p> <p>Defines a Boolean value 0 or 1.</p> <p>The value 1 indicates that the MultiPrivNIC and private IP address information will be configured for Oracle Clusterware.</p> <p>The value 0 indicates that the MultiPrivNIC and private IP address information will not be configured for Oracle Clusterware.</p>                                                                                                                                                                                                                          |
| \$CFG{multiprivnic_resname} | Scalar         | <p>Required</p> <p>Defines the MultiPrivNIC resource name in the <code>main.cf</code> file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| \$CFG{nic_add_ip_to_files}  | Scalar         | <p>Required</p> <p>Defines a Boolean value 0 or 1.</p> <p>The value 1 indicates that the IP addresses are added to the <code>/etc/hosts</code> file or the <code>/etc/inet/ipnodes</code> file.</p> <p><b>Note:</b> Make sure that the IP addresses for the NIC resource are not already present in the files or set the <code>\$CFG{nic_reuseip}</code> and <code>\$CFG{nic_reusealias}</code> variables, otherwise the network configuration step fails.</p> <p>The value 0 indicates that the IP addresses may already be present in the file.</p> |

**Table 28-17** Variables for configuring the private IP addresses and the MultiPrivNIC resource under VCS (*continued*)

| Variable                                                  | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{ <i>host1</i> }->{ <i>NIC1</i> }->{multiprivnicip}  | List           | <p>Required</p> <p>Defines the list of IP addresses for the MultiPrivNIC resource.</p> <p><b>Note:</b> The private IP addresses must be configured for each node and each interface in the cluster.</p> <p>For example, if you have two nodes galaxy and nebula in the cluster:</p> <pre>\$CFG{galaxy}{bge1} {multiprivnicip}="192.168.12.1"; \$CFG{galaxy}{bge2} {multiprivnicip}=="192.168.2.1"; \$CFG{nebula}{bge1} {multiprivnicip}="192.168.12.2"; \$CFG{nebula}{bge2} {multiprivnicip}="192.168.2.2";</pre>                                                                                                                            |
| \$CFG{ <i>host1</i> }->{ <i>NIC1</i> }->{hostname_for_ip} | List           | <p>Required</p> <p>Defines the list of private node names (<i>hostname_for_ip</i>) for the IP addresses configured in the MultiPrivNIC resource for the interface (<i>inf</i>) on the node (<i>system</i>) .</p> <p><b>Note:</b> The private IP address must be configured for each node and each interface in the cluster.</p> <p>For example, if you have two nodes galaxy and nebula in the cluster:</p> <pre>\$CFG{galaxy}{bge1} {hostname_for_ip}="galaxy-priv"; \$CFG{galaxy}{bge2} {hostname_for_ip}="galaxy-priv1"; \$CFG{nebula}{bge1} {hostname_for_ip}="nebula-priv"; \$CFG{nebula}{bge2} {hostname_for_ip}="nebula-priv1";</pre> |

**Table 28-17** Variables for configuring the private IP addresses and the MultiPrivNIC resource under VCS (*continued*)

| Variable                                 | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{nic_netmask}                       | Scalar         | Required<br><br>Defines the netmask for the private network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| \$CFG{nic_reconfigure_existing_resource} | Scalar         | Optional<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the existing MultiPrivNIC resource in the <code>main.cf</code> file will be deleted and reconfigured.<br><br>The value 0 indicates that the existing MultiPrivNIC resource in the <code>main.cf</code> file will be reused.                                                                                                                                                                                                                                                                                                                     |
| \$CFG{nic_reuseip}                       | Scalar         | Required<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the existing IP addresses in the <code>/etc/hosts</code> or <code>/etc/inet/ipnodes</code> files will be used.<br><br>The value 0 indicates that the IP addresses will not be reused.                                                                                                                                                                                                                                                                                                                                                           |
| \$CFG{nic_reusealias}                    | Scalar         | Required<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the installer will not check the <code>/etc/hosts</code> file to determine whether the host name alias for the private IP addresses exist or not. The installer assumes that the host names alias information is present in the file. Make sure that the alias information is present in the file.<br><br>The value 0 indicates that the installer checks whether the host name alias information is present in the <code>/etc/hosts</code> file. Make sure that the alias information is present in the file otherwise the installation fails. |

[Table 28-18](#) lists the variables that are used to install Oracle Clusterware.

**Table 28-18** Variables for installing Oracle Clusterware

| Variable                          | List or Scalar | Description                                                                                                                                                                                           |
|-----------------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{install_oracle_clusterware} | Scalar         | Required<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that Oracle Clusterware will be configured.<br><br>The value 0 indicates that Oracle Clusterware will not be configured. |
| \$CFG{oracle_user}                | Scalar         | Required<br><br>Defines the name of the Oracle user.                                                                                                                                                  |
| \$CFG{oracle_group}               | Scalar         | Required<br><br>Defines the primary group of the Oracle user.                                                                                                                                         |
| \$CFG{oracle_base}                | Scalar         | Required<br><br>Defines the base directory for the Oracle RAC installation.                                                                                                                           |
| \$CFG{crs_home}                   | Scalar         | Required<br><br>Defines the Oracle Clusterware home directory.<br><br>The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle Clusterware response file.      |
| \$CFG{crs_installpath}            | Scalar         | Required<br><br>Defines the full path of the Oracle Clusterware installation binaries.                                                                                                                |
| \$CFG{oracle_version}             | Scalar         | Required<br><br>Defines the version of the Oracle RAC binaries (for example, 11.2.0.1.0). This definition is overridden if a different Oracle RAC version is detected during the installation.        |

**Table 28-18** Variables for installing Oracle Clusterware (*continued*)

| Variable                | List or Scalar | Description                                                                    |
|-------------------------|----------------|--------------------------------------------------------------------------------|
| \$CFG{crs_responsefile} | Scalar         | Required<br><br>Defines the full path of the Oracle Clusterware response file. |

[Table 28-19](#) lists the variables that are used to install Oracle database.

**Table 28-19** Variables for installing Oracle database

| Variable                       | List or Scalar | Description                                                                                                                                                                                                     |
|--------------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{install_oracle_database} | Scalar         | Required<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the Oracle RAC database will be configured.<br><br>The value 0 indicates that the Oracle RAC database will not be configured. |
| \$CFG{oracle_user}             | Scalar         | Required<br><br>Defines the name of the Oracle user.                                                                                                                                                            |
| \$CFG{oracle_group}            | Scalar         | Required<br><br>Defines the primary group of the Oracle user.                                                                                                                                                   |
| \$CFG{oracle_base}             | Scalar         | Required<br><br>Defines the base directory for the Oracle RAC installation.                                                                                                                                     |
| \$CFG{crs_home}                | Scalar         | Required<br><br>Defines the Oracle Clusterware home directory.<br><br>The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle Clusterware response file.                |

**Table 28-19** Variables for installing Oracle database (*continued*)

| Variable               | List or Scalar | Description                                                                                                                                                                                        |
|------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{db_home}         | Scalar         | Required<br><br>Defines the Oracle RAC database home directory.<br><br>The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle RAC database response file. |
| \$CFG{db_installpath}  | Scalar         | Required<br><br>Defines the full path of the Oracle RAC database installation binaries.                                                                                                            |
| \$CFG{oracle_version}  | Scalar         | Required<br><br>Defines the version of the Oracle RAC binaries (for example, 11.2.0.1.0). This definition is overridden if a different Oracle RAC version is detected during the installation.     |
| \$CFG{db_responsefile} | Scalar         | Required<br><br>Defines the full path of the Oracle database response file.                                                                                                                        |

[Table 28-20](#) lists the variables that are used to configure CSSD resource.

**Table 28-20** Variables for configuring CSSD resource

| Variable                 | List or Scalar | Description                                                                                                                                                                                                                                               |
|--------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{config_cssd_agent} | Scalar         | Required<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the CSSD agent will be configured after Oracle RAC installation.<br><br>The value 0 indicates that the CSSD agent will not be configured after Oracle RAC installation. |

**Table 28-20** Variables for configuring CSSD resource *(continued)*

| Variable                         | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{reconfigure_cssd_resource} | Scalar         | Required<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the SF Oracle RAC installer deletes the existing CSSD resource from the <code>main.cf</code> file and reconfigures it.<br><br>The value 0 indicates that the SF Oracle RAC installer does not delete and reconfigure the resource. The installer exits the task with the message that the resource exists. |

[Table 28-21](#) lists the variables that are used to relink Oracle RAC libraries.

**Table 28-21** Variables for relinking Oracle RAC libraries

| Variable                      | List or Scalar | Description                                                                                                                                                                                                                                                                                                                               |
|-------------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{relink_oracle_database} | Scalar         | Required<br><br>Defines a Boolean value 0 or 1.<br><br>The value 1 indicates that the SF Oracle RAC libraries will be relinked with the Oracle RAC database after Oracle RAC installation.<br><br>The value 0 indicates that the SF Oracle RAC libraries will not be relinked with the Oracle RAC database after Oracle RAC installation. |
| \$CFG{oracle_user}            | Scalar         | Required<br><br>Defines the name of the Oracle user.                                                                                                                                                                                                                                                                                      |
| \$CFG{oracle_group}           | Scalar         | Required<br><br>Defines the primary group of the Oracle user.                                                                                                                                                                                                                                                                             |



**Table 28-21** Variables for relinking Oracle RAC libraries (*continued*)

| Variable              | List or Scalar | Description                                                                                                                                                                                               |
|-----------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{crs_home}       | Scalar         | <p>Required</p> <p>Defines the Oracle Clusterware home directory.</p> <p>The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle Clusterware response file.</p>   |
| \$CFG{db_home}        | Scalar         | <p>Required</p> <p>Defines the Oracle RAC database home directory.</p> <p>The value in this variable must be the same as that of the 'ORACLE_HOME' variable in the Oracle RAC database response file.</p> |
| \$CFG{oracle_version} | Scalar         | <p>Required</p> <p>Defines the version of the Oracle RAC binaries (for example, 11.2.0.1.0). This definition is overridden if a different Oracle RAC version is detected during the installation.</p>     |



# Adding or removing nodes from an SF Oracle RAC cluster

- [Chapter 29. Adding a node to SF Oracle RAC clusters](#)
- [Chapter 30. Removing a node from SF Oracle RAC clusters](#)



# Adding a node to SF Oracle RAC clusters

This chapter includes the following topics:

- [About adding a node to a cluster](#)
- [Before adding a node to a cluster](#)
- [Preparing to add a node to a cluster](#)
- [Adding a node to a cluster](#)
- [Preparing the new node for installing Oracle RAC](#)
- [Adding the new node to Oracle RAC](#)
- [Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node](#)
- [Sample configuration file for adding a node to the cluster](#)

## About adding a node to a cluster

After you install SF Oracle RAC and create a cluster, you can add and remove nodes from the cluster. You can create clusters of up to 16 nodes.

You can add a node:

- Using the product installer
- Using the Web installer
- Manually

The example procedures describe how to add a node to an existing cluster with two nodes.

## Before adding a node to a cluster

Before preparing to add the node to an existing SF Oracle RAC cluster, verify the following:

- Hardware and software requirements are met.  
See [“Meeting hardware and software requirements”](#) on page 494.
- Hardware is set up for the new node.  
See [“Setting up the hardware”](#) on page 494.
- The existing cluster is an SF Oracle RAC cluster and that SF Oracle RAC is running on the cluster.
- The new system has the same identical operating system versions and patch levels as that of the existing cluster.


## Meeting hardware and software requirements

The system you add to the cluster must meet the hardware and software requirements.

## Setting up the hardware

[Figure 29-1](#) shows that before you configure a new system on an existing cluster, you must physically add the system to the cluster.

**Figure 29-1** Adding a node to a two-node cluster using two switches



### To set up the hardware

#### 1 Connect the SF Oracle RAC private Ethernet controllers.

Perform the following tasks as necessary:

- When you add nodes to a cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a cluster, so you might have to swap out the cable for a switch or hub.
- If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.

[Figure 29-1](#) illustrates a new node being added to an existing two-node cluster using two independent hubs.

#### 2 Make sure that you meet the following requirements:

- The node must be connected to the same shared storage devices as the existing nodes.
- The node must have private network connections to two independent switches for the cluster.

For more information, see the *Veritas Cluster Server Installation Guide*.

- The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster.

## Preparing to add a node to a cluster

Complete the following preparatory steps on the new node before you add the node to an existing SF Oracle RAC cluster.

**To prepare the new node**

- 1 Verify that the new node meets installation requirements.

```
# ./installsf rac -precheck saturn
```

- 2 Install SF Oracle RAC on the new system.

---

**Note:** Use the `-install` option to install SF Oracle RAC. Do not configure SF Oracle RAC after the installation.

```
Would you like to configure SF Oracle RAC on saturn [y, n, q] (n)
```

You can configure the new node later using the configuration from the existing cluster nodes.

---

See [“About installing SF Oracle RAC ”](#) on page 91.

## Adding a node to a cluster

You can use one of the following methods to add a node to an existing SF Oracle RAC cluster:

|                         |                                                                                                 |
|-------------------------|-------------------------------------------------------------------------------------------------|
| SF Oracle RAC installer | See <a href="#">“Adding a node to a cluster using the SF Oracle RAC installer”</a> on page 496. |
| Manual                  | See <a href="#">“Adding the node to a cluster manually”</a> on page 500.                        |

---

**Note:** Before you add the node, make sure that SF Oracle RAC is not configured on the node.

---

### Adding a node to a cluster using the SF Oracle RAC installer

You can add a node using the `-addnode` option with the SF Oracle RAC installer. The SF Oracle RAC installer performs the following tasks:



- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and packages installed on the new node.
- Discovers the network interfaces on the new node and checks the interface settings.
- Creates the following files on the new node:
  - `/etc/llttab`
  - `/etc/VRTSvcs/conf/sysname`
- Updates the following configuration files and copies them on the new node:
  - `/etc/llthosts`
  - `/etc/gabtab`
  - `/etc/VRTSvcs/conf/config/main.cf`
- Copies the following files from the existing cluster to the new node:
  - `/etc/vxfenmode`
  - `/etc/vxfendg`
  - `/etc/vcsmmtab`
  - `/etc/vx/.uuids/clusuuid`
  - `/etc/default/llt`
  - `/etc/default/gab`
  - `/etc/default/vxfen`
  - `/etc/default/vcsmm`
- Configures security on the new node if the existing cluster is a secure cluster.

---

**Warning:** If the root broker system has failed, then you must recover or reconfigure the root broker system before you add a new node to the cluster.

---

- Configures disk-based or server-based fencing depending on the fencing mode in use on the existing cluster.
- Adds the new node to the CVM, ClusterService, and VxSS service groups in the VCS configuration.

---

**Note:** For other service groups configured under VCS, update the configuration for the new node manually.

---

- Starts SF Oracle RAC processes and configures CVM and CFS on the new node.

At the end of the process, the new node joins the SF Oracle RAC cluster.

---

**Note:** If you have configured server-based fencing on the existing cluster, make sure that the CP server does not contain entries for the new node. If the CP server already contains entries for the new node, remove these entries before adding the node to the cluster, otherwise the process may fail with an error.

---

---

**Caution:** If you plan to use the SF Oracle RAC installer for completing the Oracle pre-installation tasks on the new node, do not quit the installer after adding the node to the cluster. If you quit the installer, you must perform the Oracle pre-installation tasks manually.

---

#### To add the node to an existing cluster using the installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the SF Oracle RAC installer with the `-addnode` option.

```
# cd /opt/VRTS/install  
  
# ./installsffrac -addnode
```

The installer displays the copyright message and the location where it stores the temporary installation logs.

- 3 Enter the name of a node in the existing SF Oracle RAC cluster. The installer uses the node information to identify the existing cluster.

```
Enter one node of the SF Oracle RAC cluster to which  
you would like to add one or more nodes: galaxy
```

- 4 Review and confirm the cluster information.
- 5 Enter the name of the systems that you want to add as new nodes to the cluster.

```
Enter the system names separated by spaces  
to add to the cluster: saturn
```

The installer checks the installed products and packages on the nodes and discovers the network interfaces.

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

---

**Note:** The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster. The LLT configuration for the new node must be the same as that of the existing cluster.

---

```
Enter the NIC for the first private heartbeat
link on saturn: [b,q,?] bge1
```

- 7 Enter **y** to configure a second private heartbeat link.

---

**Note:** At least two private heartbeat links must be configured for high availability of the cluster.

---

```
Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
```

- 8 Enter the name of the network interface that you want to configure as the second private heartbeat link.

```
Enter the NIC for the second private heartbeat link
on saturn: [b,q,?] bge2
```

- 9 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

- 10 Review and confirm the information.

- 11 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on saturn: bge3
```

- 12** If the existing cluster uses server-based fencing in secure mode, provide responses to the following installer prompts.

If you are using different root brokers for the CP server and the client SF Oracle RAC cluster, enter **y** to confirm the use of different root brokers. The installer attempts to establish trust between the new node being added to the cluster and the authentication broker of the CP server.

```
Are you using different Root Brokers for the CP Server(s) and the
client cluster? (If so then installer will try to establish trust
between the new node(s) being added and CP Server's
Authentication Broker) [y,n,q] (n) y
```

Enter the host name of the authentication broker used for any one of the CP servers.

```
Enter hostname of the Authentication Broker being used for any one
of the CP Server(s): [b] mycps1.symantecexample.com
```

Enter the port number where the authentication broker for the CP server listens to establish trust with the new node:

```
Enter the port where the Authentication Broker
mycps1.symantecexample.com for the CP Server(s) is listening
for establishing trust: [b] (2821)
```

The installer then starts all the required Veritas processes and joins the new node to cluster.

---

**Note:** Do not quit the installer if you want to perform the Oracle pre-installation tasks using the SF Oracle RAC installer.

---

- 13** Confirm using `lltstat -n` and `gabconfig -a`.

## Adding the node to a cluster manually

Perform this procedure after you install SF Oracle RAC only if you plan to add the node to the cluster manually.

### To add the node manually to the cluster

- 1 Start the Volume Manager.  
 See [“Starting Volume Manager on the new node”](#) on page 501.
- 2 Configure LLT, GAB, and VCSMM.  
 See [“Configuring LLT, GAB, and VCSMM on the new node”](#) on page 502.
- 3 If the existing cluster is a secure cluster, set up the new node to run in secure mode.  
 See [“Setting up the node to run in secure mode”](#) on page 504.
- 4 If the existing cluster is configured to use server-based I/O fencing, configure server-based I/O fencing on the new node.  
 See [“Starting fencing on the new node”](#) on page 509.
- 5 Start VCS.  
 See [“To start VCS on the new node”](#) on page 511.
- 6 Configure CVM and CFS.  
 See [“Configuring CVM and CFS on the new node”](#) on page 511.
- 7 If the ClusterService group is configured on the existing cluster, add the node to the group.  
 See [“Configuring the ClusterService group for the new node”](#) on page 512.

## Starting Volume Manager on the new node

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `installsf` program.

### To start Volume Manager on the new node

- 1 To start Veritas Volume Manager on the new node, use the `vxinstall` utility:  

```
# vxinstall
```
- 2 Enter **n** when prompted to set up a system wide disk group for the system.  
 The installation completes.
- 3 Verify that the daemons are up and running. Enter the command:

```
# vxdisk list
```

Make sure the output displays the shared disks without errors.

## Configuring LLT, GAB, and VCSMM on the new node

Perform the steps in the following procedure to configure LLT, GAB, and VCSMM on the new node.

### To configure LLT, GAB, and VCSMM on the new node

- 1 Edit the `/etc/llthosts` file on the existing nodes. Using vi or another text editor, add the line for the new node to the file. The file resembles:

```
0 galaxy
1 nebula
2 saturn
```

- 2 Copy the `/etc/llthosts` file from one of the existing systems over to the new system. The `/etc/llthosts` file must be identical on all nodes in the cluster.
- 3 Create an `/etc/llttab` file on the new system. For example:

```
set-node saturn
set-cluster 101

link bge1 /dev/bge:1 - ether - -
link bge2 /dev/bge:2 - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

- 4 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

```
/sbin/gabconfig -c -nN
```

Where N represents the number of systems in the cluster. For a three-system cluster, N would equal 3.

- 5 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 6 Copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/default/llt
/etc/default/gab
/etc/default/vcs
/etc/vcsmmtab
```

- 7 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 8 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \  
-from_sys galaxy -to_sys saturn
```

- 9 Start the LLT, GAB, VCSMM, LMX, and ODM drivers on the new node:

For Solaris 9:

```
# /etc/init.d/llt start  
# /etc/init.d/gab start  
# /etc/init.d/vcsmm start  
# /etc/init.d/lmx start  
# /etc/init.d/odm restart
```

For Solaris 10:

```
# svcadm enable llt  
# svcadm enable gab  
# svcadm enable vcsmm  
# svcadm enable lmx  
# svcadm restart odm
```

10 On the new node, verify that the GAB port memberships are a, d, and o:

```
# gabconfig -a

GAB Port Memberships
=====
Port a gen      df204 membership 012
Port d gen      df20a membership 012
Port o gen      df207 membership 012

Port h may also be present. In that case, run:

# hastop -local
```

Setting up the node to run in secure mode

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

Table 29-1 uses the following information for the following command examples.

Table 29-1 The command examples definitions

| Name   | Fully-qualified host name (FQHN) | Function                                         |
|--------|----------------------------------|--------------------------------------------------|
| saturn | saturn.nodes.example.com         | The new node that you are adding to the cluster. |
| RB1    | RB1.brokers.example.com          | The root broker for the cluster                  |
| RB2    | RB2.brokers.example.com          | Another root broker, not the cluster's RB        |

To verify the existing security setup on the node

- 1
- If node saturn is configured as an authentication broker (AB) belonging to a root broker, perform the following steps. Else, proceed to configuring the authentication broker on node saturn.



- 2 Find out the root broker to which the node saturn belongs using the following command.

```
# vssregctl -l -q -b \
"Security\Authentication\Authentication Broker" \
-k "BrokerName"
```

- 3 If the node saturn already belongs to root broker RB1, it is configured as part of the cluster. Proceed to setting up VCS related security configuration.
- 4 If the node saturn belongs to a different root broker (for example RB2), perform the following steps to remove the security credentials from node saturn.

- Kill /opt/VRTSat/bin/vxatd process.
- Remove the credential that RB2 has given to AB on node saturn.

```
# vssat deletecred --domain type:domainname \
--prplname prplname
```

For example:

```
# vssat deletecred --domain vx:root@RB2.brokers.example.com \
--prplname saturn.nodes.example.com
```

## Configuring the authentication broker on node saturn

Configure a new authentication broker (AB) on node saturn. This AB belongs to root broker RB1.

### To configure the authentication broker on node saturn

- 1 Create a principal for node saturn on root broker RB1. Execute the following command on root broker RB1.

```
# vssat addprpl --pdrtype root --domain domainname \
--prplname prplname --password password \
--prpltype service
```

For example:

```
# vssat addprpl --pdrtype root \
--domain root@RB1.brokers.example.com \
--prplname saturn.nodes.example.com \
--password flurbdicate --prpltype service
```

- 2 Ensure that there is no clock skew between the times on node saturn and RB1.

**3** Copy the `/opt/VRTSat/bin/root_hash` file from RB1 to node saturn.

**4** Configure AB on node saturn to talk to RB1.

```
# vxatd -o -a -n prplname -p password -x vx -y domainname -q \
rootbroker -z 2821 -h roothash_file_path
```

For example:

```
# vxatd -o -a -n saturn.nodes.example.com -p flurbdicate \
-x vx -y root@RB1.brokers.example.com -q RB1 \
-z 2821 -h roothash_file_path
```

**5** Verify that AB is configured properly.

```
# vssat showbrokermode
```

The command should return 1, indicating the mode to be AB.

### Setting up SF Oracle RAC related security configuration

Perform the following steps to configure SF Oracle RAC related security settings.

#### Setting up SF Oracle RAC related security configuration

**1** Start `/opt/VRTSat/bin/vxatd` process.

**2** Create `HA_SERVICES` domain for SF Oracle RAC.

```
# vssat createpd --pdrtype ab --domain HA_SERVICES
```

**3** Add SF Oracle RAC and webserver principal to AB on node saturn.

```
# vssat addprpl --pdrtype ab --domain HA_SERVICES --prplname
webserver_VCS_prplname --password new_password --prpltype
service --can_proxy
```

**4** Create `/etc/VRTSvcS/conf/config/.secure` file.

```
# touch /etc/VRTSvcS/conf/config/.secure
```

### Adding a node in a VxSS group

Perform the following procedure when adding a node in a VxSS group.

### To add a node in the VxSS group using the CLI

- 1 Make a backup copy of the main.cf file. For example:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.2node
```

- 2 On one of the nodes in the existing cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Add the new node to the VCS configuration:

```
# hasys -add saturn
```

- 4 Add the node saturn to the existing VxSS group.

```
# hagrps -modify VxSS SystemList -add saturn 2
# hagrps -modify VxSS AutoStartList -add saturn
```

- 5 Save the configuration by running the following command from any node in the cluster:

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

## Configuring server-based fencing on the new node

Perform this step if your existing cluster uses server-based I/O fencing.

### To configure server-based fencing on the new node

- 1 Log in to each CP server as the root user.
- 2 Update each CP server configuration with the new node information:

```
# /opt/VRTScps/bin/cpsadm -s thunderbolt \
-a add_node -c clus1 -u {f0735332-1dd1-11b2} -h saturn -n2
Node 2 (saturn) successfully added
```

- 3 Verify that the new node is added to the CP server configuration:

```
# /opt/VRTScps/bin/cpsadm -s thunderbolt -a list_nodes
```

The new node must be listed in the output.

- 4 Add the VCS user `cpsclient@saturn` to each CP server:

```
# /opt/VRTScps/bin/cpsadm -s thunderbolt \  
-a add_user -e cpsclient@saturn \  
-f cps_operator -g vx  
User cpsclient@saturn successfully added
```

#### To configure server-based fencing with security on the new node

- 1 As the root user, create the VCS user and the domain on the new node:

- Create a dummy configuration file `/etc/VRTSvcs/conf/config/main.cf` that resembles the following example:

```
# cat main.cf  
include "types.cf"  
cluster clus1 {  
    SecureClus = 1  
}  
system saturn {  
}
```

- Start VCS in one node mode on the new node:

```
# /opt/VRTSvcs/bin/hastart -onenode
```

- 2 Verify that the VCS user and the domain are created on the new node:

```
# /opt/VRTScps/bin/cpsat showcred | grep _HA_VCS_  
# /opt/VRTScps/bin/cpsat listpd -t local | grep HA_SERVICES
```

- 3 Stop VCS if the VCS user and domain are created successfully on the new node:

```
# /opt/VRTSvcs/bin/hastop
```

- 4 If the root broker for the CP server and the new node are different, run the following command to establish trust between the authentication broker of the CP Server and the new node:

```
# /usr/bin/echo y | /opt/VRTScps/bin/cpsat setuptrust \
-b thunderbolt -s high
```

- 5 Log in to each CP server as the root user.
- 6 Update each CP server configuration with the new node information:

```
# /opt/VRTScps/bin/cpsadm -s thunderbolt \
-a add_node -c clus1 -u {f0735332-1dd1-11b2} -h saturn -n2
Node 2 (saturn) successfully added
```

- 7 Verify that the new node is added to the CP server configuration:

```
# /opt/VRTScps/bin/cpsadm -s thunderbolt -a list_nodes
```

The new node must be listed in the output.

- 8 Add the VCS user `_HA_VCS_saturn@HA_SERVICES@saturn.veritas.com` to each CP server:

```
# /opt/VRTScps/bin/cpsadm -s thunderbolt \
-a add_user -e _HA_VCS_saturn@HA_SERVICES@saturn.veritas.com \
-f cps_operator -g vx
User _HA_VCS_saturn@HA_SERVICES@saturn.veritas.com successfully added
```

## Starting fencing on the new node

Perform the following steps to start fencing on the new node.

### To start fencing on the new node

- 1 If you are using disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/default/vxfen
/etc/vxfendg
/etc/vxfenmode
```

If you are using pure CP server-based fencing on the existing cluster, then only the `/etc/vxfenmode` file needs to be copied on the new node.

- 2 Start fencing on the new node:

For Solaris 10:

```
# svcadm enable vxfen
```

For Solaris 9:

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

- 3 On the new node, verify that the GAB port memberships are a, b, d, h and o:

```
# gabconfig -a
```

```
GAB Port Memberships
=====
Port a gen      df204 membership 012
Port b gen      df20d membership 012
Port d gen      df20a membership 012
Port o gen      df207 membership 012
```

### After adding the new node

If you added the new node to the cluster manually, before you starting VCS you must create the file `cssd-pretend-offline` on the new node and make the `cssd` resource non-critical. Failing this, the `cssd` resource lapses into an UNKNOWN state until Oracle Clusterware is installed on the new node, thus preventing the `cvm` group from coming online.

If you used the product installer to add the new node to the cluster, you will not need to perform this step.

---

**Note:** The `cssd` resource will remain in FAULTED/OFFLINE state till Oracle Clusterware is installed on the new node.

---

Start VCS on the new node.

### To start VCS on the new node

**1** If you installed the new node manually:

- On one of the nodes in the existing cluster, configure the `cssd` resource as a non-critical resource:

```
# haconf -makerw
# hares -modify cssd Critical 0
# haconf -dump -makero
```

- Create the file `cssd-pretend-offline` on the new node:

```
# touch /var/VRTSvcs/lock/cssd-pretend-offline
```

**2** Start VCS on the new node:

```
# hstart
```

VCS brings the CVM and CFS groups online.

**3** Verify that the CVM and CFS groups are online:

```
# hagrps -state
```

## Configuring CVM and CFS on the new node

Modify the existing cluster configuration to configure CVM and CFS for the new node.

### To configure CVM and CFS on the new node

**1** Make a backup copy of the `main.cf` file on the existing node, if not backed up in previous procedures. For example:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.2node
```

**2** On one of the nodes in the existing cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

**3** Add the new node to the VCS configuration, if not already added:

```
# hasys -add saturn
```

- 4 To enable the existing cluster to recognize the new node, run the following commands on one of the existing nodes:

```
# hagrps -modify cvm SystemList -add saturn 2

# hagrps -modify cvm AutoStartList -add saturn

# hares -modify cvm_clus CVMNodeId -add saturn 2

# haconf -dump -makero

# /etc/vx/bin/vxclustadm -m vcs reinit

# /etc/vx/bin/vxclustadm nidmap
```

- 5 On the remaining nodes of the existing cluster, run the following commands:

```
# /etc/vx/bin/vxclustadm -m vcs reinit

# /etc/vx/bin/vxclustadm nidmap
```

- 6 Copy the configuration files from one of the nodes in the existing cluster to the new node:

```
# rcp /etc/VRTSvcs/conf/config/main.cf \
saturn:/etc/VRTSvcs/conf/config/main.cf
# rcp /etc/VRTSvcs/conf/config/CFSTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CFSTypes.cf
# rcp /etc/VRTSvcs/conf/config/CVMTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CVMTypes.cf
```

## Configuring the ClusterService group for the new node

If the ClusterService group is configured on the existing cluster, add the node to the group by performing the steps in the following procedure on one of the nodes in the existing cluster.



### To configure the ClusterService group for the new node

- 1 On an existing node, for example galaxy, write-enable the configuration:

```
# haconf -makerw
```

- 2 Add the node saturn to the existing ClusterService group.

```
# hagrps -modify ClusterService SystemList -add saturn 2
```

```
# hagrps -modify ClusterService AutoStartList -add saturn
```

- 3 Modify the IP address and NIC resource in the existing group for the new node.

```
# hares -modify gcoip Device bge0 -sys saturn
```

```
# hares -modify gconic Device bge0 -sys saturn
```

- 4 Save the configuration by running the following command from any node.

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

## Preparing the new node for installing Oracle RAC

You must complete certain pre-installation tasks before you add the node to Oracle RAC.

Use one of the following ways to complete the preparatory tasks:

Using the SF Oracle RAC installer See [“Preparing the new nodes for installing Oracle RAC using the SF Oracle RAC installer”](#) on page 513.

Manual See [“Preparing the new node manually for installing Oracle RAC”](#) on page 521.

---

**Note:** Some of the pre-installation tasks can be completed using the SF Oracle RAC installer while some of the tasks must be completed manually as indicated in the procedures.

---

## Preparing the new nodes for installing Oracle RAC using the SF Oracle RAC installer

The SF Oracle RAC installer performs the following tasks:

- Creates the Oracle user and groups on the new node
- Configures the private IP addresses and the PrivNIC or MultiPrivNIC resources (if they are configured in the existing cluster).
- If the CFSSMount and CVMVolDg resources for OCR and voting disk are configured under the `cvm` service group, the installer brings them online after adding the node to the cluster.

If the resources are configured in any other service group, make sure that you modify the service group to include the new node and bring the service group online.

---

**Note:** If OCR and voting disk are not configured under VCS, manually mount the OCR and voting disk after you finish the steps in the following procedure.

---

- Starts the CVM group on the new node.

**To prepare to install Oracle RAC on the new node using the SF Oracle RAC installer**

- 1 After you configure SF Oracle RAC on the new node, the installer displays the following options for configuring Oracle RAC:

- 1) Create Oracle User and Group
- 2) Configure private IP addresses (PrivNIC configuration)
- 3) Configure private IP addresses (MultiPrivNIC configuration)
- 4) Finish

---

**Note:** The options 2 and 3 are displayed only if you have configured private IP addresses as PrivNIC or MultiPrivNIC resources in the existing cluster.

---

Enter **1** to select the option **Create Oracle User and Group** from the SF Oracle RAC installer menu.

See [“Creating Oracle user and groups on the new node”](#) on page 516.

- 2 Configure the private IP addresses and the PrivNIC resource for Oracle Clusterware (only if the IP addresses on the existing cluster are configured as PrivNIC resources).

You must manually update the PrivNIC resource configuration in the following cases:

- If the PrivNIC resource on the existing cluster is configured under a group other than `cvm`.
- If the `Device` attribute in the PrivNIC resource configuration on the existing cluster is not configured for each node as follows:

```
Device@galaxy = {bge1=0, bge2=1}
Device@nebula = {bge1=0, bge2=1}
```

Enter **2** to select the option **Configure private IP addresses (PrivNIC configuration)**.

See [“Configuring the private IP addresses and PrivNIC resource for Oracle Clusterware”](#) on page 517.

- 3** Configure private IP addresses and the MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC (only if the IP addresses on the existing cluster are configured as MultiPrivNIC resources).

You must manually update the MultiPrivNIC resource configuration in the following cases:

- If the MultiPrivNIC resource on the existing cluster is configured under a group other than `cvm`.
- If the `Device` attribute in the MultiPrivNIC resource configuration on the existing cluster is not configured for each node as follows:

```
Device@galaxy = {bge1=0, bge2=1, bge3=2}
Device@nebula = {bge1=0, bge2=1, bge3=2}
```

Enter **3** to select the option **Configure private IP addresses (MultiPrivNIC configuration)**.

See [“Configuring the private IP addresses and MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC”](#) on page 519.

- 4** Select **Finish** to start the `cvm` group on the new node.

---

**Note:** The `cssd` resource appears FAULTED until the new node is added to Oracle Clusterware.

---

- 5** If the `cssd` resource is configured as a critical resource, the `cvm` group will be brought offline on the new node. Modify the configuration to make the `cssd` resource non-critical and bring the `cvm` group online.

- On one of the nodes in the existing cluster, configure the `cssd` resource as a non-critical resource:

```
# haconf -makerw
# hares -modify cssd Critical 0
# haconf -dump -makero
```

- Bring the `cvm` group online:

```
# hares -online cvm -sys saturn
```

**6** Verify that all the GAB ports are up:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen ada401 membership 012
Port b gen ada40d membership 012
Port d gen ada409 membership 012
Port f gen ada41c membership 012
Port h gen ada40f membership 012
Port o gen ada406 membership 012
Port u gen ada416 membership 012
Port v gen ada416 membership 012
Port w gen ada418 membership 012
```

**7** Complete the following additional preparatory tasks using the instructions in the chapter "Preparing to install Oracle RAC":

- Identify public virtual IP addresses for use by Oracle RAC.
- Set the kernel parameters.
- Verify that the user "nobody" exists.
- Set up Oracle user equivalence for remote shell and remote copy environments.
- Edit the Oracle user profile.
- If the OCR and voting disk resources are not configured under VCS, mount the OCR and voting disk manually.

**8** Create Oracle Clusterware and Oracle RAC database home directories manually.

## Creating Oracle user and groups on the new node

Perform the steps in the following procedure to create Oracle user and groups on the new node.

---

**Note:** Set the password of the Oracle user manually before you configure secure shell or remote shell connection on the node.

---

### To create Oracle user and groups on the new node

- 1 Enter the Oracle user name that is used for Oracle RAC operations on the existing cluster.

---

**Note:** If the Oracle user and groups already exist on the new node, make sure that the UID and GID of the Oracle user and groups are the same as that on the current cluster.

---

```
Enter Oracle UNIX user name: [b] Oracle
```

The installer obtains the existing group and identifier information based on the Oracle user name.

- 2 Review and confirm the information.

The installer adds the user and groups to the new node.

- 3 Press **Return** to continue with the other configuration tasks.

### Configuring the private IP addresses and PrivNIC resource for Oracle Clusterware

Perform this step only if the private IP addresses are configured as PrivNIC resources in the existing cluster.

---

**Note:** Make sure that the network interface names of the private interconnects on the new node are the same as those of the existing cluster. For maximum failover options, all available LLT links are used for PrivNIC configuration.

---

Review the pre-configuration information displayed by the installer and ensure that you meet the requirements.

**To configure the private IP addresses and PrivNIC resource for Oracle Clusterware**

- 1 Enter the name for the PrivNIC resource. The installer displays the names of the existing PrivNIC resources. Specify the name of an existing resource.

```
Enter the PrivNIC resource name: [b] (ora_priv)
```

- 2 Enter **y** if you want the installer to add the IP address to the `/etc/hosts` file.

```
Do you want the Installer to add IP addresses in /etc/hosts  
and /etc/inet/ipnodes files? [y,n,q] (y)
```

---

**Note:** The `/etc/inet/ipnodes` file is used only for IPv6 address updates. Ignore the reference to the `/etc/inet/ipnodes` file in the prompt as this release does not support IPv6 addresses.

---

You can also add the IP address to the file manually after the configuration. The installer displays the existing PrivNIC resource configuration on the cluster.

```
Resource name: ora_priv  
System: galaxy  
      Private Interfaces: bge1 bge2  
      Private IP address: 192.168.12.1  
      Alias for above IP: galaxy-priv  
System: nebula  
      Private Interfaces: bge1 bge2  
      Private IP address: 192.168.12.2  
      Alias for above IP: nebula-priv  
Is this information correct? [y,n,q] (y)
```

- 3 Review and confirm the information.

- 4 Enter the private IP address and its private node name for the new node.

```
Enter the private IP for saturn: [b] 192.168.12.5
Enter Hostname alias for the above IP address: [b] saturn-priv
```

The installer displays the resource configuration for the new node.

```
Resource name: ora_priv
System: saturn
    Private Interfaces: bge1 bge2
    Private IP address: 192.168.12.5
    Alias for above IP: saturn-priv
Is this information correct? [y,n,q] (y)
```

- 5 Review and confirm the information.

The installer updates the existing PrivNIC resource with the resource configuration for the new node and updates the `/etc/hosts` file on the new node as well as on the existing nodes (if you chose to update the file through the installer).

- 6 Press **Return** to continue with the other configuration tasks.

## Configuring the private IP addresses and MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC

Perform this step only if the private IP addresses are configured as MultiPrivNIC resources in the existing cluster.

---

**Note:** Make sure that you configure the same interfaces (as those on the existinnebulag cluster) for private interconnects on the new node. For maximum failover options, all available LLT links are used for MultiPrivNIC configuration.

---

Review the pre-configuration information displayed by the installer and ensure that you meet the requirements.

### To configure the private IP addresses and MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC

- 1 Enter the name for the MultiPrivNIC resource. The installer displays the names of the existing MultiPrivNIC resources. Specify the name of an existing resource.

```
Enter the MultiPrivNIC resource name: [b] (multi_priv)
```

- 2 Enter **y** if you want the installer to add the IP address to the `/etc/hosts` file.

```
Do you want the Installer to add IP addresses in /etc/hosts  
and /etc/inet/ipnodes files? [y,n,q] (y)
```

---

**Note:** The `/etc/inet/ipnodes` file is used only for IPv6 address updates. Ignore the reference to the `/etc/inet/ipnodes` file in the prompt as this release does not support IPv6 addresses.

---

You can also add the IP address to the file manually after the configuration. The installer displays the existing MultiPrivNIC resource configuration on the cluster.

```
Resource name: multi_priv  
System: galaxy  
    Private Interfaces: bge1 bge2  
    Private IPs on bge1: 192.168.12.1  
    Aliases for above IPs: galaxy-priv  
    Private IPs on bge2: 192.168.2.1  
    Aliases for above IPs: galaxy-priv1  
System: nebula  
    Private Interfaces: bge1 bge2  
    Private IPs on bge1: 192.168.12.2  
    Aliases for above IPs: nebula-priv  
    Private IPs on bge2: 192.168.2.2  
    Aliases for above IPs: nebula-priv1  
Is this information correct? [y,n,q] (y)
```

- 3 Review and confirm the information.



- 4 Enter the private IP address and the corresponding private node name for the bge1 interface on the new node.

```
Enter IP addresses for saturn for bge1
separated by space: [b,q,?] 192.168.12.5
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] saturn-priv
```

- 5 Enter the private IP address and the corresponding private node name for the bge2 interface on the new node.

```
Enter IP addresses for saturn for bge2
separated by space: [b,q,?] 192.168.2.6
Enter Hostname aliases for the above IP addresses
separated by space: [b,q,?] saturn-priv1
```

The installer displays the resource configuration for the new node.

```
Resource name: multi_priv
System: saturn
    Private Interfaces: bge1 bge2
    Private IPs on bge1: 192.168.12.5
    Aliases for above IPs: saturn-priv
    Private IPs on bge2: 192.168.2.6
    Aliases for above IPs: saturn-priv1
Is this information correct? [y,n,q] (y)
```

- 6 Review and confirm the information.

The installer updates the existing MultiPrivNIC resource with the resource configuration for the new node and updates the `/etc/hosts` file on the new node as well as on the existing nodes (if you chose to update the file through the installer).

- 7 Press **Return** to continue with the other configuration tasks.

## Preparing the new node manually for installing Oracle RAC

Complete the following preparatory tasks manually before you install Oracle RAC on the new node.

### To prepare to install Oracle RAC on the new node

- 1 Create Oracle user and groups.  
See [“Creating the Oracle user and groups manually”](#) on page 333.
- 2 Configure private IP addresses and the PrivNIC resource for Oracle Clusterware.  
See [“Configuring private IP address and PrivNIC resource for Oracle Clusterware”](#) on page 522.
- 3 Configure private IP addresses and the MultiPrivNIC resource for Oracle Clusterware and Oracle UDP IPC.  
See [“Configuring private IP addresses and MultiPrivNIC resource for Oracle Clusterware and UDP IPC”](#) on page 523.
- 4 Start VCS on the new node.  
See [“Starting VCS on the new node”](#) on page 524.
- 5 Create the Oracle Clusterware/Grid Infrastructure and Oracle RAC database home directories for installation.  
See [“Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories on the new node ”](#) on page 525.
- 6 Complete the following additional preparatory tasks using the instructions in the chapter "Preparing to install Oracle RAC":
  - Identify public virtual IP addresses for use by Oracle RAC.
  - Set the kernel parameters.
  - Verify that the user "nobody" exists.
  - Set up Oracle user equivalence for remote shell and remote copy environments.
  - Edit the Oracle user profile.

## Configuring private IP address and PrivNIC resource for Oracle Clusterware

This section provides instructions for configuring private IP address and PrivNIC resource for Oracle Clusterware.

Identify a private IP address that you want to use for the new node. Make sure that the IP address is in the same subnet as the existing cluster.

The procedure uses the following IP address for the new node:

On saturn: 192.168.12.5

### To configure private IP addresses for Oracle Clusterware

- 1 Make a backup copy of the main.cf file. For example:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.2node
```

- 2 Add the private IP address to the the ora\_priv resource on the active node:

```
# haconf -makerw
# hares -modify ora_priv Device -add bge1 0 -sys saturn
# hares -modify ora_priv Device -add bge2 1 -sys saturn
# hares -modify ora_priv Address "192.168.12.5" -sys saturn
# haconf -dump -makero
```

## Configuring private IP addresses and MultiPrivNIC resource for Oracle Clusterware and UDP IPC

This section provides instructions for configuring private IP addresses and MultiPrivNIC resource for Oracle Clusterware and UDP IPC.

Identify the private IP addresses that you want to use for the new node. Make sure that the IP addresses are in the same subnet as the existing cluster.

The procedure uses the following IP addresses for the new node:

On saturn: 192.168.12.5  
 192.168.2.6

### To configure private IP addresses for Oracle Clusterware

- ◆ Use the following commands to add private IP into multi\_priv resource on the active node:

```
# haconf -makerw

# hares -modify multi_priv Device -add bge1 0 -sys saturn

# hares -modify multi_priv Device -add bge2 1 -sys saturn

# hares -modify multi_priv Address -add 192.168.12.5 0 -sys saturn

# hares -modify multi_priv Address -add 192.168.2.6 1 -sys saturn

# haconf -dump -makero
```

### Starting VCS on the new node

Before you start VCS, create the file `cssd-pretend-offline` on the new node and make the `cssd` resource non-critical. Failing this, the `cssd` resource lapses into an UNKNOWN state until Oracle Clusterware is installed on the new node, thus preventing the `cvm` group from coming online.

---

**Note:** The `cssd` resource will remain in FAULTED/OFFLINE state till Oracle Clusterware is installed on the new node.

---

### To start VCS on the new node

- 1 On one of the nodes in the existing cluster, configure the `cssd` resource as a non-critical resource:

```
# haconf -makerw
# hares -modify cssd Critical 0
# haconf -dump -makero
```

- 2 Create the file `cssd-pretend-offline` on the new node:

```
# touch /var/VRTSvcs/lock/cssd-pretend-offline
```

- 3 Start VCS on the new node:

```
# hstart
```

## Creating Oracle Clusterware/Grid Infrastructure and Oracle database home directories on the new node

The Oracle Clusterware/Grid Infrastructure and Oracle database home directories must be located on the same storage as that on the existing nodes.

Depending on the storage in the existing cluster, use one of the following options to create the directories:

|                     |                                                                                                                |
|---------------------|----------------------------------------------------------------------------------------------------------------|
| Local file system   | See <a href="#">“To create the directories on the local file system”</a> on page 525.                          |
| Veritas File System | See <a href="#">“To create the file system and directories on the Veritas File system (VxFS)”</a> on page 525. |
| Cluster File System | See <a href="#">“To create the file system and directories on cluster file system”</a> on page 530.            |

### To create the directories on the local file system

- 1 Log in as the root user on each node.
- 2 Create the directories and set the ownership and permissions:

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/10.2.0/crsHOME
# mkdir -p /u01/app/oracle/product/10.2.0/dbhome_1
# chown -R oracle:oinstall /u01/app
# chmod -R 775 /u01/app
```

For Oracle RAC 11g Release 2:

```
# mkdir -p /u01/app/gridbase
# mkdir -p /u01/app/grid/product/11.2.0/gridhome
# chown -R grid:oinstall /u01/app
# chmod -R 775 /u01/app

# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

### To create the file system and directories on the Veritas File system (VxFS)

The sample commands in the procedure are for node saturn. Repeat the steps on each node of the cluster.

- 1 As the root user, create a VxVM local diskgroup `bindg_hostname` on each node.

For Oracle RAC 11g Release 2, make sure that the size of the disk group is approximately 22 GB; if not, add another disk to the disk group.

For Oracle RAC 10g Release 2 and Oracle RAC 11g Release 1, make sure that the size of the disk group is approximately 12 GB; if not, add another disk to the disk group.

```
# vxdg init bindg_saturn Disk_1
```

- 2 Create separate volumes for Oracle Clusterware/Oracle Grid Infrastructure (`crsbinvol`) binaries and Oracle binaries (`orabinvol`):

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# vxassist -g bindg_saturn make crsbinvol 5G
# vxassist -g bindg_saturn make orabinvol 7G
```

For Oracle RAC 11g Release 2:

```
# vxassist -g bindg_saturn make crsbinvol 10G
# vxassist -g bindg_saturn make orabinvol 12G
```

- 3 Create the directories for Oracle RAC.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/10.2.0/crshome
# mkdir -p /u01/app/oracle/product/10.2.0/dbhome_1
```

For Oracle RAC 11g Release 2:

```
# mkdir -p /u01/app/gridbase
# mkdir -p /u01/app/grid/product/11.2.0/gridhome
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
```

- 4 Create file systems with the volumes `crsbinvol` and `orabinvol`:

```
# mkfs -F vxfs /dev/vx/rdisk/bindg_saturn/crsbinvol
# mkfs -F vxfs /dev/vx/rdisk/bindg_saturn/orabinvol
```

## 5 Mount the file systems on the new node.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mount -F vxfs /dev/vx/dsk/bindg_saturn/crsbinvol \
/u01/app/oracle/product/10.2.0/crshome
# mount -F vxfs /dev/vx/dsk/bindg_saturn/orabinvol \
/u01/app/oracle/product/10.2.0/dbhome_1
```

For Oracle RAC 11g Release 2:

```
# mount -F vxfs /dev/vx/dsk/bindg_saturn/crsbinvol \
/u01/app/grid/product/11.2.0/gridhome
# mount -F vxfs /dev/vx/dsk/bindg_saturn/orabinvol \
/u01/app/oracle/product/11.2.0/dbhome_1
```

## 6 Set appropriate ownership and permissions for the directories.

---

**Note:** For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1: The ownership and permissions must be changed on all nodes of the cluster because `/u01/app/oracle` must be owned by `oracle:oinstall`, otherwise `/u01/app/oracle/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

For Oracle RAC 11g Release 2: The ownership and permissions must be changed on all nodes of the cluster because `/u02/app` must be owned by `grid:install`, otherwise `/u02/app/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

---

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# chown -R oracle:oinstall /u01/app
# chmod -R 775 /u01/app
```

For Oracle RAC 11g Release 2:

```
# chown -R grid:install /u01/app
# chmod -R 775 /u01/app
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

## 7 Add the resources to the VCS configuration.

See [“To add the storage resources created on VxFS to the VCS configuration”](#) on page 528.

To add the storage resources created on VxFS to the VCS configuration

- 1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

- 2 Configure the VxVM volumes under VCS:

```
# hares -add bin_dg DiskGroup cvm
# hares -modify bin_dg DiskGroup bindg_saturn -sys saturn
# hares -modify bin_dg Enabled 1
```



### 3 Set up the file system under VCS:

```
# hares -add crsbin_mnt Mount cvm
```

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# hares -modify crsbin_mnt MountPoint \  
"/u01/app/oracle/product/10.2.0/crshome"
```

For Oracle RAC 11g Release 2:

```
# hares -modify crsbin_mnt MountPoint \  
"/u01/app/grid/product/11.2.0/gridhome"
```

The following commands are common to all Oracle RAC versions:

```
# hares -modify crsbin_mnt BlockDevice \  
"/dev/vx/dsk/bindg_saturn/crsbinvol" -sys saturn  
# hares -modify crsbin_mnt FSType vxfs  
# hares -modify crsbin_mnt FsckOpt "-n"  
# hares -modify crsbin_mnt Enabled 1  
# hares -add orabin_mnt Mount cvm
```

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# hares -modify orabin_mnt MountPoint \  
"/u01/app/oracle/product/10.2.0/dbhome_1"
```

For Oracle RAC 11g Release 2:

```
# hares -modify orabin_mnt MountPoint \  
"/u01/app/oracle/product/11.2.0/dbhome_1"
```

The following commands are common to all Oracle RAC versions:

```
# hares -modify orabin_mnt BlockDevice \  
"/dev/vx/dsk/bindg_saturn/orabinvol" -sys saturn  
# hares -modify orabin_mnt FSType vxfs  
# hares -modify orabin_mnt FsckOpt "-n"  
# hares -modify orabin_mnt Enabled 1
```

### 4 Link the parent and child resources:

```
# hares -link crsbin_mnt bin_dg  
# hares -link orabin_mnt bin_dg
```

**To create the file system and directories on cluster file system**

- 1** As the root user, create a VxVM shared disk group `bindg`:

```
# vxvg -s init bindg Disk_1
```

- 2** Create separate volumes for Oracle Clusterware/Oracle Grid Infrastructure (`crsbinvol`) and Oracle database (`orabinvol`):

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# vxassist -g bindg make crsbinvol 5G
# vxassist -g bindg make orabinvol 7G
```

For Oracle RAC 11g Release 2:

```
# vxassist -g bindg make crsbinvol 10G
# vxassist -g bindg make orabinvol 12G
```

- 3** Create the following directories for Oracle. The file system and directories created on shared storage in this procedure are based on the following layout:

|                                                                                  |                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>\$ORACLE_BASE</b>                                                             | <i>/u01/app/oracle</i><br><br>/u01/app/oracle is on local storage.<br><br>For Oracle RAC 11g Release 2, create a separate Oracle base directory for the grid user, in addition to the directory for the Oracle user:<br><br><i>/u01/app/gridbase</i><br><br>/u01/app/gridbase is on local storage. |
| <b>\$CRS_HOME</b><br><br>(For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1) | <i>/u01/app/oracle/product/10.2.0/crshome</i><br><br>/u01/app/oracle/product/10.2.0 is on local storage.<br><br>/u01/app/oracle/product/10.2.0/crshome is on shared storage.                                                                                                                       |
| <b>\$GRID_HOME</b><br><br>(For Oracle RAC 11g Release 2)                         | <i>/u01/app/grid/product/11.2.0/gridhome</i><br><br>/u01/app/grid/product/11.2.0 is on local storage.<br><br>/u01/app/grid/product/11.2.0/gridhome is on shared storage.                                                                                                                           |
| <b>\$ORACLE_HOME</b>                                                             | <i>/u01/app/oracle/product/release/dbhome_1</i><br><br>where <i>release</i> is 10.2.0, 11.1.0, or 11.2.0 depending on the Oracle RAC version<br><br>/u01/app/oracle/product/ <i>release</i> is on local storage.<br><br>/u01/app/oracle/product/ <i>release</i> /dbhome_1 is on shared storage.    |

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/10.2.0/crshome
# mkdir -p /u01/app/oracle/product/10.2.0/dbhome_1
```

For Oracle RAC 11g Release 2:

```
# mkdir -p /u01/app/gridbase
# mkdir -p /u01/app/grid/product/11.2.0/gridhome
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
```

**4** Create file systems with the volumes `crsbinvol` and `orabinvol`:

```
# mkfs -F vxfs /dev/vx/rdisk/bindg/crsbinvol
# mkfs -F vxfs /dev/vx/rdisk/bindg/orabinvol
```

**5** Mount the file systems on the new node.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/crsbinvol \
/u01/app/oracle/product/10.2.0/crshome
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/orabinvol \
/u01/app/oracle/product/10.2.0/dbhome_1
```

For Oracle RAC 11g Release 2:

```
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/crsbinvol \
/u01/app/grid/product/11.2.0/gridhome
# mount -F vxfs -o cluster /dev/vx/dsk/bindg/orabinvol \
/u01/app/oracle/product/11.2.0/dbhome_1
```

## 6 Change the ownership and permissions on the new node.

---

**Note:** For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1: The ownership and permissions must be changed on all nodes of the cluster because `/u01/app/oracle` must be owned by `oracle:oinstall`, otherwise `/u01/app/oracle/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

For Oracle RAC 11g Release 2: The ownership and permissions must be changed on all nodes of the cluster because `/u01/app` must be owned by `grid:install`, otherwise `/u01/app/oraInventory` does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

---

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# chown -R oracle:oinstall /u01/app
# chmod -R 775 /u01/app
```

For Oracle RAC 11g Release 2:

```
# chown -R grid:install /u01/app
# chmod -R 775 /u01/app
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

## 7 Add the CVMVolDg and CFSSMount resources to the VCS configuration.

See [“To add the CFSSMount and CVMVolDg resources to the VCS configuration using CLI”](#) on page 372.

# Adding the new node to Oracle RAC

Install Oracle Clusterware and Oracle RAC database on the node using the Oracle RAC add node procedure.

For instructions, see the Oracle RAC documentation.

After installing Oracle Clusterware and Oracle RAC database, perform the following post-installation tasks:

1. Delete the file `/var/VRTSvcs/lock/cssd-pretend-offline` to bring the CSSD resource online.
2. If CSSD is not configured under the `cvm` group, add the new node information to the service group containing the CSSD resource.

3. Add new Oracle RAC database instances for the new node.  
For instructions, see the Oracle RAC documentation.
4. Add Oracle UDP IPC private IP addresses to the Oracle `init.ora` file.
5. Update the Oracle RAC database service groups to include the new database instances in the VCS configuration file.
6. For other service groups that are configured under VCS, manually update the service group configuration for the new node.

## Updating the Storage Foundation for Databases (SFDB) repository after adding a node

If you are using Database Checkpoints, Database Flashsnap, or Adding a Node in your configuration, update the SFDB repository to enable access for the new node after it is added to the cluster.

### To update the SFDB repository after adding a node

1. Run the following to change permission, owner, group of various SFDB directories on the newly added node:

```
# sfua_db_config
```

2. Run the `dbed_update` command on any one node in the cluster. For example:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

This completes the addition of the node to the SFDB repository.

For information on using SFDB tools features:

See the Storage Foundation guide: *Storage Foundation: Storage and Availability Management for Oracle Databases*.

## Sample configuration file for adding a node to the cluster

You may use this sample file as reference information to understand the configuration changes that take place when you add a node to a cluster.

The existing sample configuration before adding the node `saturn` is as follows:

- The existing cluster `rac_cluster101` comprises two nodes `galaxy` and `nebula` and hosts a single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- Only one private IP address is configured for Oracle Clusterware. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

The following sample configuration file shows the changes (in **bold**) effected in the configuration after adding a node "saturn" to the cluster.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)

system galaxy (
)
system nebula (
)
system saturn (
)
```

---

**Note:** In the following group `oradb1_grp`, the saturn node has been added.

---

```
group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn = 2 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)
```

---

**Note:** In the following Oracle resource, the saturn node information has been added.

---

```
Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Sid @saturn = vrts3
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

requires group cvm online local firm
oral requires oradata_mnt
oradata_mnt requires oradata_voldg
```

---

**Note:** In the following CVM and CVMCluster resources, the saturn node information has been added.

---

```
group cvm (
    SystemList = { galaxy = 0, nebula = 1, saturn =2 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)
```



```

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline"
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
    OnlineRetryLimit = 20
)

CFMount ocrvote_mnt (
    Critical = 0
    MountPoint = "/ocrvote"
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"
    MountOpt= "mincache=direct"
)

CVMVolDg ocrvote_voldg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvotevol }
    CVMActivation = sw
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1, saturn =2 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

```

---

**Note:** In the following PrivNIC resource, the saturn node information has been added.

---

```
PrivNIC ora_priv (  
    Critical = 0  
    Device@galaxy = { bge1 = 0, bge2 = 1}  
    Device@nebula = { bge1 = 0, bge2 = 1}  
    Device@saturn = { bge1 = 0, bge2 = 1}  
    Address@galaxy = "192.168.12.1"  
    Address@nebula = "192.168.12.2"  
    Address@saturn = "192.168.12.5"  
    NetMask = "255.255.255.0"  
)  
  
cssd requires ocrvote_mnt  
cssd requires ora_priv  
ocrvote_mnt requires ocrvote_voldg  
ocrvote_mnt requires vxfsckd  
ocrvote_voldg requires cvm_clus  
vxfsckd requires cvm_clus  
cvm_clus requires cvm_vxconfigd
```

# Removing a node from SF Oracle RAC clusters

This chapter includes the following topics:

- [About removing a node from a cluster](#)
- [Removing a node from a cluster](#)
- [Modifying the VCS configuration files on existing nodes](#)
- [Removing the node configuration from the CP server](#)
- [Removing security credentials from the leaving node](#)
- [Updating the Storage Foundation for Databases \(SFDB\) repository after removing a node](#)
- [Sample configuration file for removing a node from the cluster](#)

## About removing a node from a cluster

You can remove one or more nodes from an SF Oracle RAC cluster if the node no longer needs to be part of the cluster.

Removing a node from a cluster requires:

- Taking Oracle service groups offline.
- Stopping applications that use File System or Cluster File System mount points not configured under VCS.
- Removing database instances and software from the node.
- Stopping VCS on the node to be removed.
- Remove Oracle Clusterware from the node.

- Unmounting the File System and Cluster File System file systems not configured under VCS.
- Uninstalling SF Oracle RAC from the node.  
Modifying the VCS configuration files on the existing nodes.
- Removing the node configuration from the CP server if it is configured.
- Removing the security credentials from the node if it is part of a secure cluster.
- Updating the SFDB repository if you use SFDB tools.

The Veritas product installer does not support removing a node. You must remove a node manually. The example procedures describe how to remove a node from a cluster with three nodes.

## Removing a node from a cluster

Perform the following steps to remove a node from a cluster. The procedure can be done from any node remaining in the cluster or from a remote host.

### To remove a node from a cluster

- 1 Take the Oracle service groups offline (if under VCS control) on the node you want to remove.

```
# hagrps -offline oracle_group -sys saturn
```

- 2 Stop the applications that use VxFS/CFS mount points and are not configured under VCS. Use native application commands to stop the applications.

- 3 Remove the Oracle RAC database instance from the node.

For instructions, see the Oracle RAC documentation.

- 4 Remove the Oracle RAC database software from the node.

For instructions, see the Oracle RAC documentation.

- 5 Stop VCS on the node:

```
# hastop -local
```

- 6 Remove Oracle Clusterware from the node.

For instructions, see the Oracle RAC document.

- 7 Unmount the VxFS/CFS file systems that are not configured under VCS.

```
# umount mount_point
```

- 8 Uninstall SF Oracle RAC from the node using the SF Oracle RAC installer.

```
# cd /opt/VRTS/install  
  
# ./uninstallsfrac saturn
```

The installer stops all SF Oracle RAC processes and uninstalls the SF Oracle RAC packages.

- 9 Modify the VCS configuration files on the existing nodes to remove references to the deleted node.

See [“Modifying the VCS configuration files on existing nodes”](#) on page 541.

## Modifying the VCS configuration files on existing nodes

Modify the configuration files on the remaining nodes of the cluster to remove references to the deleted nodes.

The process involves:

- [Editing the /etc/llthosts file](#)
- [Editing the /etc/gabtab file](#)
- [Modifying the VCS configuration to remove the node](#)

For an example main.cf:

### Editing the /etc/llthosts file

On each of the existing nodes, edit the `/etc/llthosts` file to remove lines that contain references to the removed nodes.

For example, if saturn is the node removed from the cluster, remove the line "2 saturn" from the file:

```
0 galaxy  
1 nebula  
2 saturn
```

Change to:

```
0 galaxy  
1 nebula
```

## Editing the /etc/gabtab file

Modify the following command in the `/etc/gabtab` file to reflect the number of systems after the node is removed:

```
/sbin/gabconfig -c -nN
```

where N is the number of remaining nodes in the cluster.

For example, with two nodes remaining, the file resembles:

```
/sbin/gabconfig -c -n2
```

## Modifying the VCS configuration to remove the node

Modify the VCS configuration file `main.cf` to remove all references to the deleted node.

Use one of the following methods to modify the configuration:

- Edit the `/etc/VRTSvcs/conf/config/main.cf` file  
This method requires application down time.
- Use the command line interface  
This method allows the applications to remain online on all remaining nodes.

The following procedure uses the command line interface and modifies the sample VCS configuration to remove references to the deleted node. Run the steps in the procedure from one of the existing nodes in the cluster. The procedure allows you to change the VCS configuration while applications remain online on the remaining nodes.

**To modify the VCS configuration using the CLI**

- 1 Back up the `/etc/VRTSvcs/conf/config/main.cf` file.

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.3node.bak
```

- 2 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Remove the node from the `AutoStartList` attribute of the service group by specifying the remaining nodes in the desired order:

```
# hagrps -modify cvm AutoStartList galaxy nebula
```

- 4 Remove the node from the SystemList attribute of the service group:

```
# hagrpf -modify cvm SystemList -delete saturn
```

- 5 Remove the node from the CVMNodeId attribute of the service group:

```
# hares -modify cvm_clus CVMNodeId -delete saturn
```

- 6 If you have the other service groups (such as the database service group or the ClusterService group) that have the removed node in their configuration, perform step 4 and step 5 for each of them.

- 7 Remove the deleted node from the NodeList attribute of all CFS mount resources:

```
# hares -modify CFSMount NodeList -delete saturn
```

- 8 Remove the deleted node from the system list of any other service groups that exist on the cluster. For example, to delete the node saturn:

```
# hagrpf -modify crsgrp SystemList -delete saturn
```

- 9 Remove the deleted node from the cluster system list:

```
# hasys -delete saturn
```

- 10 Save the new configuration to disk:

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

- 11 Verify that the node is removed from the VCS configuration.

```
# grep -i saturn main.cf
```

If the node is not removed, use the VCS commands as described in this procedure to remove the node.

## Removing the node configuration from the CP server

After removing a node from a SF Oracle RAC cluster, perform the steps in the following procedure to remove that node's configuration from the CP server.

---

**Note:** The `cpsadm` command is used to perform the steps in this procedure. For detailed information about the `cpsadm` command, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

---

**To remove the node configuration from the CP server**

- 1 Log into the CP server as the root user.
- 2 View the list of VCS users on the CP server, using the following command:

```
# cpsadm -s cp_server -a list_users
```

Where *cp\_server* is the virtual IP/ virtual hostname of the CP server.

- 3 Remove the VCS user associated with the node you previously removed from the cluster.

For CP server in secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e _HA_VCS_saturn@HA_SERVICES@saturn.nodes.example.com \  
-f cps_operator -g vx
```

For CP server in non-secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e cpsclient@saturn -f cps_operator -g vx
```

- 4 Remove the node entry from the CP server:

```
# cpsadm -s cp_server -a rm_node -h saturn -c rac_cluster101 -n 2
```

- 5 View the list of nodes on the CP server to ensure that the node entry was removed:

```
# cpsadm -s cp_server -a list_nodes
```

## Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node saturn. Perform the following steps.



### To remove the security credentials

- 1 Kill the `/opt/VRTSat/bin/vxatd` process.
- 2 Remove the root credentials on node saturn.

```
# vssat deletecred --domain type:domainname --prplname prplname
```

## Updating the Storage Foundation for Databases (SFDB) repository after removing a node

If you are using Database Checkpoints, Database Flashsnap, or SmartTier for Oracle in your configuration, update the SFDB repository to remove the reference for the node after removing the node from the cluster.

---

**Note:** If you have not created an SFDB repository, you do not need to perform the following steps.

---

### To update the SFDB repository after removing a node

- 1 As Oracle user, list the nodes in the cluster:

```
$ /opt/VRTSdbed/bin/dbed_rept_node -S $ORACLE_SID -o list
```

- 2 Run the following command after physically removing the node from the cluster.

For example:

```
$ /opt/VRTSdbed/bin/dbed_rept_node -S $ORACLE_SID -n NODE -o remove
```

This completes the removal of the node from the SFDB repository.

## Sample configuration file for removing a node from the cluster

You may use this sample file as reference information to understand the configuration changes involved when you remove a node from a cluster.

The existing sample configuration before removing the node `saturn` is as follows:

- The existing cluster `rac_cluster101` comprises three nodes `galaxy`, `nebula`, and `saturn` and hosts a single database.
- The Oracle database is stored on CFS.

**Sample configuration file for removing a node from the cluster**

- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- Only one private IP address is configured for Oracle Clusterware. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

---

**Note:** The following sample file shows in **bold** the configuration information that is removed when the node "saturn" is removed from the cluster.

---

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)

system galaxy (
)
system nebula (
)
system saturn (
)
```

---

**Note:** In the following group oradb1\_grp, the saturn node must be removed.

---

```
group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn = 2 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)
```

---

**Note:** In the following Oracle resource, the saturn node information must be removed.

---

```
Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Sid @saturn = vrts3
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

requires group cvm online local firm
oral requires oradata_mnt
oradata_mnt requires oradata_voldg
```

---

**Note:** In the following CVM and CVMCluster resources, the saturn node information must be removed.

---

```
group cvm (
    SystemList = { galaxy = 0, nebula = 1, saturn =2}
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)
```

```
Application cssd (  
    Critical = 0  
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"  
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline"  
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"  
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"  
    OnlineRetryLimit = 20  
)  
  
CFSMount ocrvote_mnt (  
    Critical = 0  
    MountPoint = "/ocrvote"  
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"  
    MountOpt= "mincache=direct"  
)  
  
CVMVolDg ocrvote_voldg (  
    Critical = 0  
    CVMDiskGroup = ocrvotedg  
    CVMVolume = { ocrvotevol }  
    CVMActivation = sw  
)  
  
CFSfsckd vxfsckd (  
)  
  
CVMCluster cvm_clus (  
    CVMClustName = rac_cluster101  
    CVMNodeId = { galaxy = 0, nebula = 1, saturn =2 }  
    CVMTransport = gab  
    CVMTimeout = 200  
)  
  
CVMVxconfigd cvm_vxconfigd (  
    Critical = 0  
    CVMVxconfigdArgs = { syslog }  
)
```

---

**Note:** In the following PrivNIC resource, the saturn node information must be removed.

---

```
PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { bge1 = 0, bge2 = 1}
    Device@nebula = { bge1 = 0, bge2 = 1}
    Device@saturn = { bge1 = 0, bge2 = 1}
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    Address@saturn = "192.168.12.5"
    NetMask = "255.255.255.0"
)
```

```
cssd requires ocrvote_mnt
cssd requires ora_priv
ocrvote_mnt requires ocrvote_voldg
ocrvote_mnt requires vxfsckd
ocrvote_voldg requires cvm_clus
vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd
```



# Configuration of disaster recovery environments

- [Chapter 31. Configuring a campus cluster setup for disaster recovery](#)
- [Chapter 32. Setting up a replicated global cluster](#)
- [Chapter 33. Configuring a global cluster using VVR](#)





# Configuring a campus cluster setup for disaster recovery

This chapter includes the following topics:

- [Requirements for an SF Oracle RAC campus cluster](#)
- [Setting up an SF Oracle RAC campus cluster for disaster recovery](#)
- [Tuning guidelines for campus clusters](#)
- [Best practices for a campus cluster](#)

## Requirements for an SF Oracle RAC campus cluster

You will need the following to setup an SF Oracle RAC campus cluster:

- Hardware requirements for SF Oracle RAC:  
See [“Hardware requirements”](#) on page 46.
- License requirements: in addition to SF Oracle RAC with HA/DR, you will need:
  - FlashSnap license
  - Site awareness license

With keyless licensing, SF Oracle RAC enterprise keys enable all of the above features.

# Setting up an SF Oracle RAC campus cluster for disaster recovery

Perform the following high-level tasks to setup an SF Oracle RAC campus cluster:

- Preparing to set up campus cluster configuration
- Configuring I/O fencing to prevent data corruption
- Preparing to install Oracle RAC Clusterware and database binaries
- Configuring VxVM disk groups for campus cluster
- Installing Oracle RAC Clusterware and database binaries
- Configuring VCS service groups

The sample configuration used to illustrate the configuration procedures includes a four-node SF Oracle RAC campus cluster with two nodes at each site. Each node has SF Oracle RAC 5.1 Service Pack 1 on Solaris.

**Table 31-1**            Sample setup for campus cluster

| Site           | Hardware description                                                                                                   |
|----------------|------------------------------------------------------------------------------------------------------------------------|
| Site 1: SFRAC1 | Servers: galaxy and nebula<br>Shared LUNs:<br>disk01<br>disk02<br>disk03<br>disk04 (used as coordinator disk)<br>disk5 |
| Site 2: SFRAC2 | Servers: mercury and jupiter<br>Shared LUNs:<br>disk06<br>disk07<br>disk08<br>disk09 (used as coordinator disk)        |
| Site 3         | Shared LUN disk10 (used as coordinator disk)                                                                           |

## Preparing to set up a campus cluster

The following preparation must be completed before configuring the campus cluster.

### To prepare to set up a campus cluster

- 1 Configure the physical infrastructure for campus cluster:
  - Set up access to the local storage arrays and to remote storage arrays on each node. The storage link will extend to the third site as well.
  - Set up the private heartbeat network  
 See the *Veritas Cluster Server Administrator's Guide*
- 2 Install the operating system on all the nodes of the cluster.  
 See your operating system documentation.
- 3 Install and configure SF Oracle RAC on all nodes on both the sites.  
 See [“About configuring SF Oracle RAC”](#) on page 145.  
 In the sample setup, install and configure SF Oracle RAC 5.1 Service Pack 1 on all four nodes.

## Configuring VxVM disk groups for Oracle RAC on a campus cluster

After configuring I/O fencing for data integrity, you must configure the VxVM disk groups for a campus cluster before installing Oracle RAC by configuring VxVM disk groups for remote mirroring.

### To configure VxVM disk groups for Oracle RAC on a campus cluster

- 1 Initialize the disks as CDS disks

```
# vxdisksetup -i disk01 format=cdsdisk
# vxdisksetup -i disk02 format=cdsdisk
# vxdisksetup -i disk03 format=cdsdisk
# vxdisksetup -i disk05 format=cdsdisk
# vxdisksetup -i disk06 format=cdsdisk
# vxdisksetup -i disk07 format=cdsdisk
# vxdisksetup -i disk08 format=cdsdisk
```

- 2 Set the site name for each host:

```
# vxdctl set site=sitename
```

The site name is stored in the `/etc/vx/volboot` file. To display the site names:

```
# vxdctl list | grep siteid
```

For example, for a four node cluster with two nodes at each site, mark the sites as follows:

On the nodes at first site:

```
# vxdctl set site= SFRAC1
```

On the nodes at second site:

```
# vxdctl set site= SFRAC2
```

- 3 Set the site name for all the disks in an enclosure.

```
# vxdisk settag site=sitename encl:enclosure
```

- 4 Run the following command if you want to tag only the specific disks:

```
# vxdisk settag site=sitename disk
```

For example:

```
# vxdisk settag site=SFRAC1 disk01
# vxdisk settag site=SFRAC1 disk02
# vxdisk settag site=SFRAC1 disk03
# vxdisk settag site=SFRAC2 disk06
# vxdisk settag site=SFRAC2 disk08
```

- 5 Verify that the disks are registered to a site.

```
# vxdisk listtag
```

For example:

```
# vxdisk listtag
```

| DEVICE | NAME | VALUE  |
|--------|------|--------|
| disk01 | site | SFRAC1 |
| disk02 | site | SFRAC1 |
| disk03 | site | SFRAC1 |
| disk06 | site | SFRAC2 |
| disk08 | site | SFRAC2 |

- 6 Create a disk group for OCR and Vote Disks and another for Oracle data, with disks picked from both the sites. While the example below shows a single disk group, you can create as many as you need.

```
# vxdg -s init ocrvotedg disk05 disk07
# vxdg -s init oradatadg disk01 disk06
```

For the disks used in ocrvotedg, do not set site consistency to **on**.

- 7 Enable site-based allocation on the disk groups for each site.

```
# vxdg -g oradata addsite SFRAC1
# vxdg -g oradata addsite SFRAC2
```

- 8 Set site consistency on the disk groups.

```
# vxdg -g oradatadg set siteconsistent=on
# vxedit -g oradatadg set dgdetppolicy=global
# vxedit -g oradatadg set dgfailpolicy=leave
# vxedit -g ocrvotedg set dgdetppolicy=global
# vxedit -g ocrvotedg set dgfailpolicy=leave
```

- 9 If you are using an enclosure, set the tag on the enclosure for both sites.

```
# vxdg -o retain -g oradatadg setting encl:3pardata0 site=SFRAC1
# vxdg -o retain -g oradatadg setting encl:3pardata1 site=SFRAC2
```

- 10 Create one or more mirrored volumes in the disk group.

```
# vxassist -g ocrvotedg make ocrvotevol 2048m nmirror=2
# vxassist -g oradatadg make oradatavol 10200m nmirror=2
```

With the Site Awareness license installed on all hosts, the volume created has the following characteristics by default.

- The all sites attribute is set to ON; the volumes have at least one mirror at each site.
- The volumes are automatically mirrored across sites.
- The read policy (rdpol) is set to siteread.
- The volumes inherit the site consistency value that is set on the disk group.

- 11 From the CVM master, start the volumes for all the disk groups.

```
# vxvol -g ocrvotedg startall
# vxvol -g oradatadg startall
```

- 12 Create a file system on each volume and mount the same.

```
# mkfs -F vxfs /dev/vx/rdsk/ocrvotedg/ocrvotevol
```

```
# mkfs -F vxfs /dev/vx/rdsk/oradatadg/oradatavol

# mount -F vxfs -o cluster /dev/vx/dsk/ocrvotedg/ocrvotevol /ocr

# mount - vxfs -o cluster /dev/vx/dsk/oradatadg/oradatavol /oradata
```

13 Touch two files, one for OCR and another for Vote Disk.

```
# touch /ocr/ocr /ocr/vote
```

---

**Note:** One Vote Disk is sufficient since it is already mirrored by VxVM.

---

## Tuning guidelines for campus clusters

An important consideration while tuning an SF Oracle RAC campus cluster is setting the LLT peerinact time. Follow the guidelines below to determine the optimum value of peerinact time:

- Calculate the roundtrip time using `ltping` (1M).
- Evaluate LLT heartbeat time as half of the round trip time.
- Set the LLT peer trouble time as 2-4 times the heartbeat time.
- LLT peerinact time should be set to be more than 4 times the heart beat time.

## Best practices for a campus cluster

The following best practices ensure a robust SF Oracle RAC campus cluster:

- Tag all the mirrored volumes in the campus cluster with appropriate site names. VxVM allocates storage from the correct site when creating or resizing a volume and when changing a volume's layout if the volume is tagged with site name.
- All volumes that have data required by the application must be evenly mirrored. Each site must have at least one mirror of all volumes hosting application data, including the FlashSnap log volume.
- Do not enable site consistency on VxVM snapshot volumes.
- Use redundant links for storage and private interconnects. DWDM can be used for storage and heartbeat together. Another redundant DWDM link can be used to prevent single point of failure. Separate switches and multiplexer / de-multiplexer devices should be used.

- Use Coordination Point Server as the third coordination point.
- Use the procedure for online replacement of coordination points, to replace disk based or Coordination Point Server based coordination points.





# Setting up a replicated global cluster

This chapter includes the following topics:

- [Replication in the SF Oracle RAC environment](#)
- [Requirements for SF Oracle RAC global clusters](#)
- [About setting up a global cluster in an SF Oracle RAC environment](#)
- [Configuring a cluster at the primary site](#)
- [Configuring a cluster at the secondary site](#)
- [Configuring replication on clusters at both sites](#)
- [Modifying the ClusterService group for global clusters](#)
- [Defining the remote cluster and heartbeat objects](#)
- [Configuring the VCS service groups for global clusters](#)

## Replication in the SF Oracle RAC environment

You can set up a primary SF Oracle RAC cluster for replication to a secondary SF Oracle RAC cluster by configuring global VCS service groups and using a replication technology. The Oracle RAC cluster at the secondary site can be a single node cluster. For example, you can have a two-node cluster on the primary site and a two-node or single-node cluster on the secondary site.

You can use one of the following replication technologies:

- Veritas Volume Replicator (VVR), which provides host-based volume replication. Using VVR you can replicate data volumes on a shared disk group in SF Oracle RAC.
- Supported hardware-based replication technologies. Using hardware-based replication you can replicate data from a primary array to a secondary array.
- Using SF Oracle RAC with VVR you can run a fire drill to verify the disaster recovery capability of your configuration.  
See the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

## Requirements for SF Oracle RAC global clusters

Review the requirements information to make sure your configuration is supported for SF Oracle RAC.

For product licensing information:

See “[About Veritas product licensing](#)” on page 79.

### Supported software and hardware for SF Oracle RAC

- For supported hardware and software:
- See “[Important preinstallation information](#)” on page 45.
  - See the current compatibility list in the Veritas Technical Support website to confirm the compatibility of your hardware:  
<http://entsupport.symantec.com/docs/283161>

### Supported replication technologies for SF Oracle RAC

SF Oracle RAC supports the following replication technologies through the use of Veritas replication agents:

| Table 32-1 Supported replication options for SF Oracle RAC global clusters                                                                                    |                                                                                                              |                        |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|------------------------|
| Replication technology                                                                                                                                        | Supported modes                                                                                              | Supported software     |
| Veritas Volume Replicator (VVR)<br>Supporting agents <ul style="list-style-type: none"><li>■ RVGShared</li><li>■ RVGSharedPri</li><li>■ RVGLogOwner</li></ul> | <ul style="list-style-type: none"><li>■ Asynchronous replication</li><li>■ Synchronous replication</li></ul> | Host-based replication |

**Table 32-1** Supported replication options for SF Oracle RAC global clusters  
(continued)

| Replication technology                            | Supported modes                                                                                                                                         | Supported software                                                                                |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| EMC SRDF<br>Supporting agent: SRDF                | <ul style="list-style-type: none"> <li>■ Asynchronous replication</li> <li>■ Synchronous replication</li> </ul>                                         | All versions of Solutions Enabler                                                                 |
| Hitachi True Copy<br>Supporting agent: HTC        | <ul style="list-style-type: none"> <li>■ Asynchronous replication</li> <li>■ Synchronous replication</li> </ul>                                         | All versions of the Hitachi CCI                                                                   |
| IBM Metro Mirror<br>Supporting agent: MetroMirror | Synchronous replication                                                                                                                                 | All versions of IBM DSCLI.<br><br>The MetroMirror agent is supported for DS6000 and DS8000 arrays |
| IBM SVC<br>SVC CopyServices                       | <ul style="list-style-type: none"> <li>■ Asynchronous replication</li> <li>■ Synchronous replication</li> </ul>                                         | SSH access to the SVC                                                                             |
| EMC Mirror View<br>Supporting agent: MirrorView   | <ul style="list-style-type: none"> <li>■ Asynchronous replication</li> <li>■ Synchronous replication: only individual LUNs may be replicated</li> </ul> | All versions of NaviCLI                                                                           |

**Note:** Check your vendor's compatibility list for the supported software versions. The support listed above only exists if the host, HBA, and array combination is in your vendor's hardware compatibility list. Check your array documentation.

**Note:** All arrays must support SCSI-3 persistent reservations for SF Oracle RAC.

You can use the Veritas replication agents listed in the table above for global clusters that run SF Oracle RAC. The Veritas replication agents provide application failover and recovery support to your replication configuration. The agents provide this support for environments where data is replicated between clusters.

VCS agents control the direction of replication. They do not monitor the progress or status of replication. The replication agents manage the state of replicated

devices that are attached to SF Oracle RAC nodes. The agents make sure that the system which has the resource online also has safe and exclusive access to the configured devices.

This information is current at the time this document is released. For more current information on the replicated agents, see:

- *Veritas Cluster Server Agent for EMC SRDF Installation and Configuration Guide*
- *Veritas Cluster Server Agent for Hitachi TrueCopy Installation and Configuration Guide*
- *Veritas Cluster Server Agent for IBM Metro Mirror Installation and Configuration Guide*
- *Veritas Cluster Server Agent for IBM SVC Installation and Configuration Guide*
- *Veritas Cluster Server Agent for EMC MirrowView Installation and Configuration Guide*
- *Veritas Cluster Server Agent for Oracle Data Guard Installation and Configuration Guide*
- Technical Support TechNote for the latest updates or software issues for replication agents:  
<http://entsupport.symantec.com/docs/282004htm>

## About setting up a global cluster in an SF Oracle RAC environment

Configuring a global cluster for Oracle RAC requires the coordination of many component setup tasks. The procedures provided in this document are guidelines.

The tasks required to set up a global cluster:

- Configure an SF Oracle RAC cluster at the primary site
- Configure an SF Oracle RAC cluster at the secondary site
- Configure replication on clusters at both sites
- Configure VCS service groups for replication
- Test the HA/DR configuration
- Upon successful testing, bring the environment into production

Some SF Oracle RAC HA/DR configuration tasks may require adjustments depending upon your particular starting point, environment, and configuration. Review the installation requirements and sample cluster configuration files for primary and secondary clusters.

For requirements:

For instructions for configuring AT in a global cluster:

See the *Veritas Cluster Server Administrator's Guide*

## Configuring a cluster at the primary site

You can use an existing SF Oracle RAC cluster or you can install a new SF Oracle RAC cluster for your primary site.

For planning information:

See [“Typical configuration of SF Oracle RAC global clusters for disaster recovery”](#) on page 43.

If you are using an existing cluster as the primary and you want to set up a global cluster, skip the steps below and proceed to configure your secondary cluster.

See [“Configuring a cluster at the secondary site”](#) on page 567.

---

**Note:** You must have a GCO license enabled for a global cluster. If you are using VVR for replication, you must have a VVR license enabled.

---

If you do not have an existing cluster and you are setting up two new sites for an SF Oracle RAC global cluster, follow the steps below.

### To set up the cluster and database at the primary site

- 1 Install and configure servers and storage.
- 2 If you are using hardware-based replication, install the software for managing your array.
- 3 Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster. If you are using VVR for replication, you must have it enabled.
- 4 Install and configure SF Oracle RAC. Prepare for your installation according to your configuration needs.

For preparation:

For installation:

See [“About installing SF Oracle RAC ”](#) on page 91.

- 5 For a multi-node cluster, configure I/O fencing.

- Verify the shared storage on the secondary site supports SCSI-3 reservations.
- Set up coordinator disks
- Configure I/O fencing

For instructions for setting up fencing:

See [“About planning to configure I/O fencing”](#) on page 67.

- 6 Verify the CVM group is online on all nodes in the primary cluster:

```
# hagrps -state cvm
```

- 7 Prepare systems and storage for a global cluster. Identify the hardware and storage requirements before installing Oracle RAC Clusterware and RDBMS software.

You will need to set up:

- Local storage for Oracle RAC and CRS binaries
- Shared storage for OCR and Vote disk which is not replicated
- Replicated storage for database files

- 8 Install and configure the Oracle RAC binaries:

See [“About preparing to install Oracle RAC”](#) on page 327.

See [“About installing Oracle RAC”](#) on page 381.

---

**Note:** OCR and Vote disk must be on non-replicated shared storage.

---

After successful Oracle RAC installation and configuration, verify that CRS daemons and resources are up on all nodes.

```
$ crs_stat -t
```

- 9 Identify the disks that will be replicated, create the required CVM disk group, volume, and file system.
- 10 Create the database on the file system you created in the previous step.
- 11 Configure the VCS service groups for the database.
- 12 Verify that all VCS service groups are online.

## Configuring a cluster at the secondary site

To set up a multi-node or single-node cluster on the secondary site:

- Set up the cluster
- Set up the database

The setup requirements for the secondary site parallel the requirements for the primary site with a few additions or exceptions as noted below.

Important requirements for global clustering:

- Cluster names on the primary and secondary sites must be unique.
- Make sure that you use the same OS user and group IDs for Oracle for installation and configuration on both the primary and secondary clusters.

## Setting up the cluster on the secondary site

**To set up the cluster on secondary site**

- 1 Install and configure servers and storage.
- 2 If you are using hardware-based replication, install the software for managing your array.
- 3 Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster. If you are using VVR for replication, you must have it enabled.
- 4 Install and configure SF Oracle RAC. Prepare for your installation according to your configuration needs.

For preparation:

For installation:

See [“About installing SF Oracle RAC ”](#) on page 91.

- 5 For a multi-node cluster, configure I/O fencing.
  - Verify the shared storage on the secondary site supports SCSI-3 reservations.
  - Set up coordinator disks
  - Configure I/O fencing

For instructions for setting up fencing:

See [“About planning to configure I/O fencing”](#) on page 67.

- 6 For a single-node cluster, do not enable I/O fencing. Fencing will run in disabled mode.
- 7 Prepare systems and storage for a global cluster. Identify the hardware and storage requirements before installing Oracle RAC Clusterware and RDBMS software.

You will need to set up:

- Local storage for Oracle RAC and CRS binaries
- Shared storage for OCR and Vote disk which is not replicated
- Replicated storage for database files

- 8 Install and configure the Oracle RAC binaries:

See [“About preparing to install Oracle RAC”](#) on page 327.

See [“About installing Oracle RAC”](#) on page 381.

---

**Note:** OCR and Vote disk must be on non-replicated shared storage.

---

After successful Oracle RAC installation and configuration, verify that CRS daemons and resources are up on all nodes.

```
$ crs_stat -t
```

## Setting up the database for the secondary site

### To set up the database for the secondary site

- 1 Do not create the database. The database will be replicated from the primary site.
  - If you are using hardware-based replication, the database, disk group, and volumes will be replicated from the primary site.  
Create the directory for the CFS mount point which will host the database data and control files.
  - If you are using VVR for replication, create an identical disk group and volumes for the replicated content with the same names and size as listed on the primary site.



Create the directories for the CFS mount points as they are on the primary site. These will be used to host the database and control files when the failover occurs and the secondary is promoted to become the primary site.

- 2 Copy the `init$ORACLE_SID.ora` file from `$ORACLE_HOME/dbs` at the primary to `$ORACLE_HOME/dbs` at the secondary.
- 3 Create the following subdirectories on the secondary site as you did on the primary site:

For Oracle RAC 10g:

```
$ mkdir -p /$ORACLE_BASE/admin/adump
$ mkdir -p /$ORACLE_BASE/admin/database_name/bdump
$ mkdir -p /$ORACLE_BASE/admin/database_name/cdump
$ mkdir -p /$ORACLE_BASE/admin/database_name/dpdump
$ mkdir -p /$ORACLE_BASE/admin/database_name/hdump
$ mkdir -p /$ORACLE_BASE/admin/database_name/udump
$ mkdir -p /$ORACLE_BASE/admin/database_name/pfile
```

For Oracle 11gR1:

```
$ mkdir -p $ORACLE_BASE/admin
$ mkdir -p $ORACLE_BASE/admin/$database_name
$ cd $ORACLE_BASE/admin/$database_name
$ mkdir adump dpdump hdump pfile
```

For oracle 11gR2 release only, on both the primary and secondary sites, edit the file:

```
$ORACLE_HOME/dbs/init$ORACLE_SID.ora
```

as

```
remote_listener = 'SCAN_NAME:1521'
SPFILE=<SPFILE NAME>
```

- 4 Configure listeners on the secondary site with same name as on primary. You can do this by one of the following methods:
  - Copy the `listener.ora` and `tnsnames.ora` files from the primary site and update the names as appropriate for the secondary site.
  - Use Oracle's `netca` utility to configure the `listener.ora` and `tnsnames.ora` files on the secondary site.

- 5 On the secondary site, register the database using the `srvctl` command as Oracle user.

```
$ srvctl add database -d database_name -o oracle_home -p sp_file
```

To prevent automatic database instance restart, change the Management policy for the database (automatic, manual) to MANUAL using the `srvctl` command:

```
$ srvctl modify database -d database_name -y manual
```

- 6 Register the instances using `srvctl` command. Execute the following command on each node:

```
$ srvctl add instance -d database_name \  
-i instance_name -n node-name
```

## Configuring replication on clusters at both sites

You must configure replication for the database files. Once replication is configured, make sure it is functioning correctly by testing before proceeding.

### To configure replication at both sites

- 1 If you are using hardware-based replication, make sure that the replication management software for managing replication is installed on all nodes in both clusters.
- 2 At both sites, identify the disks on which the database resides at the primary site and associate them with the corresponding disks at the secondary site.

For VVR:

See [“Setting up replication using VVR on the primary site”](#) on page 578.

For Hardware-based replication:

See your hardware documentation for details on setting up replication between the two sites.

- 3 Start replication between the sites.

For VVR:

See [“Starting replication of Oracle RAC database volume”](#) on page 585.

See your hardware documentation for the appropriate procedures.

## Modifying the ClusterService group for global clusters

You have configured VCS service groups for SF Oracle RAC on each cluster. Each cluster requires an additional virtual IP address associated with the cluster for cross-cluster communication. The VCS installation and creation of the ClusterService group typically involves defining this IP address.

Configure a global cluster by setting:

- Heartbeat
- Wide area cluster (wac)
- GCO IP (gcoip)
- remote cluster resources

See the *Veritas Cluster Server User's Guide* for complete details on global clustering.

## Modifying the global clustering configuration using the wizard

The global clustering wizard completes the following tasks:

- Validates the ability of the current configuration to support a global cluster environment.
- Creates the components that enable the separate clusters, each of which contains a different set of GAB memberships, to connect and operate as a single unit.
- Creates the ClusterService group, or updates an existing ClusterService group.

Run the global clustering configuration wizard on each of the clusters; you must have the global clustering license in place on each node in the cluster.

**To modify the ClusterService group for global clusters using the global clustering wizard**

- 1 On the primary cluster, start the GCO Configuration wizard:

```
# /opt/VRTSvcs/bin/gcoconfig
```

- 2 The wizard discovers the NIC devices on the local system and prompts you to enter the device to be used for the global cluster. Specify the name of the device and press Enter.
- 3 If you do not have NIC resources in your configuration, the wizard asks you whether the specified NIC will be the public NIC used by all the systems. Enter y if it is the public NIC; otherwise enter n. If you entered n, the wizard prompts you to enter the names of NICs on all systems.

- 4 Enter the virtual IP address for the local cluster.
- 5 If you do not have IP resources in your configuration, the wizard prompts you for the netmask associated with the virtual IP. The wizard detects the netmask; you can accept the suggested value or enter another one.

The wizard starts running commands to create or update the ClusterService group. Various messages indicate the status of these commands. After running these commands, the wizard brings the ClusterService group online.

## Defining the remote cluster and heartbeat objects

After configuring global clustering, add the remote cluster object to define the IP address of the cluster on the secondary site, and the heartbeat object to define the cluster-to-cluster heartbeat.

Heartbeats monitor the health of remote clusters. VCS can communicate with the remote cluster only after you set up the heartbeat resource on both clusters.

### To define the remote cluster and heartbeat

- 1 On the primary site, enable write access to the configuration:

```
# haconf -makerw
```

- 2 Define the remote cluster and its virtual IP address.

In this example, the remote cluster is rac\_cluster102 and its IP address is 10.11.10.102:

```
# haclus -add rac_cluster102 10.11.10.102
```

- 3 Complete step 1 and step 2 on the secondary site using the name and IP address of the primary cluster.

In this example, the primary cluster is rac\_cluster101 and its IP address is 10.10.10.101:

```
# haclus -add rac_cluster101 10.10.10.101
```

- 4 On the primary site, add the heartbeat object for the cluster. In this example, the heartbeat method is ICMP ping.

```
# hahb -add Icmp
```

- 5 Define the following attributes for the heartbeat resource:

- ClusterList lists the remote cluster.
- Arguments enables you to define the virtual IP address for the remote cluster.

For example:

```
# hahb -modify Icmp ClusterList rac_cluster102
# hahb -modify Icmp Arguments 10.11.10.102 -clus rac_cluster102
```

- 6 Save the configuration and change the access to read-only on the local cluster:

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

- 7 Complete step 4-6 on the secondary site using appropriate values to define the cluster on the primary site and its IP as the remote cluster for the secondary cluster.

- 8 Verify cluster status with the `hastatus -sum` command on both clusters.

```
# hastatus -sum
```

## 9 Display the global setup by executing `haclus -list` command.

```
# haclus -list
    rac_cluster101
    rac_cluster102
```

Example of heartbeat additions to the `main.cf` file on the primary site:

```
.
.
remotecluster rac_cluster102 (
Cluster Address = "10.11.10.102"
)
heartbeat Icmp (
    ClusterList = { rac_cluster102 }
    Arguments @rac_cluster102 = { "10.11.10.102" }
)

system galaxy (
)

.
.
```

Example heartbeat additions to the `main.cf` file on the secondary site:

```
.
.
remotecluster rac_cluster101 (
    Cluster Address = "10.10.10.101"
)

heartbeat Icmp (
    ClusterList = { rac_cluster101 }
    Arguments @rac_cluster101 = { "10.10.10.101" }
)

system mercury (
)

.
.
```

See the *Veritas Cluster Server User's Guide* for details for configuring the required and optional attributes of the heartbeat object.

# Configuring the VCS service groups for global clusters

## To configure VCS service groups for global clusters

- 1 Configure and enable global groups for databases and resources.
    - Configure VCS service groups at both sites.
    - Configure the replication agent at both sites.
    - Make the RAC service group a global service group, enabling failover across clusters.
    - For example:
  - 2 To test real data in an environment where HA/DR has been configured, schedule a planned migration to the secondary site for testing purposes.  
 For example:  
 See [“Migrating the role of primary site to the secondary site”](#) on page 601.  
 See [“Migrating the role of new primary site back to the original primary site”](#) on page 602.
  - 3 Upon successful testing, bring the environment into production.  
 For more information see:
    - *Veritas Cluster Server Installation Guide*
    - *Veritas Cluster Server Agent for EMC SRDF Installation and Configuration Guide*
    - *Veritas Cluster Server Agent for Hitachi TrueCopy Installation and Configuration Guide*
    - *Veritas Cluster Server Agent for IBM Metro Mirror Installation and Configuration Guide*
    - *Veritas Cluster Server Agent for EMC MirrowView Installation and Configuration Guide*
    - *Veritas Cluster Server Agent for IBM SVC Installation and Configuration Guide*
    - *Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide*
    - Technical Support TechNote for the latest updates or software issues for replication agents:  
<http://entsupport.symantec.com/docs/282004htm>
- For complete details on VVR in a shared disk environment:  
 See the *Veritas Volume Replicator Administrator's Guide*.





# Configuring a global cluster using VVR

This chapter includes the following topics:

- [About configuring global clustering using VVR](#)
- [Setting up replication using VVR on the primary site](#)
- [Setting up replication using VVR on the secondary site](#)
- [Starting replication of Oracle RAC database volume](#)
- [Configuring VCS to replicate the database volume using VVR](#)
- [Using VCS commands on SF Oracle RAC global clusters](#)
- [Using VVR commands on SF Oracle RAC global clusters](#)

## About configuring global clustering using VVR

Before configuring clusters for global clustering, make sure both clusters have product and database software installed and configured.

Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster and VVR enabled.

See [“About Veritas product licensing”](#) on page 79.

After setting up two clusters running SF Oracle RAC, you can configure a global cluster environment with VVR. You must modify both cluster configurations to support replication in the global cluster environment.

Configuring SF Oracle RAC for global clusters requires:

- Setting up both clusters as part of a global cluster environment.  
See [“About setting up a global cluster in an SF Oracle RAC environment”](#) on page 564.
- Setting up replication for clusters at both sites.  
See [“Setting up replication using VVR on the primary site”](#) on page 578.  
See [“Setting up replication using VVR on the secondary site”](#) on page 581.
- Starting replication of the database.  
See [“Starting replication of Oracle RAC database volume”](#) on page 585.
- Configuring VCS for replication on clusters at both sites.  
See [“Configuring VCS to replicate the database volume using VVR”](#) on page 587.

## Setting up replication using VVR on the primary site

Setting up replication with VVR in a global cluster environment involves the following tasks:

- If you have not already done so, create a disk group on the storage on the primary site.  
For example:
- Creating the Storage Replicator Log (SRL) in the disk group for the database.  
See [“Creating the SRL volume on the primary site”](#) on page 578.
- Creating the Replicated Volume Group (RVG) on the primary site.  
See [“Setting up the Replicated Volume Group \(RVG\) on the primary site”](#) on page 579.

### Creating the SRL volume on the primary site

Create the SRL. The SRL is a volume in the RVG. The RVG also holds the data volumes for replication.

- The data volume on the secondary site has the same name and the same size as the data volume on the primary site.
- The SRL on the secondary site has the same name and the same size as the SRL on the primary site.
- The data volume and SRL volume should exist in the same disk group.
- If possible, create SRLs on disks without other volumes.
- Mirror SRLs and data volumes in the absence of hardware-based mirroring.

**To create the SRL volume on the primary site**

- 1 On the primary site, determine the size of the SRL volume based on the configuration and amount of use.  
  
See the Veritas Volume Replicator documentation for details.
- 2 Using the following command, determine whether a node is the master or the slave:

```
# vxctl -c mode
```

- 3 From the master node, issue the following command:

```
# vxassist -g oradatadg make rac1_srl 1500M nmirror=2 disk4 disk5  
# vxassist -g sybdata_101 make syb_srl 1500M nmirror=2 disk4 disk5
```

- 4 Using the following command, start the SRL volume by starting all volumes in the disk group:

```
# vxvol -g oradatadg startall
```

## Setting up the Replicated Volume Group (RVG) on the primary site

Before creating the RVG on the primary site, make sure the volumes and CVM group are active and online.

### To review the status of replication objects on the primary site

- 1 Verify the volumes you intend to include in the group are active.
- 2 Review the output of the `hagrp -state cvm` command to verify that the CVM group is online.
- 3 On each site, verify `vradmin` is running:

```
# ps -ef |grep vradmin
root 536594 598036 0 12:31:25 0 0:00 grep vradmin
```

If `vradmin` is not running start it:

```
# vxstart_vvr
VxVM VVR INFO V-5-2-3935 Using following ports:
heartbeat: 4145
vradmin: 8199
vxrsyncd: 8989
data: Anonymous-Ports
To change, see vrport(1M) command
# ps -ef |grep vradmin
root 536782 1 0 12:32:47 - 0:00 /usr/sbin/vradmin
root 1048622 598036 0 12:32:55 0 0:00 grep vradmin
# netstat -an |grep 4145
tcp4 0 0 *.4145 *.* LISTEN
udp4 0 0 *.4145 *.*
```

### To create the RVG

The command to create the primary RVG takes the form:

```
vradmin -g disk_group createpri rvg_name data_volume srl_volume
```

where:

- `disk_group` is the name of the disk group containing the database
- `rvg_name` is the name for the RVG
- `data_volume` is the volume that VVR replicates
- `srl_volume` is the volume for the SRL

For example, to create the `rac1_rvg` RVG, enter:

```
# vradmin -g oradatadg createpri rac1_rvg rac1_vol rac1_srl
```

The command creates the RVG on the primary site and adds a Data Change Map (DCM) for each data volume. In this case, a DCM exists for `rac1_vol`.

## Setting up replication using VVR on the secondary site

To create objects for replication on the secondary site, use the `vradmin` command with the `addsec` option. To set up replication on the secondary site, perform the following tasks:

- If you have not already done so, create a disk group to hold data volume, SRL, and RVG on the storage on the secondary site.  
For example:
- Create volumes for the database and SRL on the secondary site.  
See [“Creating the data and SRL volumes on the secondary site”](#) on page 581.
- Edit the `/etc/vx/vras/.rdg` file on the secondary site.  
See [“Editing the `/etc/vx/vras/.rdg` files”](#) on page 582.
- Use resolvable virtual IP addresses that set network RLINK connections as host names of the primary and secondary sites.  
See [“Setting up IP addresses for RLINKs on each cluster”](#) on page 582.
- Create the replication objects on the secondary site.  
See [“Setting up the disk group on secondary site for replication”](#) on page 583.

### Creating the data and SRL volumes on the secondary site

Note the following when creating volumes for the data and SRL:

- The sizes and names of the volumes must reflect the sizes and names of the corresponding volumes in the primary site.
- Create the data and SRL volumes on different disks in the disk group. Use the `vxdisk -g diskgroup list` command to list the disks in the disk group.
- Mirror the volumes.

To create the data and SRL volumes on the secondary site

- 1 In the disk group created for the Oracle RAC database, create a data volume of same size as that in primary for data; in this case, the `rac_vol1` volume on the primary site is 6.6 GB:

```
# vxassist -g oradatadg make rac_vol1 6600M nmirror=2 disk1 disk2
```

- 2 Create the volume for the SRL, using the same name and size of the equivalent volume on the primary site. Create the volume on different disks from the disks for the database volume, but on the same disk group that has the data volume:

```
# vxassist -g oradatadg make rac1_srl 1500M nmirror=2 disk4 disk6
```

## Editing the /etc/vx/vras/.rdg files

Editing the /etc/vx/vras/.rdg file on the secondary site enables VVR to replicate the disk group from the primary site to the secondary site. On each node, VVR uses the /etc/vx/vras/.rdg file to check the authorization to replicate the RVG on the primary site to the secondary site. The file on each node in the secondary site must contain the primary disk group ID, and likewise, the file on each primary system must contain the secondary disk group ID.

### To edit the /etc/vx/vras/.rdg files

- 1 On a node in the primary site, display the primary disk group ID:

```
# vxprint -l diskgroup

.....
```

- 2 On each node in the secondary site, edit the /etc/vx/vras/.rdg file and enter the primary disk group ID on a single line.
- 3 On each cluster node of the primary cluster, edit the /etc/vx/vras/.rdg file and enter the secondary disk group ID on a single line.

## Setting up IP addresses for RLINKs on each cluster

Creating objects with the vradm command requires resolvable virtual IP addresses that set network RLINK connections as host names of the primary and secondary sites.

### To set up IP addresses for RLINKs on each cluster

- 1 For each RVG running on each cluster, set up a virtual IP address on one of the nodes of the cluster. These IP addresses are part of the RLINK.

The example assumes for the cluster on the primary site:

- The public network interface is bge0:1
- The virtual IP address is 10.10.9.101

- The net mask is 255.255.255.0
  - ```
# ifconfig bge0:1 plumb
# ifconfig bge0:1 inet 10.10.9.101 netmask 255.255.255.0
# ifconfig bge0:1 up
```
- 2** Use the same commands with appropriate values for the interface, IP address, and net mask on the secondary site.
- The example assumes for the secondary site:
- The public network interface is bge0:1
  - virtual IP address is 10.11.9.102
  - net mask is 255.255.255.0
- 3** Define the virtual IP addresses to correspond to a virtual cluster host name on the primary site and a virtual cluster host name on the secondary site.
- Update the `/etc/hosts` file on all the nodes on both the primary and secondary sites.
- The examples assume:
- `rac_cluster101` has IP address 10.10.9.101
  - `rac_cluster102` has IP address 10.11.9.102
- 4** Use the ping command to verify the links are functional.

## Setting up the disk group on secondary site for replication

Create the replication objects on the secondary site from the master node of the primary site, using the `vradmin` command.

### To set up the disk group on the secondary site for replication

- 1** Issue the command in the following format from the cluster on the primary site:

```
# vradmin -g dg_pri addsec rvg_pri pri_host sec_host
```

where:

- `dg_pri` is the disk group on the primary site that VVR will replicate. For example: `rac1_vol`
- `rvg_pri` is the RVG on the primary site. For example: `rac1_rvg`

- `pri_host` is the virtual IP address or resolvable virtual host name of the cluster on the primary site.  
For example: `rac_cluster101_1`
- `sec_host` is the virtual IP address or resolvable virtual host name of the cluster on the secondary site.  
For example: `rac_cluster102_1`

For example, the command to add the cluster on the primary site to the Replicated Data Set (RDS) is:

```
vradmin -g oradatadg addsec rac1_rvg \  
rac_cluster101_1  
rac_cluster102_1
```

On the secondary site, the above command performs the following tasks:

- Creates an RVG within the specified disk group using the same name as the one for the primary site
- Associates the data and SRL volumes that have the same names as the ones on the primary site with the specified RVG
- Adds a data change map (DCM) for the data volume
- Creates cluster RLINKS for the primary and secondary sites with the default names; for example, the "primary" RLINK created for this example is `rlk_rac_cluster101_1_rac1_rvg` and the "secondary" RLINK created is `rlk_rac_cluster102_1_rac1_rvg`.

- 2 Verify the list of RVGs in the RDS by executing the following command.

```
# vradmin -g oradatadg -l printrvg
```

For example:

```
Replicated Data Set: rac1_rvg  
Primary:  
HostName: 10.180.88.187 <localhost>  
RvgName: rac1_rvg  
DgName: rac1_vol  
datavol_cnt: 1  
vset_cnt: 0  
srl: rac1_srl  
RLinks:  
name=rlk_rac_cluster102_1_rac1_rvg, detached=on,
```



```
synchronous=off
Secondary:
HostName: 10.190.99.197
RvgName: rac1_rvg
DgName: oradatadg
datavol_cnt: 1
vset_cnt: 0
srl: rac1_srl
RLinks:
name=rlk_rac_cluster101_1_rac1_rvg, detached=on,
synchronous=off
```

---

**Note:** Once the replication is started the value of the detached flag will change the status from OFF to ON.

---

## Starting replication of Oracle RAC database volume

When you have both the primary and secondary sites set up for replication, you can start replication from the primary site to the secondary site.

Start with the default replication settings:

- Mode of replication: `synchronous=off`
- Latency Protection: `latencyprot=off`
- SRL overflow protection: `srlprot_autodcm`
- Packet size: `packet_size=8400`
- Network protocol: `protocol=UDP`

Method of initial synchronization:

- Automatic synchronization
- Full synchronization with Checkpoint

For guidelines on modifying these settings and information on choosing the method of replication for the initial synchronization:

See the *Veritas Volume Replicator Administrator's Guide*

## Starting replication using automatic synchronization

Use the `vradmin` command to start replication or the transfer of data from the primary site to the secondary site over the network. Because the cluster on the

secondary site uses only one host name, the command does not require the *sec\_host* argument.

#### To start replication using automatic synchronization

- ◆ From the primary site, use the following command to automatically synchronize the RVG on the secondary site:

```
vradmin -g disk_group -a startrep pri_rvg sec_host
```

where:

- *disk\_group* is the disk group on the primary site that VVR will replicate
- *pri\_rvg* is the name of the RVG on the primary site
- *sec\_host* is the virtual host name for the secondary site

For example:

```
# vradmin -g oradatadg -a startrep rac1_rvg  
          rac_cluster102
```

## Starting replication using full synchronization with Checkpoint

Use the `vradmin` command with the Checkpoint option to start replication using full synchronization with Checkpoint.

#### To start replication using full synchronization with Checkpoint

- 1 From the primary site, synchronize the RVG on the secondary site with full synchronization (using the `-c checkpoint` option):

```
vradmin -g disk_group -full -c ckpt_name syncrvg pri_rvg sec_host
```

where:

- *disk\_group* is the disk group on the primary site that VVR will replicate
- *ckpt\_name* is the name of the checkpoint on the primary site
- *pri\_rvg* is the name of the RVG on the primary site
- *sec\_host* is the virtual host name for the secondary site

For example:

```
# vradmin -g oradatadg -c rac1_ckpt syncrvg rac1_rvg
rac_cluster102
```

- 2 To start replication after full synchronization, enter the following command:

```
# vradmin -g oradatadg -c rac1_ckpt startrep rac1_rvg
rac_cluster102
```

## Verifying replication status

Verify that replication is properly functioning.

### To verify replication status

- 1 Check the status of VVR replication:

```
# vradmin -g disk_group_name repstatus rvg_name
```

- 2 Review the `flags` output for the status. The output may appear as `connected` and `consistent`. For example:

```
# vxprint -g oradatadg -l rlk_rac_cluster102_1_rac1_rvg
Rlink: rlk_rac_cluster102_1_rac1_rvg
info: timeout=500 packet_size=8400 rid=0.1078
      latency_high_mark=10000 latency_low_mark=9950
      bandwidth_limit=none
state: state=ACTIVE
      synchronous=off latencyprot=off srlprot=autodcm
.
.
protocol: UDP/IP
checkpoint: rac1_ckpt
flags: write enabled attached consistent connected
asynchronous
```

## Configuring VCS to replicate the database volume using VVR

After configuring both clusters for global clustering and setting up the Oracle RAC database for replication, configure VCS to provide high availability for the database. Specifically, configure VCS agents to control the cluster resources, including the replication resources.

To view the sample main.cf files on your system:

```
# cd /etc/VRTSvcs/conf/sample_rac
```

The following files include CVM/VVR configuration examples:

```
sfrac_07_main.cf      sfrac_08_main.cf
```

The following sample main.cf files illustrate the VCS configuration changes after setting up an existing Oracle RAC database for replication:

See “[sfrac07\\_main.cf and sample08\\_main.cf files](#)” on page 683.

## About modifying the VCS configuration for replication

The following resources must be configured or modified for replication:

- Log owner group
- RVG group
- CVMVolDg resource
- RVGSharedPri resource
- Oracle RAC database service group

For detailed examples of service group modification:

For more information on service replication resources:

See the *Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide*.

### Log owner group

Create a log owner group including the RVGLogowner resources. The RVGLogowner resources are used by:

- RLINKs for the RVG
- RVGLogowner resource. The RVG and its associated disk group are defined as attributes for the RVGLogowner resource.

The RVG log owner service group has an online local firm dependency on the service group containing the RVG.

The VCS uses the following agents to control the following resources:

- RVGLogowner agent to control the RVGLogowner resource
- RVGShared agent to control the RVGShared resource

## RVG group

Create an RVG group that includes the RVGShared resource replication objects. Define the RVGShared resource and CVMVolDg resource together within a parallel service group. The group is defined as parallel because it may be online at the same time on all cluster nodes.

## CVMVolDg resource

The CVMVolDg resource does not have volumes specified for the CVMVolume attribute; the volumes are contained in the RVG resource. The CVMVolume attribute for the CVMVolDg resource is empty because all volumes in the RVG are defined by the RVG attribute of the RVGShared resource. The RVG service group has an online local firm dependency on the CVM service group.

For a detailed description of the CVMVolDg agent in this guide:

See “[CVMVolDg agent](#)” on page 732.

## RVGSharedPri resource

Add the RVGSharedPri resource to the existing Oracle RAC database service group. The CVMVolDg resource must be removed from the existing Oracle RAC database service group.

## Oracle RAC database service group

The existing Oracle RAC database service group is a parallel group consisting of the Oracle RAC database resource, CVMVolDg resource, and CFSMount resource (if the database resides in a cluster file system). Define the Oracle RAC service group as a global group by specifying the clusters on the primary and secondary sites as values for the ClusterList group attribute.

## Configuration examples before and after modification

Review the following illustrations that display the changes to the VCS configuration, after setting up replication on the existing Oracle RAC database.

- Configuration before modification:

[Figure 33-1](#)


- Configuration after modification:

[Figure 33-2](#)

All of the dependencies between parent and child groups are online local firm. The CVM service group is the same in all illustrations because its definition requires no changes.


Configuration before modification for replication:

**Figure 33-1** Illustration of dependencies before modification for replication (Oracle RAC 10g)



Configuration after modification for replication:

**Figure 33-2** Illustration of dependencies after modification for replication (Oracle RAC 10g)



## Modifying the VCS Configuration on the Primary Site

The following are the procedural highlights required to modify the existing VCS configuration on the primary site:

- Configure two service groups:
  - A log owner group including the RVGLogowner resource.
  - An RVG group including the RVGShared resource replication objects.
- Add the RVGSharedPri resource to the existing Oracle RAC database service group and define this group as a global group by setting the ClusterList and ClusterFailOverPolicy attributes.
- Move the CVMVolDg resource from the existing Oracle RAC database service group to the newly created RVG group.

### To modify VCS on the primary site

- 1 Log into one of the nodes on the primary cluster.
- 2 Use the following command to save the existing configuration to disk, and make the configuration read-only while you make changes:

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

- 3 Use the following command to make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

- 4 Use vi or another text editor to edit the main.cf file. Review the sample configuration file after the SF Oracle RAC installation.

Add a failover service group using the appropriate values for your cluster and nodes. Include the following resources:

- RVGLogowner resource. The node on which the group is online functions as the log owner (node connected to the second cluster for the purpose of replicating data).
- IP resource
- NIC resources

The following are examples of RVGLogowner service group for the different platforms.

```
group rlogowner (
    SystemList = { galaxy = 0, nebula = 1 }
```



```

AutoStartList = { galaxy,nebula }
)

IP logowner_ip (
    Device = bge0
    Address = "10.10.9.101"
    NetMask = "255.255.255.0"
)

NIC nic (
    Device = bge0
    NetworkType = ether
    NetworkHosts = "10.10.8.1"
)

RVGLogowner logowner (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)
requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic

```

## 5 Add the RVG service group using the appropriate values for your cluster and nodes.

**Example RVGgroup service group:**

```

group RVGgroup (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    AutoStartList = { galaxy,nebula }
)

RVGShared racdata_rvg (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)
CVMVolDg racdata_voldg (
    CVMDiskGroup = oradatadg
    CVMActivation = sw
)
requires group cvm online local firm
racdata_rvg requires racdata_voldg

```

- 6 Modify the Oracle RAC service group using the appropriate values for your cluster and nodes:
  - Define the Oracle RAC service group as a global group by specifying the clusters on the primary and secondary sites as values for the ClusterList group attribute. See the bolded attribute in the example that follows.
  - Add the ClusterFailOverPolicy cluster attribute. Symantec recommends using the Manual value. See the bolded attribute in the example.
  - Add the RVGSharedPri resource to the group configuration.
  - Remove the CVMVolDg resource, if it has been configured in your previous configuration. This resource is now part of the RVG service group.
  - Specify the service group (online, local, firm) to depend on the RVG service group.
  - Remove the existing dependency of the Database service group on the CVM service group. Remove the line:

```
requires group CVM online local firm
```

- Remove the existing dependency between the CFSMount for the database and the CVMVoldg for the Oracle RAC database. Remove the line:

```
oradata_mnt requires oradata_voldg
```

The following is an example of an Oracle RAC database service group configured for replication:

```
group database_grp (
    SystemList = { galaxy = 0, nebula = 1 }
    ClusterList = { rac_cluster101 = 0, rac_cluster102 = 1 }
    Parallel = 1
    ClusterFailOverPolicy = Manual
    Authority = 1
    AutoStartList = { galaxy,nebula }
)

CFSMount oradata_mnt (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/racdb_vol"
)

RVGSharedPri ora_vvr_shpri (
    RvgResourceName = racdata_rvg
```

```

        OnlineRetryLimit = 0
    )

Oracle rac_db (
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Owner = Oracle
    Home = "/oracle/orahome"

    StartUpOpt = SRVCTLSTART
    ShutDownOpt = SRVCTLSTOP
)

requires group RVGgroup online local firm
oradata_mnt requires ora_vvr_shpri
rac_db requires oradata_mnt

```

- 7 Save and close the main.cf file.
- 8 Use the following command to verify the syntax of the /etc/VRTSvcs/conf/config/main.cf file:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

- 9 Stop and restart VCS.

```
# hstop -all -force
```

Wait for port h to stop on all nodes, and then restart VCS with the new configuration on all primary nodes:

```
# hstart
```

## Modifying the VCS Configuration on the Secondary Site

The following are highlights of the procedure to modify the existing VCS configuration on the secondary site:

- Add the log owner and RVG service groups.
- Add a service group to manage the Oracle RAC database and the supporting resources.
- Define the replication objects and agents, such that the cluster at the secondary site can function as a companion to the primary cluster.

The following steps are similar to those performed on the primary site.

### To modify VCS on the secondary site

- 1 Log into one of the nodes on the secondary site as root.
- 2 Use the following command to save the existing configuration to disk, and make the configuration read-only while making changes:

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

- 3 Use the following command to make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

- 4 Use vi or another text editor to edit the main.cf file. Edit the CVM group on the secondary site.

Review the sample configuration file after the SF Oracle RAC installation to see the CVM configuration.

In our example, the secondary site has rac\_cluster102 consisting of the nodes mercury and jupiter. To modify the CVM service group on the secondary site, use the CVM group on the primary site as your guide.

- 5 Add a failover service group using the appropriate values for your cluster and nodes. Include the following resources:
  - RVGLogowner resource. The node on which the group is online functions as the log owner (node connected to the second cluster for the purpose of replicating data).
  - IP resource
  - NIC resources

**Example RVGLogowner service group:**

```
group rlogowner (
    SystemList = { mercury = 0, jupiter = 1 }
    AutoStartList = { mercury, jupiter }
)

IP logowner_ip (
    Device = bge0
    Address = "10.11.9.102"
    NetMask = "255.255.255.0"
)
```

```
NIC nic (
    Device = bge0
    NetworkHosts = { "10.10.8.1" }
    NetworkType = ether
)

RVGLogowner logowner (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)

requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic
```

- 6** Add the RVG service group using the appropriate values for your cluster and nodes.

The following is an example RVGgroup service group:

```
group RVGgroup (
    SystemList = { mercury = 0, jupiter = 1 }
    Parallel = 1
    AutoStartList = { mercury, jupiter }
)

RVGShared racdata_rvg (
    RVG = rac1_rvg
    DiskGroup = oradatadg
)

CVMVolDg racdata_voldg
    CVMDiskGroup = oradatadg
    CVMActivation = sw
)

requires group cvm online local firm
racdata_rvg requires racdata_voldg
```

- 7** Add an Oracle RAC service group. Use the Oracle RAC service group on the primary site as a model for the Oracle RAC service group on the secondary site.

- Define the Oracle RAC service group as a global group by specifying the clusters on the primary and secondary sites as values for the ClusterList group attribute.
- Assign this global group the same name as the group on the primary site; for example, *database\_grp*.
- Include the ClusterList and ClusterFailOverPolicy cluster attributes. Symantec recommends using the Manual value.
- Add the RVGSharedPri resource to the group configuration.
- Remove the CVMVolDg resource, if it has been configured in your previous configuration. This resource is now part of the RVG service group.
- Specify the service group to depend (online, local, firm) on the RVG service group.

Example of the Oracle RAC group on the secondary site:

```
group database_grp (
    SystemList = { mercury = 0, jupiter = 1 }
    ClusterList = { rac_cluster102 = 0, rac_cluster101 = 1 }
    Parallel = 1
    OnlineRetryInterval = 300
    ClusterFailOverPolicy = Manual
    Authority = 1
    AutoStartList = { mercury, jupiter }
)

RVGSharedPri ora_vvr_shpri (
    RvgResourceName = racdata_rvg
    OnlineRetryLimit = 0
)

CFSMount oradata_mnt (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/racdb_vol"
    Critical = 0
)

RVGSharedPri ora_vvr_shpri (
    RvgResourceName = racdata_rvg
    OnlineRetryLimit = 0
)
```

```

Oracle rac_db (
    Sid @mercury = vrts1
    Sid @jupiter = vrts2
    Owner = Oracle
    Home = "/oracle/orahome"
    Pfile @mercury = "/oracle/orahome/dbs/initvrts1.ora"
    Pfile @jupiter = "/oracle/orahome/dbs/initvrts2.ora"
    StartUpOpt = SRVCTLSTART
    ShutDownOpt = SRVCTLSTOP
)

requires group RVGgroup online local firm
oradata_mnt requires ora_vvr_shpri
rac_db requires oradata_mnt

```

**8** Save and close the `main.cf` file.

**9** Use the following command to verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

**10** Stop and restart VCS.

```
# hstop -all -force
```

Wait for port `h` to stop on all nodes, and then restart VCS with the new configuration on all primary nodes:

```
# hstart
```

- 11 Verify that VCS brings all resources online. On one node, enter the following command:

```
# hagrps -display
```

The Oracle RAC, RVG, and CVM groups are online on both nodes of the primary site. The RVGLogOwner group is online on one node of the cluster. If either the RVG group or the RVGLogOwner group is partially online, manually bring the groups online using the `hagrps -online` command. This information applies to the secondary site, except for the Oracle RAC group which must be offline.

- 12 Verify the service groups and their resources that are brought online. On one node, enter the following command:

```
# hagrps -display
```

The Oracle RAC service group is offline on the secondary site, but the CVM, RVG log owner, and RVG groups are online.

This completes the setup for an SF Oracle RAC global cluster using VVR for replication. Symantec recommends testing a global cluster before putting it into production.

## Using VCS commands on SF Oracle RAC global clusters

For information on the VCS commands for global clusters:

See the *Veritas Cluster Server Administrator's Guide*.

## Using VVR commands on SF Oracle RAC global clusters

If you have two SF Oracle RAC clusters configured to use VVR for replication, the following administrative functions are available:

- Migration of the role of the primary site to the remote site
- Takeover of the primary site role by the secondary site

### About migration and takeover of the primary site role

Migration is a planned transfer of the role of primary replication host from one cluster to a remote cluster. This transfer enables the application on the remote



cluster to actively use the replicated data. The former primary cluster becomes free for maintenance or other activity.

Takeover occurs when an unplanned event (such as a disaster) causes a failure, making it necessary for the applications using the replicated data to be brought online on the remote cluster.

## Migrating the role of primary site to the secondary site

After configuring the replication objects within VCS, you can use VCS commands to migrate the role of the cluster on the primary site to the remote cluster. In the procedure below, VCS takes the replicated database service group, *database\_grp*, offline on the primary site and brings it online on the secondary site; the secondary site now assumes the role of the primary site.

---

**Note:** The `hagrp -switch` command cannot migrate a parallel group within a cluster or between clusters in a global cluster environment.

---

### To migrate the role of primary site to the remote site

- 1 From the primary site, use the following command to take the Oracle service group offline on all nodes.

```
# hagrp -offline database_grp -any
```

Wait for VCS to take all Oracle service groups offline on the primary site.

- 2 Verify that the RLINK between the primary and secondary is up to date. Use the `vxrlink -g` command with the status option and specify the RLINK for the primary cluster. You can use the command from any node on the primary cluster.

For example:

```
# vxrlink -g data_disk_group status rlk_rac_cluster101_priv_rac1_rvg
```

Where `rlk_rac_cluster101_priv_rac1_rvg` is the RLINK.

- 3 On the secondary site, which is now the new primary site, bring the Oracle service group online on all nodes:

```
# hagrp -online database_grp -any
```

## Migrating the role of new primary site back to the original primary site

After migrating the role of the primary site to the secondary site, you can use VCS commands to migrate the role of the cluster on the new primary site to the original primary site. In the procedure below, VCS takes the replicated database service group, *database\_grp*, offline on the new primary (former secondary) site and brings it online on the original primary site; the original primary site now resumes the role of the primary site.

---

**Note:** The `hagrp -switch` command cannot migrate a parallel group within a cluster or between clusters in a global cluster environment.

---

### To migrate the role of new primary site back to the original primary site

- 1 Make sure that all CRS resources are online, and switch back the group *database\_grp* to the original primary site.

Issue the following command on the remote site:

```
# hagrp -offline database_grp -any
```

- 2 Verify that the RLINK between the primary and secondary is up to date. Use the `vxrlink -g` command with the status option and specify the RLINK for the primary cluster. You can use the command from any node on the primary cluster.

For example:

```
# vxrlink -g data_disk_group status rlk_rac_cluster101_priv_rac1_rvg
```

Where `rlk_rac_cluster101_priv_rac1_rvg` is the RLINK.

- 3 Make sure that *database\_grp* is offline on the new primary site. Then, execute the following command on the original primary site to bring the *database\_grp* online:

```
# hagrp -online database_grp -any
```

## Taking over the primary role by the remote cluster

Takeover occurs when the remote cluster on the secondary site starts the application that uses replicated data. This situation may occur if the secondary

site perceives the primary site as dead, or when the primary site becomes inaccessible (perhaps for a known reason). For a detailed description of concepts of taking over the primary role:

See the *Veritas Volume Replicator Administrator's Guide*.

Before enabling the secondary site to take over the primary role, the administrator on the secondary site must "declare" the type of failure at the remote (primary, in this case) site and designate the failure type using one of the options for the `haclus` command.

Takeover options are:

- [Disaster](#)
- [Outage](#)
- [Disconnect](#)
- [Replica](#)

## Disaster

When the cluster on the primary site is inaccessible and appears dead, the administrator declares the failure type as "disaster." For example, fire may destroy a data center, including the primary site and all data in the volumes. After making this declaration, the administrator can bring the service group online on the secondary site, which now has the role as "primary" site.

## Outage

When the administrator of a secondary site knows the primary site is inaccessible for a known reason, such as a temporary power outage, the administrator may declare the failure as an "outage." Typically, an administrator expects the primary site to return to its original state.

After the declaration for an outage occurs, the RVGSharedPri agent enables DCM logging while the secondary site maintains the primary replication role. After the original primary site becomes alive and returns to its original state, DCM logging makes it possible to use fast fail back resynchronization when data is resynchronized to the original cluster.

Before attempting to resynchronize the data using the fast fail back option from the current primary site to the original primary site, take the precaution at the original primary site of making a snapshot of the original data. This action provides a valid copy of data at the original primary site for use in the case the current primary site fails before the resynchronization is complete.

See ["Examples for takeover and resynchronization"](#) on page 604.

See “[Replica](#)” on page 604.

## Disconnect

When both clusters are functioning properly and the heartbeat link between the clusters fails, a split-brain condition exists. In this case, the administrator can declare the failure as "disconnect," which means no attempt will occur to take over the role of the primary site at the secondary site. This declaration is merely advisory, generating a message in the VCS log indicating the failure results from a network outage rather than a server outage.

## Replica

In the rare case where the current primary site becomes inaccessible while data is resynchronized from that site to the original primary site using the fast fail back method, the administrator at the original primary site may resort to using a data snapshot (if it exists) taken before the start of the fast fail back operation. In this case, the failure type is designated as "replica".

## Examples for takeover and resynchronization

The examples illustrate the steps required for an outage takeover and resynchronization.

### To take over after an outage

- 1 From any node of the secondary site, issue the `haclus` command:

```
# haclus -declare outage -clus rac_cluster101
```

- 2 After declaring the state of the remote cluster, bring the Oracle service group online on the secondary site. For example:

```
# hagrps -online -force database_grp -any
```

### To resynchronize after an outage

- 1 On the original primary site, create a snapshot of the RVG before resynchronizing it in case the current primary site fails during the resynchronization. Assuming the disk group is *data\_disk\_group* and the RVG is *rac1\_rvg*, type:

```
# vxrvg -g data_disk_group -F snapshot rac1_rvg
```

See the *Veritas Volume Replicator Administrator's Guide* for details on RVG snapshots.

- 2 Resynchronize the RVG. From the CVM master node of the current primary site, issue the `hares` command and the `-action` option with the `fbsync` action token to resynchronize the `RVGSharedPri` resource. For example:

```
# hares -action ora_vvr_shpri fbsync -sys mercury
```

To determine which node is the CVM master node, type:

```
# vxdctl -c mode
```

- 3 Perform one of the following commands, depending on whether the resynchronization of data from the current primary site to the original primary site is successful:
  - If the resynchronization of data is successful, use the `vxrvg` command with the `snapback` option to reattach the snapshot volumes on the original primary site to the original volumes in the specified RVG:

```
# vxrvg -g data_disk_group snapback rac1_rvg
```

- A failed attempt at the resynchronization of data (for example, a disaster hits the primary RVG when resynchronization is in progress) could generate inconsistent data.  
You can restore the contents of the RVG data volumes from the snapshot taken in step 1:

```
# vxrvg -g data_disk_group snaprestore rac1_rvg
```

## Troubleshooting CVM and VVR components of SF Oracle RAC

The following topic is useful for troubleshooting the VVR component of SF Oracle RAC.

### Updating the rlink

If the rlink is not up to date, use the `hares -action` command with the `resync` action token to synchronize the RVG.

The following command example is issued on any node (`galaxy`, in this case) in the primary cluster, specifying the RVGSharedPri resource, `ora_vvr_shpri`:

```
# hares -action ora_vvr_shpri resync -sys galaxy
```

## VCS agents to manage wide-area failover

VCS agents now manage external objects that are part of wide-area failover. These objects include replication, DNS updates, and so on. These agents provide a robust framework for specifying attributes and restarts, and can be brought online upon fail over.

VCS provides agents for other array-based or application-based solutions. This section covers the replication agents that is bundled with VVR. See the VCS replication agent documentation for more details.

---

**Note:** See the Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide for more information about the RVG and RVGPrimary agents.

---

---

**Note:** The RVGSnapshot agent is not supported for SF Oracle RAC.

---

### DNS agent

The DNS agent updates the canonical name-mapping in the domain name server after a wide-area failover. See the Veritas Cluster Server Bundled Agents Reference Guide for more information about the agent.

### RVG agent

The RVG agent manages the Replicated Volume Group (RVG). Specifically, it brings the RVG online, monitors read-write access to the RVG, and takes the RVG offline. Use this agent when using VVR for replication.

## **RVGPrimary agent**

The RVGPrimary agent attempts to migrate or take over a Secondary to a Primary following an application failover. The agent has no actions associated with the offline and monitor routines.





## Uninstallation of SF Oracle RAC

- [Chapter 34. Preparing to uninstall SF Oracle RAC from a cluster](#)
- [Chapter 35. Uninstalling SF Oracle RAC from a cluster](#)



# Preparing to uninstall SF Oracle RAC from a cluster

This chapter includes the following topics:

- [About uninstalling SF Oracle RAC from a cluster](#)
- [Options for uninstalling SF Oracle RAC](#)
- [Preparing to uninstall SF Oracle RAC from a cluster](#)

## About uninstalling SF Oracle RAC from a cluster

You can uninstall SF Oracle RAC using the `uninstallsrac` program.


---

**Note:** After you uninstall SF Oracle RAC, you cannot access the Oracle database as Veritas Volume Manager and Veritas File System are uninstalled from the cluster. Make sure that you back up the Oracle database before you uninstall SF Oracle RAC.

---

[Figure 34-1](#) illustrates the steps that are required to uninstall SF Oracle RAC from a cluster.

Figure 34-1 SF Oracle RAC uninstallation



# Options for uninstalling SF Oracle RAC

Table 34-1 lists the available options for uninstalling SF Oracle RAC:

**Table 34-1** Options for uninstalling SF Oracle RAC

| Options                              | Description                                                                                                                                                                          |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SF Oracle RAC uninstallation program | Use the <code>uninstallsrac</code> program to uninstall SF Oracle RAC.<br>See <a href="#">“Preparing to uninstall SF Oracle RAC from a cluster”</a> on page 613.                     |
| Response file                        | Use a response file to automate or perform an unattended uninstallation of SF Oracle RAC.<br><br>See <a href="#">“Uninstalling SF Oracle RAC using a response file”</a> on page 631. |

# Preparing to uninstall SF Oracle RAC from a cluster

Perform the steps in the following procedure before you uninstall SF Oracle RAC from a cluster.

## To prepare to uninstall SF Oracle RAC from a cluster

- 1 Stop Oracle instances.  
See [“Stopping Oracle instances”](#) on page 614.
- 2 Back up the Oracle database.  
See [“Backing up the Oracle database”](#) on page 615.
- 3 Unlink the SF Oracle RAC libraries from Oracle RAC libraries (optional).  
See [“Unlinking the SF Oracle RAC libraries from Oracle RAC”](#) on page 615.
- 4 Uninstalling Oracle RAC (optional)  
See [“Uninstalling Oracle RAC \(optional\)”](#) on page 617.
- 5 Remove root disk encapsulation.  
See [“Removing root disk encapsulation”](#) on page 617.
- 6 Stop the applications that use CFS (outside of VCS control).  
See [“Stopping the applications that use CVM or CFS \(outside of VCS control\)”](#) on page 618.
- 7 Unmount CFS file systems (outside of VCS control).  
See [“Unmounting CFS file systems \(outside of VCS control\)”](#) on page 618.
- 8 Stop VCS.  
See [“Stopping VCS”](#) on page 619.

- 9
- Stop the applications that use VxFS (outside of VCS control).
- See “Stopping the applications that use VxVM or VxFS (outside of VCS control)” on page 620.
- 10
- Unmount VxFS file systems (outside of VCS control).
- See “Unmounting VxFS file systems (outside of VCS control)” on page 620.

## Stopping Oracle instances

You need to stop Oracle Clusterware and the Oracle instances on the cluster nodes where you want to uninstall SF Oracle RAC. Before you stop the Oracle instances, stop the applications that are dependent on the service groups that contain Oracle.

The procedure in this section provides instructions to stop the instances on a two-node cluster; the nodes are galaxy and nebula. Depending on the VCS configuration, the procedure to stop Oracle instances may vary.

### To stop Oracle instances

- 1
- Log in as the superuser on one of the nodes in the cluster.
- 2
- On each node, take the Oracle resources in the VCS configuration file (main.cf) offline.

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

For example:

```
# hagrps -offline ora1 -sys galaxy
```

```
# hagrps -offline ora1 -sys nebula
```

These commands stop the Oracle resources under VCS control.

If the database is not managed by VCS, stop the Oracle database as follows:

```
$ srvctl stop database -d db_name
```

- 3
- Verify that the state of the Oracle (if the database is managed by VCS) and CVM service groups are offline and online respectively.

```
# hagrps -state
Group      Attribute      System      Value
ora1       State         galaxy     |OFFLINE|
ora1       State         nebula     |OFFLINE|
cvm        State         galaxy     |ONLINE|
cvm        State         nebula     |ONLINE|
```

## Backing up the Oracle database

If you plan to retain the Oracle database, you must back up the Oracle database.

For instructions on backing up the Oracle database, see the Oracle documentation.

## Unlinking the SF Oracle RAC libraries from Oracle RAC

If you want to reuse the Oracle RAC binaries, unlink the SF Oracle RAC libraries from Oracle RAC.

---

**Note:** The Oracle RAC binaries that are located on Veritas File System will not be accessible after the uninstallation.

---

## To unlink the SF Oracle RAC libraries from Oracle RAC libraries

### 1 Unlink the Veritas VCSMM library.

Perform this step on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, perform the step on one node only.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
$ cd $CRS_HOME/lib
```

For Oracle RAC 11g Release 2:

```
$ cd $GRID_HOME/lib
```

```
$ ln -sf $CRS_HOME/lib/libskgxns.so libskgxns2.so
```

```
$ ln -sf $GRID_HOME/lib/libskgxns.so libskgxns2.so
```

### 2 For Oracle RAC 10g: Unlink the Veritas VCS IPC library.

Perform this step on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, perform the step on one node only.

```
$ cd $ORACLE_HOME/lib
```

```
$ cp libskgxpu.so libskgxpu10.so
```

### 3 Unlink the Veritas ODM library.

Perform this step on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, perform the step on one node only.

```
$ cd $ORACLE_HOME/lib
```

For Oracle RAC 10g:

```
$ ln -sf $ORACLE_HOME/lib/libodmd10.so libodmd10.so
```

For Oracle RAC 11g:

```
$ ln -sf $ORACLE_HOME/lib/libodmd11.so libodmd11.so
```



## Uninstalling Oracle RAC (optional)

You cannot use the Oracle database after you uninstall SF Oracle RAC. However, you can continue using the Oracle RAC software, provided the Oracle Clusterware and database binaries are not installed on Veritas file system. To continue using Oracle RAC, unlink the SF Oracle RAC libraries from Oracle RAC as described in the preceding topic.

For instructions on uninstalling Oracle RAC, see the Oracle documentation.

## Removing root disk encapsulation

Perform this step only if you plan to remove the VxVM and VVR packages.

If you have VxVM and VVR installed, you need to indicate to the installer whether or not you want to remove the VxVM packages from all nodes in the cluster. If you want to remove these packages, you need to ensure that the root disk is not encapsulated. The uninstallation fails if you choose to remove these packages while the root disk is encapsulated.

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

### To remove root disk encapsulation

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

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

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

---

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

---

- 2 Convert all the encapsulated volumes in the root disk to make them accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.

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

- 3 To check if the root disk is unencapsulated:

```
# df -v /
```

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

## Stopping the applications that use CVM or CFS (outside of VCS control)

You need to stop the applications that use CVM volumes or CFS mount points not controlled by VCS.

### To stop the applications that use CVM or CFS (outside of VCS control)

- 1 Using native application commands, stop the applications that use CVM or a CFS.
- 2 Verify that no processes use the CFS mount point:

```
# fuser -c mount_point
```

## Unmounting CFS file systems (outside of VCS control)

You need to unmount CFS file systems that are not under VCS control on all nodes.

### To unmount CFS file systems not under VCS control

- 1 Determine the file systems that need to be unmounted by checking the output of the mount command. The command lists all the mounted clustered file systems. Consult the main.cf file for identifying the files that are under VCS control.

```
# mount -v | grep vxfs | grep cluster
```

- 2 Unmount each file system that is not controlled by VCS:

```
# umount mount_point
```

## Stopping VCS

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

The process also stops replication as the replication service group is also taken offline.

### To stop VCS

- 1 Log in as the superuser on one of the cluster nodes.
- 2 Stop VCS on all nodes:

```
# hastop -all
```

- 3 Verify the output of the `gabconfig -a` command to ensure that VCS has been stopped.

In this command output, the VCS engine or high availability daemon (HAD) port h is not displayed. This output indicates that VCS has been stopped.

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

## Stopping the applications that use VxVM or VxFS (outside of VCS control)

You need to stop all applications that use VxVM volumes or VxFS mount points not under VCS control.

### To stop the applications that use VxVM or VxFS (outside of VCS control)

- 1 Using native application commands, stop the applications that use VxVM or VxFS.
- 2 Verify that no processes use the VxFS mount point:

```
# fuser -c mount_point
```

## Unmounting VxFS file systems (outside of VCS control)

You need to unmount VxFS file systems that are not under VCS control on all nodes.

---

**Note:** To avoid issues on rebooting, you must remove all entries of VxFS from the `/etc/vfstab` directory.

---

### To unmount VxFS file systems not under VCS control

- 1 Determine the file systems that need to be unmounted by checking the output of the `mount` command. The command lists all the mounted file systems.

```
# mount -v | grep vxfs
```

- 2 Unmount each file system that is not under VCS control:

```
# umount mount_point
```

# Uninstalling SF Oracle RAC from a cluster

This chapter includes the following topics:

- [Uninstalling SF Oracle RAC from a cluster](#)
- [Uninstalling SF Oracle RAC with the Veritas Web-based installer](#)
- [Uninstalling SF Oracle RAC using a response file](#)

## Uninstalling SF Oracle RAC from a cluster

You can remove the SF Oracle RAC packages from all nodes in the SF Oracle RAC cluster using the `uninstallsfrac` program. The `uninstallsfrac` program can be accessed from the product disc or from the `/opt/VRTS/install` directory.

Perform the steps in the following procedure to remove SF Oracle RAC from a cluster.

### To remove SF Oracle RAC from a cluster

- 1 Remove the CP server configuration using the removal script.  
See [“Removing the CP server configuration using the removal script”](#) on page 622.
- 2 Remove the SF Oracle RAC packages. You can remove the packages using the uninstallation program or using the response file.  
Using the uninstallation program:  
See [“Removing the SF Oracle RAC packages”](#) on page 626.  
Using the response file:  
See [“Uninstalling SF Oracle RAC using a response file”](#) on page 631.

- 3 Remove other configuration files (optional).  
See [“Removing other configuration files \(optional\)”](#) on page 627.
- 4 Remove the Storage Foundation for Databases (SFDB) repository after removing the product.  
See [“Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product”](#) on page 628.
- 5 Reboot the nodes.

```
# shutdown -g0 -y -i6
```

## Removing the CP server configuration using the removal script

This section describes how to remove the CP server configuration from a node or cluster hosting the CP server.

---

**Warning:** Ensure that no SF Oracle RAC cluster is using the CP server that will have its CP server configuration removed.

---

A configuration utility that is part of VRTScps package is used to remove the CP server configuration. When using the configuration utility, a configuration removal script is run and the following tasks are performed:

- All CP server configuration files are removed
- The VCS configuration for CP server is removed

After running the utility and script, you can then uninstall VCS from the node or cluster.

---

**Note:** The configuration script has to run only once per CP server (which can be on a single node or SFHA cluster), when removing the CP server configuration.

---

The configuration utility performs the following steps to remove the CP server configuration:

- Takes the the CP server service group (CPSSG) offline, if it is online
- Removes the CPSSG service group from the VCS configuration

The following procedure describes how to remove the CP server configuration.

### To remove the CP server configuration

- 1 To run the configuration removal script, enter the following command on the node where you want to remove the CP server configuration:

```
root@mycps1.symantecexample.com # /opt/VRTScps/bin/configure_cps.pl
```

- 2 The Veritas Coordination Point Server Configuration utility appears with an option menu.

```
VERITAS COORDINATION POINT SERVER CONFIGURATION UTILITY  
=====
```

Select one of the following:

[1] Configure Coordination Point Server on single node VCS system

[2] Configure Coordination Point Server on SFHA cluster

[3] Unconfigure Coordination Point Server

- 3 Select option 3 to unconfigure the Coordination Point Server.
- 4 A warning appears and prompts you to confirm the action to unconfigure the Coordination Point Server.

Enter "y" to proceed.

```
WARNING: Unconfiguring Coordination Point Server stops the  
vxcpserv process. VCS clusters using this server for  
coordination purpose will have one less coordination point.
```

```
Are you sure you want to bring down the cp server? (y/n)  
(Default:n) :y
```

- 5 After entering "y" to proceed, messages appear informing you of the progress in removing the CP server configuration.

When the CP server configuration has been unconfigured, a success message appears.

For an example of the messages from a single node VCS cluster:

```
A single node VCS cluster is currently configured.
Stopping the CP server ...
```

```
Removing the CP Server from VCS configuration..
```

```
Removing resource dependencies...
Deleting the resources configured under CPSSG service group...
Deleting the CPSSG service group...
```

```
Successfully unconfigured the Veritas Coordination Point Server.
```

For an example of the messages from a CP server on an SFHA cluster:

```
A multinode CP Server cluster is currently configured.
Stopping the CP server ...
```

```
Removing the CP Server from VCS configuration..
```

```
Removing resource dependencies...
Deleting the resources configured under CPSSG service group...
Deleting the CPSSG service group...
```

```
Successfully unconfigured the Veritas Coordination Point Server.
```

- 6 You are then prompted to delete the CP server database. Enter "y" to delete the database. For example:

```
Do you want to delete the CP Server database? (y/n) (Default:n) :
```

- 7 Enter "y" at the prompt to confirm the deletion of the CP server database.

```
Warning: This database won't be available if CP server
is reconfigured on the cluster. Are you sure you want to
proceed with the deletion of database? (y/n) (Default:n) :
```



- 8 Enter "y" to delete the CP server configuration file and log files. For example:

```
Do you want to delete the CP Server configuration file  
(/etc/vxcps.conf) and log files (in /var/VRTScps)? (y/n)  
(Default:n) : y
```

- 9 Run the `hagrp -state` command to ensure that the CPSSG service group has been removed from the node. For example:

```
root@mycps1.symantecexample.com # hagrp -state CPSSG  
  
VCS WARNING V-16-1-40131 Group CPSSG does not exist  
in the local cluster
```

## Manually remove the CP server fencing configuration

The following procedure describes how to manually remove the CP server fencing configuration from the CP server. This procedure is performed as part of the process to stop and remove server-based IO fencing.

---

**Note:** This procedure must be performed after the SF Oracle RAC cluster has been stopped, but before the SF Oracle RAC cluster software is uninstalled.

---

This procedure is required so that the CP server database can be reused in the future for configuring server-based fencing on the same SF Oracle RAC cluster(s).

Perform the steps in the following procedure to manually remove the CP server fencing configuration.

---

**Note:** The `cpsadm` command is used in the following procedure. For detailed information about the `cpsadm` command, see the *Veritas Storage Foundation for Oracle RAC Administrator's Guide*.

---

### To manually remove the CP server fencing configuration

- 1 Unregister all SF Oracle RAC cluster nodes from all CP servers using the following command:

```
# cpsadm -s cp_server -a unreg_node -u uuid -n nodeid
```

- 2 Remove the SF Oracle RAC cluster from all CP servers using the following command:

```
# cpsadm -s cp_server -a rm_clus -u uuid
```

- 3 Remove all the SF Oracle RAC cluster users communicating to CP servers from all the CP servers using the following command:

```
# cpsadm -s cp_server -a rm_user -e user_name -g domain_type
```

- 4 Proceed to uninstall the SF Oracle RAC cluster software.

## Removing the SF Oracle RAC packages

The installer offers the option to remove the Veritas Volume Manager packages. The `uninstallsrac` program can remove these packages only if the root disk is not under VxVM control and there are no open volumes.

The installer performs the following tasks:

- Removes the SF Oracle RAC packages.
- Removes the language packages, if installed.

---

**Note:** The following directories remain after uninstallation: `/opt/VRTS`, `/opt/VRTSperl`, `/etc/VRTSvcs`, `/var/VRTSvcs`, `/var/VRTSat_lhc`, `/var/VRTSat`. They contain logs and configuration information for future reference. You may or may not remove them.

---

### To remove the SF Oracle RAC packages

- 1 Log in as the superuser on any node in the cluster.
- 2 Navigate to the directory that contains the `uninstallsrac` program:

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

### 3 Start the `uninstallsfrac` program:

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

The program displays the directory where the logs are created and the copyright message.

### 4 If you have VxVM installed, indicate whether or not you want to remove the VxVM packages from all nodes in the cluster. Enter **y** only if the root disk is outside of VxVM control.

The `uninstallsfrac` program performs the following tasks:

- Checks the operating system on each node
- Verifies the system-to-system communication
- Verifies the licenses
- Checks for the SF Oracle RAC packages installed on the nodes. This process involves identifying system uninstallation requirements and dependencies between packages to determine the safety and order of uninstalling packages.

### 5 Confirm to uninstall SF Oracle RAC.

```
All SF Oracle RAC processes that are currently running must be stopped.  
Do you want to stop SF Oracle RAC processes now? [y,n,q,?] (y) y
```

---

**Note:** If you have not already unmounted the VxFS file systems, the installer will display a message asking that the file systems be unmounted. Make sure that you unmount the file systems before you proceed.

---

The program performs the following tasks:

- Stops the agents and performs verifications on each node to proceed with uninstallation
- Stops the SF Oracle RAC processes and uninstalls the SF Oracle RAC packages
- Displays the location of the uninstallation summary and log files for reference.

## Removing other configuration files (optional)

You can remove the Veritas configuration files and the packages that are left after running the `uninstallsfrac` program.

**To remove residual Veritas configuration files (optional)**

- 1 List all VRTS packages that can be removed.

```
# pkginfo -l |grep -i vrts
```

- 2 Run the `pkgrm pkgname` command to remove the remaining VRTS packages.
- 3 Move the residual Veritas configuration files to the `vrts.bkp` directory:

```
# cd /var
# mkdir vrts.bkp
# mv *VRTS* vrts.bkp
# mv vx vrts.bkp
# cd /var/opt
# mkdir vrts.bkp
# mv *VRTS* vrts.bkp
# cd /opt
# mkdir vrts.bkp
# mv *VRTS* vrts.bkp
# cd /etc
# mkdir vrts.bkp
# mv vx *llt* *fen* *gab* *vcs* vcsmmtab vrts.bkp
```

You can remove the `vrts.bkp` directories at a later time.

## Removing the Storage Foundation for Databases (SFDB) repository after removing the product

After removing the product, you can remove the SFDB repository file and any backups.

Removing the SFDB repository file will disable the SFDB tools.

### To remove the SFDB repository

- 1 Change directories to the location of the local lookup information for the Oracle SID.

For example:

```
# cd /var/vx/vxdba/$ORACLE_SID
```

- 2 Identify the SFDB repository file and any associated links:

For example:

```
# ls -al
```

```
lrwxrwxrwx  1 oracle  oinstall      26 Jul 21 13:58 .sfdb_rept -> \
/ora_data1/TEST/.sfdb_rept
```

```
# cd /ora_data1/TEST
```

Follow the symlink of .sfdb\_rept.

- 3 Remove the repository directory containing the repository file and all backups.

For example:

```
# rm -rf .sfdb_rept
```

- 4 Remove the local lookup directory for the Oracle SID:

```
# cd /var/vx/vxdba
```

```
# rm -rf $ORACLE_SID
```

This completes the removal of the SFDB repository.

## Uninstalling SF Oracle RAC with the Veritas Web-based installer

This section describes how to uninstall with the Veritas Web-based installer.

### To uninstall SF Oracle RAC

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 3 Start the Web-based installer.  
See [“Starting the Veritas Web-based installer”](#) on page 97.
- 4 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
- 5 Select **Veritas Storage Foundation for Oracle RAC** from the Product drop-down list, and click **Next**.
- 6 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Validate**.
- 7 After the validation completes successfully, click **Next** to uninstall SF Oracle RAC on the selected system.
- 8 If there are any processes running on the target system, the installer stops the processes. Click **Next**.
- 9 After the installer stops the processes, the installer removes the products from the specified system.  
Click **Next**.
- 10 After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.
- 11 Click **Finish**.

The Web-based installer prompts you for another task.

# Uninstalling SF Oracle RAC using a response file

Perform the steps in the following procedure to uninstall SF Oracle RAC using a response file.

## To uninstall SF Oracle RAC using a response file

- 1 Make sure that you have completed the pre-uninstallation tasks.
- 2 Create a response file using one of the available options.  
For information on various options available for creating a response file:  
See [“About response files”](#) on page 425.

---

**Note:** You must replace the host names in the response file with that of the systems from which you want to uninstall SF Oracle RAC.

---

For a sample response file:

See [“Sample response file for uninstalling SF Oracle RAC”](#) on page 632.

- 3 Navigate to the directory containing the SF Oracle RAC uninstallation program:

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

- 4 Start the uninstallation:

```
# ./uninstallsfrac -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the full path name of the response file.

- 5 Remove other Veritas configuration files and packages that may be present.
- 6 Reboot the nodes:

```
# shutdown -g0 -y -i6
```

## Response file variables to uninstall SF Oracle RAC

[Table 35-1](#) lists the response file variables that you can define to uninstall SF Oracle RAC.

**Table 35-1** Response file variables specific to uninstalling SF Oracle RAC

| Variable            | List or Scalar | Description                                                                                                                                                                                                                                        |
|---------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{uninstall} | Scalar         | Uninstalls SF Oracle RAC packages.<br>(Required)                                                                                                                                                                                                   |
| CFG{systems}        | List           | List of systems on which the product is to be uninstalled.<br>(Required)                                                                                                                                                                           |
| CFG{prod}           | Scalar         | Defines the product to be uninstalled.<br>(Required)                                                                                                                                                                                               |
| CFG{opt}{keyfile}   | Scalar         | Defines the location of an ssh keyfile that is used to communicate with all remote systems.<br>(Optional)                                                                                                                                          |
| CFG{opt}{rsh}       | Scalar         | Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems.<br>(Optional)                                                                                                                              |
| CFG{opt}{logpath}   | Scalar         | Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.<br><b>Note:</b> The installer copies the response files and summary files also to the specified <i>logpath</i> location.<br>(Optional) |

## Sample response file for uninstalling SF Oracle RAC

The following sample response file uninstalls SF Oracle RAC from nodes, galaxy and nebula.

```
our %CFG;  
  
$CFG{opt}{uninstall}=1;
```



```
$CFG{prod}="SFRAC51";  
$CFG{systems}=[ qw(galaxy nebula) ];  
  
1;
```



## Installation reference

- [Appendix A. Installation scripts](#)
- [Appendix B. SF Oracle RAC installation and configuration sample values](#)
- [Appendix C. Sample configuration files](#)
- [Appendix D. Setting up inter-system communication](#)
- [Appendix E. Automatic Storage Management](#)
- [Appendix F. Creating a test database](#)
- [Appendix G. High availability agent information](#)
- [Appendix H. SF Oracle RAC deployment scenarios](#)



# Installation scripts

This appendix includes the following topics:

- [About installation scripts](#)
- [Installation script options](#)

## About installation scripts

Veritas Storage Foundation and High Availability Solutions 5.1 Service Pack 1 provides several installation scripts.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from an electronic download site, which does not include the installer, use the appropriate product installation script.

The following product installation scripts are available:

|                                                           |                            |
|-----------------------------------------------------------|----------------------------|
| Veritas Cluster Server (VCS)                              | <code>installvcs</code>    |
| Veritas Storage Foundation (SF)                           | <code>installsf</code>     |
| Veritas Storage Foundation and High Availability (SFHA)   | <code>installsfha</code>   |
| Veritas Storage Foundation Cluster File System (SFCFS)    | <code>installsfdfs</code>  |
| Veritas Storage Foundation for Oracle RAC (SF Oracle RAC) | <code>installsffrac</code> |
| Symantec Product Authentication Service (AT)              | <code>installat</code>     |
| Veritas Dynamic Multi-pathing                             | <code>installdmp</code>    |

Symantec VirtualStore

installsvs

To use the installation script, enter the script name at the prompt. For example, to install Veritas Storage Foundation, type `./installsf` at the prompt.

# Installation script options

[Table A-1](#) shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas Storage Foundation product scripts, except where otherwise noted.

See [“About installation scripts”](#) on page 637.

**Table A-1** Available command line options

| Command Line Option       | Function                                                                                                                                                                                                                                          |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>system1 system2...</i> | Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.                                                                               |
| -addnode                  | Adds a node to a high availability cluster.                                                                                                                                                                                                       |
| -allpkgs                  | Displays all packages and patches required for the specified product. The packages and patches are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network. |
| -comcleanup               | The -comcleanup option removes the ssh or rsh configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of ssh or rsh are abruptly terminated.                    |
| -configure                | Configures the product after installation.                                                                                                                                                                                                        |

Table A-1 Available command line options (*continued*)

| Command Line Option | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -copyinstallscripts | <p>Use this option when you manually install products and want to use the intallation scripts that are stored on the system to perform product configuration, uninstallation, and licensing tasks without the product media.</p> <p>Use this option to copy the installation scripts to an alternate rootpath when you use it with the <code>-rootpath</code> option.</p> <p>The following examples demonstrate the usage for this option:</p> <ul style="list-style-type: none"> <li> <pre>■ ./installer -copyinstallscripts</pre> <p>Copies the installation and uninstallation scripts for all products in the release to <code>/opt/VRTS/install</code>. It also copies the installation Perl libraries to <code>/opt/VRTSperl/lib/site_perl/release_name</code>.</p> </li> <li> <pre>■ ./installproduct_name -copyinstallscripts</pre> <p>Copies the installation and uninstallation scripts for the specified product and any subset products for the product to <code>/opt/VRTS/install</code>. It also copies the installation Perl libraries to <code>/opt/VRTSperl/lib/site_perl/release_name</code>.</p> </li> <li> <pre>■ ./installer -copyinstallscripts -rootpath alt_root_path</pre> <p>The path <i>alt_root_path</i> can be a directory like <code>/rdisk2</code>. In that case, this command copies installation and uninstallation scripts for all the products in the release to <code>/rdisk2/opt/VRTS/install</code>. CPI perl libraries are copied to <code>/rdisk2/opt/VRTSperl/lib/site_perl/release_name</code>, where the <i>release_name</i> is a string that starts with UXRT and includes the release version with no punctuation.</p> </li> </ul> |
| -fencing            | Configures I/O fencing in a running cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Table A-1** Available command line options (*continued*)

| Command Line Option                             | Function                                                                                                                                                                                                                             |
|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-hostfile <i>full_path_to_file</i></code> | Specifies the location of a file that contains a list of hostnames on which to install.                                                                                                                                              |
| <code>-ignorepatchreqs</code>                   | The <code>-ignorepatchreqs</code> option is used to allow installation or upgrading even if the prerequisite packages or patches are missed on the system.                                                                           |
| <code>-install</code>                           | The <code>-install</code> option is used to install products on systems.                                                                                                                                                             |
| <code>-installallpkgs</code>                    | Specifies that all packages are installed.                                                                                                                                                                                           |
| <code>-installminpkgs</code>                    | Specifies that the minimum package set is installed.                                                                                                                                                                                 |
| <code>-installrecpkgs</code>                    | Specifies that the required package set is installed.                                                                                                                                                                                |
| <code>-jumpstart <i>dir_path</i></code>         | Produces a sample finish file for Solaris JumpStart installation. The <i>dir_path</i> indicates the path to the directory in which to create the finish file.                                                                        |
| <code>-keyfile <i>ssh_key_file</i></code>       | Specifies a key file for secure shell (SSH) installs. This option passes <code>-i <i>ssh_key_file</i></code> to every SSH invocation.                                                                                                |
| <code>-license</code>                           | Registers or updates product licenses on the specified systems.                                                                                                                                                                      |
| <code>-listpatches</code>                       | The <code>-listpatches</code> option displays product patches in correct installation order.                                                                                                                                         |
| <code>-logpath <i>log_path</i></code>           | Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.                                                                         |
| <code>-makeresponsefile</code>                  | The <code>-makeresponsefile</code> generates a response file without doing an actual installation. Text displaying install, uninstall, start, and stop actions are simulations. These actions are not being performed on the system. |



**Table A-1** Available command line options (*continued*)

| Command Line Option          | Function                                                                                                                                                                                                                                                                                                                     |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -minpkgs                     | Displays the minimal packages and patches required for the specified product. The packages and patches are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option. |
| -patchpath <i>patch_path</i> | Designates the path of a directory that contains all patches to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.                                                                                                                                    |
| -pkginfo                     | Displays a list of packages and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the <code>-pkginfo</code> option with the <code>installvcs</code> script to display VCS packages.                                            |
| -pkgpath <i>package_path</i> | Designates the path of a directory that contains all packages to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.                                                                                                                                   |
| -pkgset                      | Discovers and displays the package group (minimum, recommended, all) and packages that are installed on the specified systems.                                                                                                                                                                                               |
| -pkgtable                    | Displays product's packages in correct installation order by group.                                                                                                                                                                                                                                                          |
| -postcheck                   | Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.                                                                                                                                                                      |
| -precheck                    | Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.                                                                                                                                                               |

**Table A-1** Available command line options (*continued*)

| Command Line Option                             | Function                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-recpkgs</code>                           | Displays the recommended packages and patches required for the specified product. The packages and patches are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.                                                        |
| <code>-redirect</code>                          | Displays progress details without showing the progress bar.                                                                                                                                                                                                                                                                                                                             |
| <code>-requirements</code>                      | The <code>-requirements</code> option displays required OS version, required patches, file system space, and other system requirements in order to install the product.                                                                                                                                                                                                                 |
| <code>-responsefile <i>response_file</i></code> | Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.                                              |
| <code>-rootpath <i>root_path</i></code>         | Specifies an alternative root directory on which to install packages.<br><br>On Solaris operating systems, <code>-rootpath</code> passes <code>-R path</code> to <code>pkgadd</code> command.                                                                                                                                                                                           |
| <code>-rsh</code>                               | Specify this option when you want to use <code>rsh</code> and <code>rcp</code> for communication between systems instead of the default <code>ssh</code> and <code>scp</code> .                                                                                                                                                                                                         |
| <code>-security</code>                          | Enable or disable Symantec Product Authentication Service in a VCS cluster that is running.<br><br>You can specify this option with the <code>installvcs</code> , <code>installsfha</code> or <code>installsfdfs</code> scripts.<br><br>For more information about Symantec Product Authentication Service in a VCS cluster, see the <i>Veritas Cluster Server Installation Guide</i> . |

**Table A-1** Available command line options (*continued*)

| Command Line Option      | Function                                                                                                                                                                                                                                                                     |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -serial                  | Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.                                            |
| -start                   | Starts the daemons and processes for the specified product.                                                                                                                                                                                                                  |
| -stop                    | Stops the daemons and processes for the specified product.                                                                                                                                                                                                                   |
| -tmppath <i>tmp_path</i> | Specifies a directory other than /var/tmp as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.                                                 |
| -uninstall               | The -uninstall option is used to uninstall products from systems.                                                                                                                                                                                                            |
| -upgrade                 | Specifies that an existing version of the product exists and you plan to upgrade it.                                                                                                                                                                                         |
| -upgrade_kernelpkgs      | The -upgrade_kernelpkgs option is used to perform rolling upgrade Phase-I. In the phase, the product kernel packages get upgraded to the latest version                                                                                                                      |
| -upgrade_nonkernelpkgs   | The -upgrade_nonkernelpkgs option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent packages upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version."                                        |
| -version                 | Checks and reports the installed products and their versions. Identifies the installed and missing packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable. |



# SF Oracle RAC installation and configuration sample values

This appendix includes the following topics:

- [About the installation and configuration worksheets](#)
- [SF Oracle RAC worksheet](#)
- [Oracle RAC worksheet](#)
- [Replicated cluster using VVR worksheet](#)
- [Replicated cluster using SRDF worksheet](#)
- [Required installation information for Oracle Clusterware/Grid Infrastructure](#)
- [Required installation information for Oracle database](#)

## About the installation and configuration worksheets

The installation programs prompt you for information during the installation and configuration of SF Oracle RAC and Oracle RAC. The installation and configuration worksheets provide sample values that you can use as examples of the information required when you run the Veritas installation programs or the Oracle Universal Installer.

Symantec recommends using the worksheets to record values for your systems before you begin the installation and configuration process.

# SF Oracle RAC worksheet

This section provides worksheets for installing and configuring SF Oracle RAC, its component products, and features.

Table B-1 contains the sample values that may be used when you install and configure SF Oracle RAC. Enter the SF Oracle RAC values for your systems in the following table:

Table B-1                      SF Oracle RAC worksheet

| Installation information             | Sample value                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Assigned value |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| Number of nodes in the cluster       | 2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                |
| Host names for Primary cluster       | galaxy and nebula<br><b>Note:</b> Do not use the underscore character in host names. Host names that use the underscore character are not compliant with RFC standards and cause issues.                                                                                                                                                                                                                                                                                             |                |
| Host names for added or removed node | saturn<br><b>Note:</b> Do not use the underscore character in host names. Host names that use the underscore character are not compliant with RFC standards and cause issues.                                                                                                                                                                                                                                                                                                        |                |
| License keys                         | License keys can be one of the following types:<br><ul style="list-style-type: none"><li>Valid license keys for each system in the cluster</li><li>Valid site license key</li><li>Valid demo license key</li></ul> If you want to configure Veritas Volume Replicator to enable disaster recovery, you must enter appropriate license keys.<br><b>Note:</b> You can choose between SF Oracle RAC and SF Oracle RAC Disaster Recovery and High Availability options for license keys. |                |

**Table B-1** SF Oracle RAC worksheet (*continued*)

| Installation information                                                  | Sample value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Assigned value |
|---------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| SF Oracle RAC packages to be installed on the cluster                     | <p>Select one of the following options:</p> <ul style="list-style-type: none"><li>■ Install minimal required Veritas Storage Foundation for Oracle RAC packages</li><li>■ Install recommended Veritas Storage Foundation for Oracle RAC packages</li><li>■ Install all Veritas Storage Foundation for Oracle RAC packages</li></ul> <p>Install only the required packages if you do not want to configure any optional components or features.</p> <p>Default option is to install the recommended packages.</p> |                |
| Primary cluster name                                                      | rac_cluster101                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                |
| Primary cluster ID number                                                 | 101                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                |
| Private network links                                                     | bge1, bge2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                |
| Cluster Manager NIC (Primary NIC)                                         | bge0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                |
| Cluster Manager IP                                                        | 10.10.12.1, 10.10.12.2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                |
| Netmask for the virtual IP address                                        | 255.255.240.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                |
| VCS user name (not required if you configure your cluster in secure mode) | <p>VCS usernames must not exceed 1024 characters.</p> <p>Example: smith</p>                                                                                                                                                                                                                                                                                                                                                                                                                                      |                |
| VCS user password                                                         | <p>VCS passwords must not exceed 512 characters.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                |
| VCS user privileges                                                       | <p>Users have three levels of privileges: A=Administrator, O=Operator, or G=Guest.</p> <p>Example: A</p>                                                                                                                                                                                                                                                                                                                                                                                                         |                |

Table B-1 SF Oracle RAC worksheet (continued)

| Installation information                               | Sample value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Assigned value |
|--------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| Domain-based address of SMTP server                    | smtp.symantecexample.com                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                |
| Email address of SMTP notification recipients          | admin@symantecexample.com                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                |
| Minimum severity of events for SMTP email notification | <p>Events have four levels of severity:</p> <ul style="list-style-type: none"><li>■ I=Information</li><li>■ W=Warning</li><li>■ E=Error</li><li>■ S=SevereError</li></ul> <p>Example: I</p> <p>The severity levels are defined as follows:</p> <ul style="list-style-type: none"><li>■ Information - Important events that exhibit normal behavior</li><li>■ Warning - Deviation from normal behavior</li><li>■ Error - A fault</li><li>■ Severe Error -Critical error that can lead to data loss or corruption</li></ul> |                |
| SNMP trap daemon port number the console               | 162                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                |
| System name for the SNMP console                       | system2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                |



**Table B-1** SF Oracle RAC worksheet (*continued*)

| Installation information                                                                                             | Sample value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Assigned value |
|----------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| Minimum severity of events for SNMP trap notification                                                                | <p>Events have four levels of severity:</p> <ul style="list-style-type: none"> <li>■ I=Information</li> <li>■ W=Warning</li> <li>■ E=Error</li> <li>■ S=SevereError</li> </ul> <p>Example: I</p> <p>The severity levels are defined as follows:</p> <ul style="list-style-type: none"> <li>■ Information - Important events that exhibit normal behavior</li> <li>■ Warning - Deviation from normal behavior</li> <li>■ Error - A fault</li> <li>■ Severe Error -Critical error that can lead to data loss or corruption</li> </ul> |                |
| Vxfen disks<br><br>These values are required if SCSI-3 disks are used as coordination points for your configuration. | c1t1d0s2 , c2t1d0s2, c3t1d0s2                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                |
| Vxfen disk group                                                                                                     | vxfencoordg                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                |

## Veritas Cluster Server component information

[Table B-2](#) displays the information that is required to configure the Veritas Cluster Server component.

**Table B-2** Veritas Cluster Server component information

| Information         | Example                                                                                                                                                                                                       | Assigned values |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Name of the cluster | <p>The name must begin with a letter of the alphabet (a-z, A-Z) and contain only the characters a through z, A through Z, and 1 through 0, hyphen (-), and underscore (_).</p> <p>Example: rac_cluster101</p> |                 |

Table B-2 Veritas Cluster Server component information (continued)

| Information                                                         | Example                                                                                                                                                                                                                                                                                                                                                                                                                                                | Assigned values |
|---------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Unique ID number for the cluster                                    | Number in the range of 0-65535. Within the site that contains the cluster, each cluster must have a unique ID.<br><br>Example: 101                                                                                                                                                                                                                                                                                                                     |                 |
| Device names of the NICs used by the private networks among systems | You can choose a network interface card that is not part of any aggregated interface, or you can choose an aggregated interface.<br><br>The interface names that are associated with each NIC for each network link must be the same on all nodes.<br><br>For example:<br><ul style="list-style-type: none"><li>■ bge1</li><li>■ bge2</li></ul><br>Do not use the network interface card that is used for the public network, which is typically bge0. |                 |

I/O fencing information

Table B-3 displays the information that is required to configure I/O fencing.

Table B-3 I/O fencing information

| Information                                                   | Sample values                                                                                                                                                               | Assigned values |
|---------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| The name of three disks that form the coordinator disk group. | The following are examples of disk names:<br><ul style="list-style-type: none"><li>■ /dev/rdsk/c1t1d0s2</li><li>■ /dev/rdsk/c2t1d0s2</li><li>■ /dev/rdsk/c3t1d0s2</li></ul> |                 |

Table B-3 I/O fencing information (continued)

| Information                                                           | Sample values                                                                                                                                                           | Assigned values |
|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| The names for each disk in the coordinator disk group (if using DMP). | The following are examples: <ul style="list-style-type: none"> <li>■ /dev/vx/rdmp/c1t1d0s2</li> <li>■ /dev/vx/rdmp/c2t1d0s2</li> <li>■ /dev/vx/rdmp/c3t1d0s2</li> </ul> |                 |

SF Oracle RAC secure mode cluster information

Table B-4 displays the information that is required to configure SF Oracle RAC clusters in secure mode.

**Note:** Configuring SF Oracle RAC in secure mode is optional.

Table B-4 SF Oracle RAC secure mode cluster information

| Information                                                                 | Examples                                                                                                                                                                                                                                                             | Assigned values |
|-----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Which mode do you want to choose to configure Authentication Service?       | Select one of the three modes: <ul style="list-style-type: none"> <li>■ Configure security completely automatically</li> <li>■ Provide AB credentials using BLOBs</li> <li>■ Provide AB credentials without using BLOBs</li> </ul> Default option is automatic mode. |                 |
| Host name of the Symantec Product Authentication Service Root Broker System | Example: venus                                                                                                                                                                                                                                                       |                 |

SF Oracle RAC add user information

Table B-5 displays the information that is required to add VCS users. If you configure SF Oracle RAC cluster in secure mode, you need to add VCS users.

**Note:** Adding VCS users is optional.

Table B-5 SF Oracle RAC add user information

| Information    | Examples                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Assigned values |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| User name      | smith                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                 |
| User password  | *****                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                 |
| User privilege | <p>Users have three levels of privileges:</p> <ul style="list-style-type: none"><li>■ A=Administrator</li><li>■ O=Operator</li><li>■ G=Guest</li></ul> <p>Example: A</p> <p>VCS privilege levels include:</p> <ul style="list-style-type: none"><li>■ Administrators— Can perform all operations, including configuration options on the cluster, service groups, systems, resources, and users.</li><li>■ Operators—Can perform specific operations on a cluster or a service group.</li><li>■ Guests—Can view specified objects.</li></ul> |                 |

Global cluster information

Table B-6 displays the information that is required to configure global clusters.

**Note:** Global clusters are an optional feature that requires a license.

Table B-6 Global cluster information

| Information            | Example                                                                                            | Assigned values |
|------------------------|----------------------------------------------------------------------------------------------------|-----------------|
| Name of the public NIC | <p>You must specify appropriate values for the NIC when you are prompted.</p> <p>Example: bge0</p> |                 |

**Table B-6** Global cluster information (*continued*)

| Information                        | Example                                                                                                          | Assigned values |
|------------------------------------|------------------------------------------------------------------------------------------------------------------|-----------------|
| Virtual IP address of the NIC      | You must specify appropriate values for the virtual IP address when you are prompted.<br><br>Example: 10.10.12.1 |                 |
| Netmask for the virtual IP address | You must specify appropriate values for the netmask when you are prompted.<br><br>Example: 255.255.255.0         |                 |

## Oracle RAC worksheet

This section provides a worksheet with sample values for installing and configuring Oracle RAC.

[Table B-7](#) displays examples of information required for Oracle RAC.

**Table B-7** Required information for Oracle RAC

| Information                                      | Sample value                                                                                | Assigned value |
|--------------------------------------------------|---------------------------------------------------------------------------------------------|----------------|
| Oracle user name                                 | oracle                                                                                      |                |
| Grid user name<br>(For Oracle RAC 11g Release 2) | grid                                                                                        |                |
| Oracle user ID                                   | 1000                                                                                        |                |
| Oracle group name                                | For inventory group as primary group: oinstall<br><br>For dba group as secondary group: dba |                |
| Oracle group ID                                  | For inventory group as primary group: 1000<br><br>For dba group as secondary group: 1001    |                |
| Oracle user home directory                       | /home/oracle                                                                                |                |
| Grid user home directory                         | /home/grid                                                                                  |                |

**Table B-7** Required information for Oracle RAC *(continued)*

| Information                           | Sample value                                                                                                                                                                                                                                                                                                                                                                                        | Assigned value |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| DISPLAY environment variable value    | 10.20.12.150:0.0<br><b>Note:</b> You can provide the host name or IP address of the system that you want to use to display the Oracle installer.                                                                                                                                                                                                                                                    |                |
| VCS resource name for PrivNIC         | ora_priv                                                                                                                                                                                                                                                                                                                                                                                            |                |
| Private IP addresses for each node    | <ul style="list-style-type: none"><li>■ On galaxy: 192.168.12.1</li><li>■ On nebula: 192.168.12.2</li></ul> <b>Note:</b> The IP addresses must not contain leading zeroes in any of the octets that comprise the address for example 0X.X.X.X or X0X.X.X, or X.X.0X.X, or X.X.X.0X. Make sure that the private IP addresses have a format as displayed in the above examples for galaxy and nebula. |                |
| Private hostnames (set in /etc/hosts) | galaxy-priv, nebula-priv                                                                                                                                                                                                                                                                                                                                                                            |                |

**Table B-7** Required information for Oracle RAC (*continued*)

| Information                          | Sample value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Assigned value |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| NIC address for network on each node | <p>You have to choose an LLT device as a device for the Oracle Clusterware heartbeat. The interfaces specified should be exactly the same in name and total number as those which have been used for LLT configuration.</p> <p>For example, if the LLT devices are bge1, bge2:</p> <ul style="list-style-type: none"> <li>■ On galaxy: bge1, bge2</li> <li>■ On nebula: bge1, bge2</li> </ul> <p>Then the PrivNIC or MultiPrivNIC device names will be as follows:</p> <ul style="list-style-type: none"> <li>■ Device@galaxy= { bge1 = 0, bge2 = 1 }</li> <li>■ Device@nebula= { bge1 = 0, bge2 = 1 }</li> </ul> <p>If aggregated device names are configured under LLT, then the aggregated names must be used in PrivNIC or MultiPrivNIC agent.</p> <p><b>Note:</b> If you configured aggregated interfaces for LLT, then you must set the Device attribute value to use the same aggregated interface names that you configured for LLT.</p> <p>For example, if LLT device names are:</p> <ul style="list-style-type: none"> <li>■ On galaxy: aggr1</li> <li>■ On nebula: aggr1</li> </ul> <p>Then the Device Attribute for the MultiPrivNIC agent would be as follows:</p> <ul style="list-style-type: none"> <li>■ Device@galaxy = { aggr1 = 0 }</li> <li>■ Device@nebula = { aggr1 = 0 }</li> </ul> |                |
| VCS resource name for MultiPrivNIC   | multi_priv                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                |

**Table B-7** Required information for Oracle RAC *(continued)*

| Information                                              | Sample value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Assigned value |
|----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| Private IP addresses for each node<br><br>(MultiPrivNIC) | <p>On galaxy:</p> <ul style="list-style-type: none"><li>■ 192.168.12.1 (Oracle Clusterware)</li><li>■ 192.168.2.1 (UDP to use for the CLUSTER_INTERCONNECT parameter)</li><li>■ 192.168.3.1 (Second UDP to use for the CLUSTER_INTERCONNECT parameter)</li></ul> <p>On nebula:</p> <ul style="list-style-type: none"><li>■ 192.168.12.2 (Oracle Clusterware)</li><li>■ 192.168.2.2 (UDP to use for the CLUSTER_INTERCONNECT parameter)</li><li>■ 192.168.3.2 (Second UDP to use for the CLUSTER_INTERCONNECT parameter)</li></ul> <p><b>Note:</b> The IP addresses must not contain leading zeroes in any of the octets that comprise the address for example 0X.X.X.X or X.0X.X.X, or X.X.0X.X, or X.X.X.0X. Make sure that the private IP addresses have a format as displayed in the above examples for galaxy and nebula.</p> <p>Symantec recommends that all Oracle Clusterware and UDP cache-fusion links be LLT links.</p> |                |
| Private hostnames (set in /etc/hosts)                    | <p>On galaxy:</p> <ul style="list-style-type: none"><li>■ galaxy-priv</li><li>■ galaxy-priv1</li></ul> <p>On nebula:</p> <ul style="list-style-type: none"><li>■ nebula-priv</li><li>■ nebula-priv1</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                |
| Netmask for cluster                                      | 255.255.255.0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                |



**Table B-7** Required information for Oracle RAC (*continued*)

| Information                                                                   | Sample value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Assigned value |
|-------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| Disk for each node that contains the Oracle Clusterware and database binaries | Disk_1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                |
| VxVM disk group                                                               | <ul style="list-style-type: none"> <li>■ On galaxy: bindg_galaxy</li> <li>■ On nebula: bindg_nebula</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                     |                |
| Volume for each local disk group                                              | <ul style="list-style-type: none"> <li>■ On galaxy: crsbinvol</li> <li>■ On nebula: orabinvol</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                           |                |
| File system on each local volume                                              | <ul style="list-style-type: none"> <li>■ On galaxy:               <br/><br/>               /dev/vx/rdisk/\               <br/>bindg_galaxy/crsbinvol               <br/><br/>               /dev/vx/rdisk/\               <br/>bindg_galaxy/orabinvol             </li> <li>■ On nebula:               <br/><br/>               /dev/vx/rdisk/\               <br/>bindg_nebula/crsbinvol               <br/><br/>               /dev/vx/rdisk/\               <br/>bindg_nebula/orabinvol             </li> </ul> |                |
| Oracle base directory (ORACLE_BASE)                                           | <ul style="list-style-type: none"> <li>■ Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:               <br/>/u01/app/oracle             </li> <li>■ Oracle RAC 11g Release 2:               <br/>/u01/app/gridbase             </li> </ul>                                                                                                                                                                                                                                                                      |                |
| Oracle home directory (ORACLE_HOME)                                           | /u01/app/oracle/product/ <i>release</i> /dbhome_1<br>where <i>release</i> is 10.2.0, 11.1.0, or 11.2.0 depending on the Oracle RAC version                                                                                                                                                                                                                                                                                                                                                                         |                |
| Oracle Clusterware home directory (CRS_HOME)                                  | /u01/app/oracle/product/10.2.0/crshome                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                |
| Grid home directory (GRID_HOME)                                               | /u01/app/grid/product/11.2.0/gridhome                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                |

**Table B-7** Required information for Oracle RAC (*continued*)

| Information                               | Sample value                                                                                                                                                                                                                                                                                                                                                                     | Assigned value |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| File system mount point (VxFS)            | <ul style="list-style-type: none"> <li>■ Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:<br/>/u01/app/oracle (for crsbinvol and orabinvol)</li> <li>■ Oracle RAC 11g Release 2:<br/>/u01/app/grid (crsbinvol)<br/>/u01/app/oracle (orabinvol)</li> </ul>                                                                                                                      |                |
| File system mount point (CFS)             | <ul style="list-style-type: none"> <li>■ Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:<br/>/u01/app/oracle/product/10.2.0/crshome (for crsbinvol)<br/>/u01/app/oracle/product/10.2.0/dbhome_1 (for orabinvol)</li> <li>■ Oracle RAC 11g Release 2:<br/>/u01/app/grid/product/11.2.0/gridhome (crsbinvol)<br/>/u01/app/oracle/product/11.2.0/dbhome_1 (orabinvol)</li> </ul> |                |
| Disk group for Oracle Database            | oradatadg (Shared)                                                                                                                                                                                                                                                                                                                                                               |                |
| Volume for Oracle Database                | oradatavol                                                                                                                                                                                                                                                                                                                                                                       |                |
| Volume options for Oracle Database        | layout=striped-mirror                                                                                                                                                                                                                                                                                                                                                            |                |
| Mount point for Oracle Database           | /oradata (Shared)                                                                                                                                                                                                                                                                                                                                                                |                |
| VCS resource names for Oracle Database    | oradata_voldg, oradata_mnt                                                                                                                                                                                                                                                                                                                                                       |                |
| Global Database Name                      | db.symantecexample.com (database name.domain)                                                                                                                                                                                                                                                                                                                                    |                |
| SID prefix                                | ora1                                                                                                                                                                                                                                                                                                                                                                             |                |
| Shared disk group for OCR and voting disk | ocrvotedg                                                                                                                                                                                                                                                                                                                                                                        |                |

**Table B-7** Required information for Oracle RAC (*continued*)

| Information                             | Sample value                                                                    | Assigned value |
|-----------------------------------------|---------------------------------------------------------------------------------|----------------|
| Volume for shared disk group            | ■ CVM raw volumes<br>OCR: ocrvol<br>Voting disk: votevol<br>■ CFS<br>ocrvotevol |                |
| Volume options for OCR and voting disk  | nmirror=2                                                                       |                |
| File system on shared volume (CFS)      | /dev/vx/rdisk/ocrvotedg/ocrvotevol                                              |                |
| Mount point for shared file system      | /ocrvote                                                                        |                |
| VCS resource names for Votedisk and OCR | ocrvote_voldg_ocrvotedg<br>ocrvote_mnt_ocrvotedg                                |                |
| Mount point for archive logs            | /oradata                                                                        |                |

## Replicated cluster using VVR worksheet

[Table B-8](#) contains the sample values that may be used when you install and configure CVM and VVR. If applicable, enter the CVM/VVR values for your systems in the following table:

**Table B-8** Replicated cluster using VVR worksheet

| Installation information         | Sample value     | Assigned value |
|----------------------------------|------------------|----------------|
| Host names for Secondary Cluster | mercury, jupiter |                |
| Secondary Cluster Name           | rac_cluster102   |                |
| Secondary cluster ID number      | 102              |                |
| Primary Cluster Address          | 10.10.10.101     |                |
| Primary Cluster Logowner IP      | 10.10.9.101      |                |
| Secondary Cluster Address        | 10.11.10.102     |                |

**Table B-8** Replicated cluster using VVR worksheet (*continued*)

| Installation information                                                      | Sample value                | Assigned value |
|-------------------------------------------------------------------------------|-----------------------------|----------------|
| Secondary Cluster Logowner IP                                                 | 10.11.9.102                 |                |
| RVG Name                                                                      | rac1_rvg                    |                |
| Global Database Name                                                          | vrts                        |                |
| Database Resource Name                                                        | ora1                        |                |
| Database Group Name<br>(depends on cvm, includes resources oracle agent etc.) | oradb1_grp                  |                |
| Srl Volume Name                                                               | rac1_srl                    |                |
| Resolvable Virtual Hostname of the cluster on Primary Site (for VVR)          | rac_clus101_priv            |                |
| Resolvable Virtual Hostname of the cluster on Secondary Site (for VVR)        | rac_clus102_priv            |                |
| Private IP addresses for Secondary Cluster                                    | 192.168.12.3 - 192.168.12.4 |                |

## Replicated cluster using SRDF worksheet

[Table B-9](#) contains the sample values that may be used when you install and configure CVM and VVR. If applicable, enter the CVM/VVR values for your systems in the following table:

**Table B-9** Replicated cluster using SRDF worksheet

| Installation information         | Sample value     | Assigned value |
|----------------------------------|------------------|----------------|
| Host names for Secondary Cluster | mercury, jupiter |                |
| Secondary Cluster Name           | rac_cluster102   |                |
| Secondary cluster ID number      | 102              |                |

**Table B-9** Replicated cluster using SRDF worksheet (*continued*)

| Installation information                                                            | Sample value                   | Assigned value |
|-------------------------------------------------------------------------------------|--------------------------------|----------------|
| Primary Cluster Address                                                             | 10.10.10.101                   |                |
| Primary Cluster Logowner IP                                                         | 10.10.9.101                    |                |
| Secondary Cluster Address                                                           | 10.11.10.102                   |                |
| Secondary Cluster Logowner IP                                                       | 10.11.9.102                    |                |
| RVG Name                                                                            | rac1_rvg                       |                |
| Global Database Name                                                                | vrts                           |                |
| Database Resource Name                                                              | ora1                           |                |
| Database Group Name<br>(depends on cvm,<br>includes resources<br>oracle agent etc.) | oradb1_grp                     |                |
| Srl Volume Name                                                                     | rac1_srl                       |                |
| Resolvable Virtual<br>Hostname of the cluster<br>on Primary Site (for<br>VVR)       | rac_clus101_priv               |                |
| Resolvable Virtual<br>Hostname of the cluster<br>on Secondary Site (for<br>VVR)     | rac_clus102_priv               |                |
| Private IP addresses for<br>Secondary Cluster                                       | 192.168.12.3 -<br>192.168.12.4 |                |

## Required installation information for Oracle Clusterware/Grid Infrastructure

This section describes the information required by the Oracle Universal Installer for installing the Oracle Clusterware/Grid Infrastructure software.

Table B-10 lists the information required by the Oracle Universal Installer when you install Oracle Clusterware.

Table B-10 Required installation information for Oracle Clusterware - Oracle RAC 10g Release 2/Oracle RAC 11g Release 1

| OUI menu                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Specify Inventory Details and Credentials | <p>Enter the full path of the inventory directory and the operating system group name.</p> <p><b>Note:</b> Retain the directory paths and operating system group name suggested by the installer. Changing the information results in failure of certain post-verification tasks.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Specify Home Details                      | <p>Enter the path to the CRS home directory.</p> <p>The Oracle Universal Installer performs product-specific prerequisite checks and verifies that your environment meets all of the minimum requirements for installing Oracle RAC. You must manually verify and confirm the items that are flagged with warnings and items that require manual checks.</p> <p>The OUI displays the full path of the oraInventory logs. Make a note of the log file path to verify the installation at a later time.</p> <p>Select the <b>Manual</b> check box in the Oracle installer window. This is a work around for an Oracle Clusterware installation issue. If you do not select the Manual check box, then the Oracle Universal Installer fails during the product-specific prerequisite check.</p> |
| Specify cluster configuration             | <p>Verify the cluster name and the nodes to be managed by Oracle Clusterware.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Table B-10

Required installation information for Oracle Clusterware - Oracle RAC 10g Release 2/Oracle RAC 11g Release 1 (continued)

| OUI menu                                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Specify network interface usage                   | <p>Identify the planned use for each interface: Public, Private, or Do Not use.</p> <p><b>Note:</b> Make sure that the same private interfaces that you specified at the time of configuring PrivNIC and MultiPrivNIC are listed on the screen.</p> <p><b>Note:</b> Mark the interfaces for the subnet containing the private IP addresses managed by the PrivNIC/MultiPrivNIC agents as 'Private'.</p> <p>The interfaces that are <b>Private</b> are stored in OCR as a 'cluster_interconnect' for database cache fusion traffic. Oracle recommends the use of a common private interface for both Oracle Clusterware and Oracle RAC database.</p> |
| Specify Oracle Cluster Registry (OCR) information | <p>Enter the full path of the location where you want to store the OCR information.</p> <p>For example, if you are storing the OCR information on CFS, enter: <code>/ocrvote/ocr</code>.</p> <p><b>Note:</b> Store the OCR information at the root of the file system. If you place the OCR information under a directory on CFS, the configuration of the CSSD agent after Oracle RAC installation fails.</p> <p>If you are storing the OCR information on raw volumes, enter: <code>/dev/vx/rdisk/ocrvotedg/ocrvol</code></p> <p><b>Note:</b> Select the option <b>External Redundancy</b>. OCR mirroring is performed by CVM.</p>                |

Table B-10

Required installation information for Oracle Clusterware - Oracle RAC 10g Release 2/Oracle RAC 11g Release 1 (continued)

| OUI menu                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Specify voting disk location | <p>Enter the full path of the location where you want to store the voting disk information.</p> <p>For example, if you are storing the voting disk information on CFS, enter: <code>/ocrvote/vote</code></p> <p><b>Note:</b> Store the voting disk information at the root of the file system. If you place the voting disk information under a directory on CFS, the configuration of the CSSD agent after Oracle RAC installation fails.</p> <p>If you are storing the voting disk information on raw volumes, enter: <code>/dev/vx/rdisk/ocrvotedg/votevol</code></p> <p><b>Note:</b> Select the option <b>External Redundancy</b>. Voting disk redundancy is provided by CVM.</p> |

Table B-11 lists the information required by the Oracle Universal Installer when you install Oracle Grid Infrastructure.

Table B-11

Required installation information for Oracle Grid Infrastructure - Oracle RAC 11g Release 2

| OUI menu                       | Description                                                                                                                                                                                                                                                                         |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Select installation option     | Select the option <b>Install and Configure Grid Infrastructure for a Cluster</b> .                                                                                                                                                                                                  |
| Select installation type       | Select the option <b>Advanced Installation</b> .                                                                                                                                                                                                                                    |
| Specify cluster configuration  | Enter the SCAN name for the cluster that will be used by the database clients to connect to databases within the cluster.                                                                                                                                                           |
| Grid Plug and Play information | <p>Provide the following information:</p> <ul style="list-style-type: none"><li>■ Name of the cluster</li><li>■ SCAN name<br/>The SCAN address on the domain name server (DNS) must resolve to three addresses (recommended) or at least one address.</li><li>■ SCAN port</li></ul> |



Table B-11

Required installation information for Oracle Grid Infrastructure - Oracle RAC 11g Release 2 (continued)

| OUI menu                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Specify network interface usage | <p>Identify the planned use for each interface: Public, Private, or Do Not use.</p> <p><b>Note:</b> Make sure that the same private interfaces that you specified at the time of configuring PrivNIC and MultiPrivNIC are listed on the screen.</p> <p><b>Note:</b> Mark the interfaces for the subnet containing the private IP addresses managed by the PrivNIC/MultiPrivNIC agents as 'Private'.</p> <p>The interfaces that are <b>Private</b> are stored in GPnP profile as a 'cluster_interconnect' for Oracle Grid Infrastructure communication and database cache fusion traffic.</p>                                                      |
| Storage option information      | Select the option <b>Shared File System</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| OCR storage option              | <p>Enter the full path of the location where you want to store the OCR information.</p> <p>For example, if you are storing the OCR information on CFS, enter: <code>/ocrvote/ocr</code>.</p> <p><b>Note:</b> Select the option <b>External Redundancy</b> when you install Oracle Clusterware/Grid Infrastructure. Mirroring is performed by CVM.</p> <p>If you are storing the OCR information on an ASM disk group that uses CVM raw volumes, enter:<br/><code>/dev/vx/rdisk/ocrvotedg/ocrvol</code></p> <p><b>Note:</b> Select the option <b>External</b> when you create the ASM disk group and instances. Mirroring is performed by CVM.</p> |

Table B-11

Required installation information for Oracle Grid Infrastructure - Oracle RAC 11g Release 2 (continued)

| OUI menu                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Voting Disk storage option    | <p>Enter the full path of the location where you want to store the voting disk information.</p> <p>For example, if you are storing the voting disk information on CFS, enter: <code>/ocrvote/vote</code></p> <p><b>Note:</b> Select the option <b>External Redundancy</b> when you install Oracle Clusterware/Grid Infrastructure. Mirroring is performed by CVM.</p> <p>If you are storing the voting disk information on an ASM disk group that uses CVM raw volumes, enter: <code>/dev/vx/rdisk/ocrvotedg/votevol</code></p> <p><b>Note:</b> Select the option <b>External</b> when you create the ASM disk group and instances. Mirroring is performed by CVM.</p> |
| Specify installation location | Enter the full path to the Oracle base directory and the Oracle Grid Infrastructure home directory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Create inventory              | Enter the full path to the Oracle inventory directory where you want to store the installation files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Required installation information for Oracle database

This section describes the information required by the Oracle Universal Installer for installing the Oracle Clusterware/Grid Infrastructure software.

[Table B-12](#) lists the information required by the Oracle Universal Installer when you install Oracle database software.

Table B-12

Required installation information for Oracle database - Oracle RAC 10g Release 2/Oracle RAC 11g Release 1

| OUI menu                 | Description                                                      |
|--------------------------|------------------------------------------------------------------|
| Select installation type | Select <b>Enterprise Edition</b> .                               |
| Specify home details     | Review or enter the ORACLE_HOME and ORACLE_BASE directory paths. |

**Table B-12** Required installation information for Oracle database - Oracle RAC 10g Release 2/Oracle RAC 11g Release 1 (*continued*)

| OUI menu                                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Specify Hardware Cluster Installation Mode</b> | <p>Select <b>Cluster Installation</b>.</p> <p>Select the nodes on which you want to install the Oracle RAC database software.</p> <p>The Oracle Universal Installer runs a product-specific prerequisite check. Any items that are flagged must be manually checked and configured. The OUI displays the full path of the oraInventory logs. Make a note of the log file path to verify the installation at a later time.</p> |
| <b>Select Configuration Option</b>                | <p>Select the option <b>Install database software only</b>.</p> <p><b>Note:</b> Do not select the option <b>Create a database</b>. Symantec recommends that you create the database later.</p>                                                                                                                                                                                                                                |

[Table B-13](#) lists the information required by the Oracle Universal Installer when you install Oracle database software.

**Table B-13** Required installation information for Oracle database - Oracle RAC 11g Release 2

| OUI menu                             | Description                                                                                                                                                                                                                  |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Select installation option</b>    | Select the option <b>Install database software only</b> .                                                                                                                                                                    |
| <b>Node selection</b>                | <p>Select <b>Real Application Clusters database installation</b>.</p> <p>Select the nodes on which the Oracle RAC database software must be installed.</p>                                                                   |
| <b>Select database edition</b>       | Select <b>Enterprise Edition</b> .                                                                                                                                                                                           |
| <b>Specify installation location</b> | <p>Review or enter the ORACLE_BASE and ORACLE_HOME directory paths.</p> <p>The Oracle Universal Installer runs product-specific prerequisite checks. Any items that are flagged must be manually checked and configured.</p> |



# Sample configuration files

This appendix includes the following topics:

- [About VCS configuration file](#)
- [About the LLT and GAB configuration files](#)
- [About I/O fencing configuration files](#)
- [Sample configuration files](#)

## About VCS configuration file

This section provides a high-level overview of the contents of the VCS configuration file after the SF Oracle RAC installation. Review the configuration file after the SF Oracle RAC installation and before the Oracle installation.

The configuration file includes the following information:

- The "include" statements list the various VCS types files (`.cf`) for SF Oracle RAC.  
The files are located in the `/etc/VRTSvcs/conf/config` directory. These files define the agents that control the resources in the cluster.
- The cluster definition, with the cluster name provided during installation (for example, `rac_cluster101`), includes the names of users and administrators of the cluster. The `UseFence = SCSI3` attribute is not automatically present; you must manually add it after the installation.  
See [“Setting up disk-based I/O fencing manually”](#) on page 178.
- The `main.cf` includes the `cvm` service group. The service group includes definitions for monitoring the CFS and the CVM resources. The `CVMCluster` agent resource definition indicates that the nodes use GAB for messaging operations.

The cvm group has the Parallel attribute set to 1. This value enables the resources to run in parallel on each node in the system list.

## About the LLT and GAB configuration files

Low Latency Transport (LLT) and Group Membership and Atomic Broadcast (GAB) are VCS communication services. LLT requires `/etc/llthosts` and `/etc/llttab` files. GAB requires `/etc/gabtab` file.

[Table C-1](#) lists the LLT configuration files and the information that these files contain.

Table C-1                    LLT configuration files

| File             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/default/llt | <p>This file stores the start and stop environment variables for LLT:</p> <ul style="list-style-type: none"><li>■ LLT_START—Defines the startup behavior for the LLT module after a system reboot. Valid values include:<ul style="list-style-type: none"><li>1—Indicates that LLT is enabled to start up.</li><li>0—Indicates that LLT is disabled to start up.</li></ul></li><li>■ LLT_STOP—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include:<ul style="list-style-type: none"><li>1—Indicates that LLT is enabled to shut down.</li><li>0—Indicates that LLT is disabled to shut down.</li></ul></li></ul> <p>The installer sets the value of these variables to 1 at the end of SF Oracle RAC configuration.</p> |
| /etc/llthosts    | <p>The file <code>llthosts</code> is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file must be identical on each node in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.</p> <p>For example, the file <code>/etc/llthosts</code> contains the entries that resemble:</p> <pre>0      galaxy 1      nebula</pre>                                                                                                                                                                                                                                                                                                         |

**Table C-1** LLT configuration files (*continued*)

| File        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/llttab | <p>The file <code>llttab</code> contains the information that is derived during installation and used by the utility <code>lltconfig(1M)</code>. After installation, this file lists the private network links that correspond to the specific system. For example, the file <code>/etc/llttab</code> contains the entries that resemble the following:</p> <p>■ For Solaris SPARC:</p> <pre> set-node galaxy set-cluster 2 link bge1 /dev/bge1 - ether - - link bge1 /dev/bge2 - ether - - </pre> <p>■ For Solaris x64:</p> <pre> set-node galaxy set-cluster 2 link e1000g1 /dev/e1000g:1 - ether - - link e1000g2 /dev/e1000g:2 - ether - - </pre> <p>The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the <code>link</code> command. These lines identify the two network cards that the LLT protocol uses.</p> <p>If you configured a low priority link under LLT, the file also includes a "link-lowpri" line.</p> <p>Refer to the <code>llttab(4)</code> manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the <code>llttab</code> file.</p> |

[Table C-2](#) lists the GAB configuration files and the information that these files contain.

Table C-2            GAB configuration files

| File             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/default/gab | <p>This file stores the start and stop environment variables for GAB:</p> <ul style="list-style-type: none"><li>■ GAB_START—Defines the startup behavior for the GAB module after a system reboot. Valid values include:<br/>1—Indicates that GAB is enabled to start up.<br/>0—Indicates that GAB is disabled to start up.</li><li>■ GAB_STOP—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include:<br/>1—Indicates that GAB is enabled to shut down.<br/>0—Indicates that GAB is disabled to shut down.</li></ul> <p>The installer sets the value of these variables to 1 at the end of SF Oracle RAC configuration.</p>                                                          |
| /etc/gabtab      | <p>After you install SF Oracle RAC, the file /etc/gabtab contains a <code>gabconfig(1)</code> command that configures the GAB driver for use.</p> <p>The file /etc/gabtab contains a line that resembles:</p> <pre>/sbin/gabconfig -c -nN</pre> <p>The <code>-c</code> option configures the driver for use. The <code>-nN</code> specifies that the cluster is not formed until at least <i>N</i> nodes are ready to form the cluster. Symantec recommends that you set <i>N</i> to be the total number of nodes in the cluster.</p> <p><b>Note:</b> Symantec does not recommend the use of the <code>-c -x</code> option for <code>/sbin/gabconfig</code>. Using <code>-c -x</code> can lead to a split-brain condition.</p> |

## About I/O fencing configuration files

Table C-3 lists the I/O fencing configuration files.



Table C-3 I/O fencing configuration files

| File               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/default/vxfen | <p>This file stores the start and stop environment variables for I/O fencing:</p> <ul style="list-style-type: none"><li>■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include:<br/>1—Indicates that I/O fencing is enabled to start up.<br/>0—Indicates that I/O fencing is disabled to start up.</li><li>■ VXFEN_STOP—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include:<br/>1—Indicates that I/O fencing is enabled to shut down.<br/>0—Indicates that I/O fencing is disabled to shut down.</li></ul> <p>The installer sets the value of these variables to 1 at the end of SF Oracle RAC configuration.</p> |
| /etc/vxfendg       | <p>This file includes the coordinator disk group information.</p> <p>This file is not applicable for server-based fencing.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table C-3 I/O fencing configuration files (continued)

| File           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/vxfenmode | <p>This file contains the following parameters:</p> <ul style="list-style-type: none"><li>■ vxfen_mode<ul style="list-style-type: none"><li>■ scsi3—For disk-based fencing</li><li>■ customized—For server-based fencing</li><li>■ disabled—To run the I/O fencing driver but not do any fencing operations.</li></ul></li><li>■ vxfen_mechanism<p>This parameter is applicable only for server-based fencing. Set the value as cps.</p></li><li>■ scsi3_disk_policy<ul style="list-style-type: none"><li>■ dmp—Configure the vxfen module to use DMP devices<br/>The disk policy is dmp by default. If you use iSCSI devices, you must set the disk policy as dmp.</li><li>■ raw—Configure the vxfen module to use the underlying raw character devices</li></ul></li></ul> <p><b>Note:</b> You must use the same SCSI-3 disk policy on all the nodes.</p> <ul style="list-style-type: none"><li>■ security<p>This parameter is applicable only for server-based fencing.</p><p>1—Indicates that Symantec Product Authentication Service is used for CP server communications. This setting is the default.</p><p>0—Indicates that communication with the CP server is in non-secure mode.</p><p><b>Note:</b> The CP server and the SF Oracle RAC clusters must have the same security setting.</p></li><li>■ List of coordination points<p>This list is required only for server-based fencing configuration.</p><p>Coordination points in a server-based fencing can include coordinator disks, CP servers, or a mix of both. If you use coordinator disks, you must create a coordinator disk group with the coordinator disk names.</p><p>Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points.</p></li><li>■ single_cp<p>This parameter is applicable only for server-based fencing which uses a single highly available CP server as its coordination point.</p></li></ul> |

Table C-3 I/O fencing configuration files (continued)

| File          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/vxfentab | <p>When I/O fencing starts, the vxfen startup script creates this <code>/etc/vxfentab</code> file on each node. The startup script uses the contents of the <code>/etc/vxfendg</code> and <code>/etc/vxfenmode</code> files. Any time a system is rebooted, the fencing driver reinitializes the <code>vxfentab</code> file with the current list of all the coordinator points.</p> <p><b>Note:</b> The <code>/etc/vxfentab</code> file is a generated file; do not modify this file.</p> <p>For disk-based I/O fencing, the <code>/etc/vxfentab</code> file on each node contains a list of all paths to each coordinator disk. An example of the <code>/etc/vxfentab</code> file in a disk-based fencing configuration on one node resembles as follows:</p> <ul style="list-style-type: none"><li>■ Raw disk:<br/><br/><code>/dev/rdisk/clt1d0s2</code><br/><code>/dev/rdisk/c2t1d0s2</code><br/><code>/dev/rdisk/c3t1d2s2</code></li><li>■ DMP disk:<br/><br/><code>/dev/vx/rdmp/clt1d0s2</code><br/><code>/dev/vx/rdmp/c2t1d0s2</code><br/><code>/dev/vx/rdmp/c3t1d0s2</code></li></ul> <p>For server-based fencing, the <code>/etc/vxfentab</code> file also includes the security settings information.</p> <p>For server-based fencing with single CP server, the <code>/etc/vxfentab</code> file also includes the <code>single_cp</code> settings information.</p> |

## Sample configuration files

SF Oracle RAC provides several sample configuration files illustrating various scenarios. You may use the sample files as a guideline for setting up your cluster environment. These sample files are located at `/etc/VRTSvcs/conf/sample_rac/`.

This section briefly describes each of the sample files and illustrates the service group configuration for each of them. The section does not include a copy of the `main.cf` files.

The following sample files are discussed in this section:

- [sfrac01\\_main.cf](#) file
- [sfrac02\\_main.cf](#) file
- [sfrac03\\_main.cf](#) file

- [sfrac04\\_main.cf file](#)
- [sfrac05\\_main.cf file](#)
- [sfrac06\\_main.cf file](#)
- [sfrac07\\_main.cf and sample08\\_main.cf files](#)
- [sfrac09\\_main.cf and sample10\\_main.cf files](#)
- [sfrac11\\_main.cf file](#)
- [sfrac12\\_main.cf and sample13\\_main.cf files](#)
- [sfrac14\\_main.cf file](#)

## sfrac01\_main.cf file


The sample configuration is applicable only for Oracle RAC 10g installations.

This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- The database uses the VCSIPC for database cache fusion.
- Only one private IP address is configured for Oracle Clusterware. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

[Figure C-1](#) illustrates the configuration.

**Figure C-1** Service group configuration for sfrac01\_main.cf file




## sfrac02\_main.cf file

This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.  
The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CVM raw volumes.

Figure C-2 illustrates the configuration.

**Figure C-2** Service group configuration for sfrac02\_main.cf file




## sfrac03\_main.cf file

This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting two databases.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- One IP address (on bge1) is shared by Oracle Clusterware and one of the databases for cache fusion.  
The second IP address (on bge2) is used by the second database for cache fusion.  
The private IP addresses are managed by the MultiPrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

Figure C-3 illustrates the configuration.

**Figure C-3** Service group configuration for sfrac03\_main.cf file



## sfrac04\_main.cf file

This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting two databases.
- The Oracle database is stored on CFS.
- Database is not managed by VCS. Oracle Clusterware starts, stops, and monitors the databases.

The CRSResource agent monitors the status of the database, the VIP resource, and the listener resource configured under Oracle Clusterware.

---

**Note:** The CFSSMount and CVMVolDg resources for Oracle database can not be set as critical resources in the group.

The CRSResource agent appears FAULTED until Oracle Clusterware brings up the database.


---

- The database uses the Oracle UDP IPC for database cache fusion.
- One IP address (on bge1) is shared by Oracle Clusterware and one of the databases for cache fusion.  
The second IP address (on bge2) is used by the second database for cache fusion.  
The private IP addresses are managed by the MultiPrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

Figure C-4 illustrates the configuration.



**Figure C-4** Service group configuration for sfrac04\_main.cf file



## sfrac05\_main.cf file


This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting single database.
- The Oracle database is stored on ASM.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.

- A common IP address is used by Oracle Clusterware and database cache fusion. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

Figure C-5 illustrates the configuration.

**Figure C-5** Service group configuration for sfrac05\_main.cf file



## sfrac06\_main.cf file

This sample file describes the following configuration:


- A two node SF Oracle RAC cluster hosting two databases.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle. The agent starts, stops, and monitors the database.
- The database uses Oracle UDP IPC for cache fusion. A dedicated link is used for each database .
- One private IP address (on bge1) is used by Oracle Clusterware.

The private IP address on bge2 is used by one of the databases for cache fusion. The private IP address on NIC3 is used by the other database for cache fusion. The private IP addresses are managed by the MultiPrivNIC agent for high availability.

- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

Figure C-6 illustrates the configuration.

**Figure C-6** Service group configuration for sfrac06\_main.cf file



## sfrac07\_main.cf and sample08\_main.cf files

The sample configuration, `sfrac07_main.cf`, describes a disaster recovery configuration for the primary site. The sample configuration, `sample08_main.cf`, describes a disaster recovery configuration for the secondary site. The configuration uses VVR for replicating data between the sites.


This sample file describes the following configuration:

- Two SF Oracle RAC clusters, comprising two nodes each, hosted at different geographical locations.

- A single Oracle database that is stored on CFS.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.  
The private IP address is managed by the PrivNIC agent for high availability.
- One virtual IP address must be configured under the `ClusterService` group on each site for inter-cluster communication.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.
- Veritas Volume Replicator (VVR) is used to replicate data between the sites.
- The shared volumes replicated across the sites are configured under the RVG group.
- The replication link used by VVR for communicating log information between sites are configured under the `rlogowner` group. This is a failover group that will be online on only one of the nodes in the cluster at each site.
- The database group will be online on the primary cluster. The `RVGSharedPri` resource determines where the database group will be brought online.
- The database group is configured as a global group by specifying the clusters on the primary and secondary sites as values for the `ClusterList` group attribute.

Figure C-7 illustrates the configuration. The service group configuration is the same on the primary and secondary site. The availability of groups (online/offline) differ between the sites.

**Figure C-7** Service group configuration for sfrac07\_main.cf and sfrac08\_main.cf files



sfrac09\_main.cf and sample10\_main.cf files

The sample configuration, **sfrac09\_main.cf**, describes a disaster recovery configuration for the primary site. The sample configuration, **sample10\_main.cf**, describes a disaster recovery configuration for the secondary site. The sample configuration uses EMC SRDF technology for replicating data between the sites.

---

**Note:** You can use other supported hardware-based replication technologies with this configuration.

---

This sample file describes the following configuration:

- Two SF Oracle RAC clusters, comprising two nodes each, hosted at different geographical locations.
- A single Oracle database that is stored on CFS.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.  
The private IP address is managed by the PrivNIC agent for high availability.
- One virtual IP address must be configured under the `ClusterService` group on each site for inter-cluster communication.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.
- EMC SRDF is used to replicate data between the sites.
- The SRDF disk groups that are replicated across the sites using SRDF technology and the replication mode are specified under the SRDF resource in the database group. The CVM disk group that comprises the SRDF disk group must be configured under the `CVMVolDg` resource in the database group.
- The database group will be online on the primary cluster. The SRDF resource determines where the database group will be brought online.
- The database group is configured as a global group by specifying the clusters on the primary and secondary sites as values for the `ClusterList` group attribute.

Figure C-8 illustrates the configuration on the primary site.

Figure C-8      Service group configuration for sfrac09\_main.cf file





Figure C-9 illustrates the configuration on the secondary site.

**Figure C-9** Service group configuration for sfrac10\_main.cf file



## sfrac11\_main.cf file

This sample file describes the following configuration:

- An SF Oracle RAC campus cluster with four nodes hosted across two sites.
- Each site comprises two nodes of the cluster hosting a single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- The IP address on bge1 is used by Oracle Clusterware.  
The second IP address on NIC2 is used for Oracle database cache fusion.




The private IP addresses are managed by the MultiPrivNIC agent for high availability.

- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.
- Group the hosts at each physical site into separate logical system zones using the SystemZones attribute.

Figure C-10 illustrates the configuration.

**Figure C-10** Service group configuration for sfrac11\_main.cf file



## sfrac12\_main.cf and sample13\_main.cf files

The sample configuration, sfrac12\_main.cf, describes a disaster recovery configuration for the primary site. The sample configuration, sample13\_main.cf, describes a disaster recovery configuration for the secondary site with fire-drill capability. The sample configuration uses Hitachi True Copy technology for replicating data between the sites.

---

**Note:** You can use other supported hardware-based replication technologies with this configuration.

---

This sample file describes the following configuration:

- Two SF Oracle RAC clusters, comprising two nodes each, hosted at different geographical locations.
- A single Oracle database that is stored on CFS.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.  
The private IP address is managed by the PrivNIC agent for high availability.
- One virtual IP address must be configured under the `ClusterService` group on each site for inter-cluster communication.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.
- Hitachi True Copy is used to replicate data between the sites.
- The HTC disk groups that are replicated across the sites using HTC technology and the replication mode are specified under the HTC resource in the database group. The CVM disk group that comprises the HTC disk group must be configured under the `CVMVolDg` resource in the database group.
- The database group will be online on the primary cluster. The HTC resource determines where the database group will be brought online.
- The database group is configured as a global group by specifying the clusters on the primary and secondary sites as values for the `ClusterList` group attribute.
- The database group `oradb_grp_fd` on the secondary is configured for fire drill.
- When the group `oradb_grp_fd` is brought online, the HTCSnap creates a snapshot of the disk group configured under the HTC resource in the database group `oradg_grp`.  
Further, the Oracle database and the associated volumes and mount points configured under the service group `oradb_grp_fd` are brought online using the snapshots created by HTCSnap.

Figure C-11 illustrates the configuration on the primary site.

Figure C-11      Service group configuration for sfrac12\_main.cf file





Figure C-12 illustrates the configuration on the secondary site.

**Figure C-12** Service group configuration for sfrac13\_main.cf file




## sfrac14\_main.cf file

This sample file describes the following configuration:

- A two node SF Oracle RAC cluster hosting single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle.  
The agent starts, stops, and monitors the database.
- The database uses the Oracle UDP IPC for database cache fusion.
- A common IP address is used by Oracle Clusterware and database cache fusion.  
The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.
- The cluster is configured as a secure cluster by setting the SecureClus attribute to 1. The VxSS group starts the authentication daemon `vxatd`.

Figure C-13 illustrates the configuration.

**Figure C-13** Service group configuration for sfrac14\_main.cf file



## Sample configuration files for CP server

The following are example main.cf files for a CP server that is hosted on a single node, and a CP server that is hosted on an SFHA cluster.

- The main.cf file for a CP server that is hosted on a single node:  
See “[Sample main.cf file for CP server hosted on a single node that runs VCS](#)” on page 694.
- The main.cf file for a CP server that is hosted on an SFHA cluster:  
See “[Sample main.cf file for CP server hosted on a two-node SFHA cluster](#)” on page 696.

---

**Note:** The CP server supports Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with SF Oracle RAC clusters. The following example main.cf files use IPv4 addresses.

---

## Sample main.cf file for CP server hosted on a single node that runs VCS

The following is an example of a single CP server node main.cf.

For this CP server single node main.cf, note the following values:

- Cluster name: cps1
- Node name: mycps1

```
include "types.cf"

// cluster name: cps1
// CP server: mycps1

cluster cps1 (
    UserNames = { admin = bMNfMHmJNiNNlVNHMK, haris = fopKojNvpHouNn,
                  "mycps1.symantecexample.com@root@vx" = aj,
                  "root@mycps1.symantecexample.com" = hq }
    Administrators = { admin, haris,
                       "mycps1.symantecexample.com@root@vx",
                       "root@mycps1.symantecexample.com" }
    SecureClus = 1
    HacliUserLevel = COMMANDROOT
)

system mycps1 (
)

group CPSSG (
    SystemList = { mycps1 = 0 }
    AutoStartList = { mycps1 }
)

IP cpsvip (
    Device @mycps1 = bge0
    Address = "10.209.3.1"
    NetMask = "255.255.252.0"
)

NIC cpsnic (
    Device @mycps1 = bge0
)
```

```
    Process vxcpserve (
        PathName = "/opt/VRTScps/bin/vxcpserve"
        ConfInterval = 30
        RestartLimit = 3
    )

cpsvip requires cpsnic
vxcpserve requires cpsvip

// resource dependency tree
//
// group CPSSG
// {
//   Process vxcpserve
//   {
//       IP cpsvip
//       {
//           NIC cpsnic
//       }
//   }
// }

group VxSS (
    SystemList = { mycps1 = 0 }
    Parallel = 1
    AutoStartList = { mycps1 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Phantom phantom_vxss (
)

ProcessOnOnly vxatd (
    IgnoreArgs = 1
    PathName = "/opt/VRTSat/bin/vxatd"
)

// resource dependency tree
```

```
//  
//    group VxSS  
//    {  
//    Phantom phantom_vxss  
//    ProcessOnOnly vxatd  
//    }
```

## Sample main.cf file for CP server hosted on a two-node SFHA cluster

The following is an example of a main.cf, where the CP server is hosted on an SFHA cluster.

For this CP server hosted on an SFHA cluster main.cf, note the following values:

- Cluster name: cps1
- Nodes in the cluster: mycps1, mycps2

```
include "types.cf"  
include "CFSTypes.cf"  
include "CVMTypes.cf"  
  
// cluster: cps1  
// CP servers:  
// mycps1  
// mycps2  
  
cluster cps1 (  
    UserNames = { admin = ajkCjeJgkFkkIskEjh,  
                  "mycps1.symantecexample.com@root@vx" = JK,  
                  "mycps2.symantecexample.com@root@vx" = dl }  
    Administrators = { admin, "mycps1.symantecexample.com@root@vx",  
                       "mycps2.symantecexample.com@root@vx" }  
    SecureClus = 1  
)  
  
system mycps1 (  
)  
  
system mycps2 (  
)  
  
group CPSSG (  

```



```
SystemList = { mycps1 = 0, mycps2 = 1 }
AutoStartList = { mycps1, mycps2 } )

DiskGroup cpsdg (
    DiskGroup = cps_dg
)

IP cpsvip (
    Device @mycps1 = bge0
    Device @mycps2 = bge0
    Address = "10.209.81.88"
    NetMask = "255.255.252.0"
)

Mount cpsmount (
    MountPoint = "/etc/VRTScps/db"
    BlockDevice = "/dev/vx/dsk/cps_dg/cps_volume"
    FSType = vxfs
    FsckOpt = "-y"
)

NIC cpsnic (
    Device @mycps1 = bge0
    Device @mycps2 = bge0
)

Process vxcperv (
    PathName = "/opt/VRTScps/bin/vxcperv"
)

Volume cpsvol (
    Volume = cps_volume
    DiskGroup = cps_dg
)
```

```
cpsmount requires cpsvol
cpsvip requires cpsnic
cpsvol requires cpsdg
vxcperv requires cpsmount
vxcperv requires cpsvip
```

```
// resource dependency tree
```

```
//
// group CPSSG
// {
//   Process vxcpserv
//   {
//     Mount cpsmount
//     {
//       Volume cpsvol
//       {
//         DiskGroup cpsdg
//       }
//     }
//     IP cpsvip
//     {
//       NIC cpsnic
//     }
//   }
// }

group VxSS (
    SystemList = { mycps1 = 0, mycps2 = 1 }
    Parallel = 1
    AutoStartList = { mycps1, mycps2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Phantom phantom_vxss (
)

ProcessOnOnly vxatd (
    IgnoreArgs = 1
    PathName = "/opt/VRTSat/bin/vxatd"
)

// resource dependency tree
//
// group VxSS
// {
//   Phantom phantom_vxss
```

```
// ProcessOnOnly vxatd  
// }
```



# Setting up inter-system communication

This appendix includes the following topics:

- [About configuring ssh or rsh using the Veritas installer](#)
- [Setting up inter-system communication](#)

## About configuring ssh or rsh using the Veritas installer

The installer can configure passwordless secure shell (ssh) or remote shell (rsh) communications among systems. The installer uses the ssh or rsh daemon that comes bundled with the operating system. During an installation, you choose the communication method that you want to use. You then provide the installer with the superuser passwords for the systems where you plan to install. The ssh or rsh communication among the systems is removed when the installation process completes, unless the installation abruptly terminates. If installation terminated abruptly, use the installation script's `-comcleanup` option to remove the ssh or rsh configuration from the systems.

See [“Installation script options”](#) on page 638.

In most installation, configuration, upgrade (where necessary), and uninstallation scenarios, the installer can configure ssh or rsh on the target systems. In the following scenarios, you need to set up ssh or rsh manually:

- When the root broker is outside of the cluster that you plan to configure.
- When you add new nodes to an existing cluster.
- When the nodes are in a sub-cluster during a phased upgrade.

- When you perform installer sessions using a response file.

## Setting up inter-system communication

If you manually need to set up a communication mode, refer to these procedures. You must have root privilege to issue ssh or rsh commands on all systems in the cluster. If ssh is used to communicate between systems, it must be configured in a way such that it operates without requests for passwords or passphrases. Similarly, rsh must be configured in such a way to not prompt for passwords.

If system communication is not possible between systems using ssh or rsh, contact Symantec Support. See <http://support.symantec.com>.

### Setting up ssh on cluster systems

Use the Secure Shell (ssh) to install SF Oracle RAC on all systems in a cluster from a system outside of the cluster. Before you start the installation process, verify that ssh is configured correctly.

Use Secure Shell (ssh) to do the following:

- Log on to another system over a network
- Execute commands on a remote system
- Copy files from one system to another

The ssh shell provides strong authentication and secure communications over channels. It is intended to replace rlogin, rsh, and rcp.

### Configuring ssh

The procedure to configure ssh uses OpenSSH example file names and commands.

---

**Note:** You can configure ssh in other ways. Regardless of how ssh is configured, complete the last step in the example to verify the configuration.

---

### To configure ssh

- 1 Log in as root on the source system from which you want to install the Veritas product.

- 2 To generate a DSA key pair on the source system, type the following:

```
# ssh-keygen -t dsa
```

System output similar to the following is displayed:

```
Generating public/private dsa key pair.  
Enter file in which to save the key (//.ssh/id_dsa):
```

- 3 Press **Enter** to accept the default location of /.ssh/id\_dsa. System output similar to the following is displayed:

```
Enter passphrase (empty for no passphrase):
```

- 4 Do not enter a passphrase. Press **Enter**. Enter same passphrase again. Press **Enter** again.

- 5 Make sure the /.ssh directory is on all the target installation systems. If that directory is absent, create it on the target system and set the write permission to root only:

```
# mkdir /.ssh  
# chmod go-w /  
# chmod 700 /.ssh  
# chmod go-rwx /.ssh
```

- 6 Make sure the secure file transfer program (SFTP) is enabled on all the target installation systems. To enable SFTP, the /etc/ssh/sshd\_config file must contain the following two lines:

```
PermitRootLogin yes  
Subsystem sftp /usr/lib/ssh/sftp-server
```

- 7 If the lines are not there, add them and restart SSH. To restart SSH on Solaris 10, type the following command:

```
# svcadm restart ssh
```

To restart on Solaris 9, type the following commands:

```
# /etc/init.d/sshd stop  
# /etc/init.d/sshd start
```

- 8** To copy the public DSA key, `/.ssh/id_dsa.pub` to each target system, type the following commands:

```
# sftp target_sys
```

If you run this step for the first time on a system, output similar to the following appears:

```
Connecting to target_sys...
The authenticity of host 'target_sys (10.182.00.00)'
can't be established. DSA key fingerprint is
fb:6f:9e:61:91:9e:44:6b:87:86:ef:68:a6:fd:87:7d.
Are you sure you want to continue connecting (yes/no)?
```

- 9** Enter **yes**. Output similar to the following is displayed:

```
Warning: Permanently added 'target_sys,10.182.00.00'
(DSA) to the list of known hosts.
root@target_sys password:
```

- 10** Enter the root password.

- 11** At the sftp prompt, type the following command:

```
sftp> put /.ssh/id_dsa.pub
```

The following output is displayed:

```
Uploading /.ssh/id_dsa.pub to /id_dsa.pub
```

- 12** To quit the SFTP session, type the following command:

```
sftp> quit
```

- 13** To begin the ssh session on the target system, type the following command:

```
# ssh target_sys
```

- 14** Enter the root password at the prompt:

```
password:
```

- 15** After you log in, enter the following command to append the authorization key to the `id_dsa.pub` file:

```
# cat /id_dsa.pub >> /.ssh/authorized_keys
```



- 16** Delete the `id_dsa.pub` public key file. Before you delete this public key file, make sure to complete the following tasks:

- The file is copied to the target (host) system
- The file is added to the authorized keys file

To delete the `id_dsa.pub` public key file, type the following command:

```
# rm /id_dsa.pub
```

- 17** To log out of the ssh session, type the following command:

```
# exit
```

- 18** When you install from a source system that is also an installation target, add the local system `id_dsa.pub` key to the local `/.ssh/authorized_key` file. The installation can fail if the installation source system is not authenticated.

- 19** Run the following commands on the source installation system. These commands bring the private key into the shell environment and makes the key globally available for the user root:

```
# exec /usr/bin/ssh-agent $SHELL
# ssh-add
Identity added: /.ssh/identity
```

This step is shell-specific and is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

- 20** To verify that you can connect to the target system, type the following command:

```
# ssh -l root target_sys uname -a
```

The commands should execute on the remote system without any requests for a passphrase or password from the system.



# Automatic Storage Management

This appendix includes the following topics:

- [About ASM in SF Oracle RAC environments](#)
- [ASM configuration with SF Oracle RAC](#)
- [Configuring ASM in SF Oracle RAC environments](#)
- [Sample configuration file Veritas CVM and ASM main.cf file](#)

## About ASM in SF Oracle RAC environments

ASM is an integrated storage management solution from Oracle RAC that combines file system and volume management capabilities.

ASM can be configured with Cluster Volume Manager (CVM) for better performance and availability. CVM mirrored volumes with dynamic multipathing improves data access performance and offers continuous data availability in large heterogeneous SAN environments. You can create CVM disk groups and volumes for use as ASM disk groups and configure the ASM disk groups to be managed by the Veritas ASMDG agent. The ASMDG agent mounts, unmounts, and monitors the ASM disk groups.

Depending on the Oracle RAC version installed, the following are supported on ASM for that version:

|                                                   |                                                                                  |
|---------------------------------------------------|----------------------------------------------------------------------------------|
| Oracle RAC 10g Release 2/Oracle RAC 11g Release 1 | Data files, control files, online redo logs, archive log files, and backup files |
|---------------------------------------------------|----------------------------------------------------------------------------------|

|                                                                                                         |                                                                                                                                                                              |
|---------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oracle RAC 11g Release 2                                                                                | Oracle Cluster Registry devices (OCR), voting disk, data files, control files, online redo logs, archive log files, and backup files                                         |
| Depending on the Oracle RAC version installed, the following are not supported on ASM for that version: |                                                                                                                                                                              |
| Oracle RAC 10g Release 2/Oracle RAC 11g Release 1                                                       | Oracle binaries, trace files, alert logs, export files, tar files, core files, Oracle Cluster Registry devices (OCR), voting disk, application binaries and application data |
| Oracle RAC 11g Release 2                                                                                | Oracle binaries, trace files, alert logs, export files, tar files, core files, application binaries and application data                                                     |

## ASM configuration with SF Oracle RAC

Configure ASM disk groups over CVM volumes in SF Oracle RAC environments. The CVM volumes are mirrored for high availability of data and leverage dynamic multipathing to access the shared storage.

Figure E-1 illustrates the configuration of ASM disk groups over CVM.

Figure E-1 ASM disk groups over CVM




Figure E-2 illustrates the following supported configuration for all supported versions of Oracle RAC:

- ASM disk groups configured over CVM volumes
- Oracle Clusterware and database binaries stored locally
- Oracle database files stored on ASM configured over CVM  
The Oracle databases are managed by Oracle Clusterware.
- Oracle Cluster Registry and voting disk stored on CFS

**Figure E-2** Supported Oracle RAC configuration (all supported Oracle RAC versions)





Figure E-3 illustrates the following supported configuration for Oracle RAC 11g Release 2:

- ASM disk groups configured over CVM volumes
- Oracle Clusterware and database binaries stored locally
- Oracle database files stored on ASM configured over CVM  
The Oracle databases are managed by Oracle Clusterware.
- Oracle Cluster Registry and voting disk stored on ASM over CVM volumes

**Figure E-3** Supported Oracle RAC configuration (Oracle RAC 11g Release 2)



# Configuring ASM in SF Oracle RAC environments

Before you configure ASM, review the planning guidelines for ASM:

See [“Planning for Oracle ASM over CVM”](#) on page 65.

**Note:** Make sure that you have installed SF Oracle RAC and Oracle RAC before configuring ASM.

Depending on the Oracle RAC version, the ASM home directory varies as follows:

|                                                   |                                                                                                  |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Oracle RAC 10g Release 2/Oracle RAC 11g Release 1 | Use a home directory distinct from the ORACLE_HOME directory.                                    |
| Oracle RAC 11g Release 2                          | The ASM home directory is the same as the Oracle Grid Infrastructure home directory (GRID_HOME). |

## To configure ASM in SF Oracle RAC installations

- 1 Set up ASM.  
See [“Setting up Automatic Storage Management”](#) on page 711.
- 2 For Oracle RAC 10g installations: Copy the VCS IPC libraries.  
See [“Copying the VCS IPC libraries”](#) on page 712.

- 3 Create the required database storage on ASM by creating CVM disk groups and volumes for use as ASM disks.  
See [“Creating database storage on ASM”](#) on page 713.
- 4 Create ASM disk groups and instances.  
See [“Creating ASM disk groups and instances”](#) on page 714.
- 5 Verify the ASM setup.  
See [“Verify the ASM setup”](#) on page 715.
- 6 Create the Oracle database. For instructions, see the Oracle RAC documentation.
- 7 Configure VCS service groups for the Oracle database.  
See [“Configuring VCS service groups for database instances on ASM”](#) on page 715.

## Setting up Automatic Storage Management

Perform this step only for Oracle RAC 10g Release 2 or Oracle RAC 11g Release 1 installations.

Set up ASM in a directory distinct from the Oracle database home (ORACLE\_HOME) directory. A separate ASM home directory enables you to upgrade or uninstall the software independent of the Oracle database home.

The following procedure describes how to set up ASM using the Oracle Universal Installer. Symantec recommends that you set up ASM locally on each node.

### To set up ASM using the Oracle Universal Installer

- 1 Log in as the Oracle user. On the first node, set the DISPLAY variable.

- For Bourne Shell (bash), type:

```
$ DISPLAY=10.20.12.150:0.0 export DISPLAY
```

- For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY 10.20.12.150:0.0
```

- 2 Start the Oracle Universal Installer.

```
$ cd /dvd_mount
```

```
$ ./runInstaller
```

- 3

Enter the following information when prompted by the Oracle Universal Installer:
- Select installation type

Specify home details

Specify Hardware Cluster Installation Mode

Select **Enterprise Edition**.

Enter the full path of the ASM home directory.  
**Note:** Oracle recommends that the ASM home directory be different from the ORACLE\_HOME directory used for installing the Oracle database.  
  
Oracle Home location: The installation destination (\$ORACLE\_HOME). Depending on where you install Oracle binaries, this location is either on shared storage or an identical location on each of the local cluster nodes.

Select **Cluster Installation**.  
Select the nodes on which you want to install the Oracle RAC database software.
- 4

Click **Next**.  
  
The Oracle Universal Installer runs a product-specific prerequisite check. Any items that are flagged must be manually checked and configured.
- 5

On the **Select Configuration Option** screen, select the option **Install database software only**.
- Note: Do not select the option **Configure Automatic Storage Management (ASM)**. Symantec recommends that you configure ASM later.
- 6

Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the installation.
- 7

Run the root.sh script as prompted by the Oracle Universal Installer.

# cd \$ORACLE\_HOME

# ./root.sh

## Copying the VCS IPC libraries

Perform this step only for Oracle RAC 10g installations and if you are using VCS IPC.



Perform the steps in the procedure on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, copy the libraries on one node only. Use the mount command to check that the file system containing the Oracle libraries are mounted.

#### To copy VCS IPC libraries

- 1 Log in as the Oracle user.
- 2 Change to the `$ORACLE_HOME/lib` directory:  
  

```
$ cd $ORACLE_HOME/lib
```
- 3 Back up Oracle's `libskgxp10` library:  
  

```
$ mv libskgxp10.so libskgxp10.so.bkp
```
- 4 Replace the file `$ORACLE_HOME/lib/libskgxp10.so` with the Veritas VCS IPC library:  
  

```
$ cp /opt/VRTSvcs/rac/lib/libskgxp10_ver25_64.so libskgxp10.so
```

## Creating database storage on ASM

This step creates the database storage on ASM using CVM volumes. To create the storage for Oracle databases on ASM, first create the required CVM disk groups and volumes. Then, use these CVM volumes to create ASM disk groups for storing the database files.

#### To create database storage on ASM

- 1 Log in as the root user to the CVM master:  
To determine the CVM master:  
  

```
# vxdctl -c mode
```
- 2 Initialize the disks as VxVM disks:  
  

```
# vxdisksetup -i c1t1d0s2
```

### 3 Create the CVM disk group and volume:

```
# vxdg -s init ora_asm_dg clt1d0s2

# vxassist -g ora_asm_dg make ora_asm_vol 2000M
```

### 4 Set the permissions for the Oracle user on the volume:

```
# vxedit -g ora_asm_dg \
set group=dba user=oracle mode=660 ora_asm_vol
```

Configure ASM using DBCA (Oracle RAC 10g Release 2/Oracle RAC 11g Release 1), ASMCA (Oracle RAC 11g Release 2), or OEM.

## Creating ASM disk groups and instances

Depending on the Oracle RAC version, use one of the following ways to create ASM disk groups:

|                                                   |                                         |
|---------------------------------------------------|-----------------------------------------|
| Oracle RAC 10g Release 2/Oracle RAC 11g Release 1 | Database Configuration Assistant (DBCA) |
|                                                   | Oracle Enterprise Manager (OEM)         |
|                                                   | Manual                                  |
| Oracle RAC 10g Release 2                          | ASM Configuration Assistant (ASMCA)     |
|                                                   | Manual                                  |

The following tasks are performed:

- The ASM disk group and instance is created.
- The ASM instance is started and the disk group is mounted on all nodes in the cluster. The default ASM instance name is +ASM $n$  where  $n$  is the instance number depending on the number of ASM instances.

---

**Note:** For ASM instances that use a pfile, if you restart a node, make sure that the underlying volumes are available before the ASM disk group is mounted. Then, update the ASM init.ora parameter with the full path that contains the VxVM volumes created for ASM.

For ASM instances that use an spfile, the parameter is updated automatically by the Oracle Configuration Assistant.

In either case, dependencies must be managed manually. Symantec recommends the use of the ASMDG agent for easier management of ASM disk groups.

---

SF Oracle RAC requires the following settings when you create ASM disk groups and instances:

Type of parameter file for the server parameter file (spfile)  
ASM instance

Disk Discovery Path      Enter the full path that contains the VxVM volumes created for ASM, for example  
/dev/vx/rdisk/ora\_asm\_dg/ora\_asm\_vol.

Redundancy      Select the option **External**. Mirroring is performed by CVM.

Select member disks      Select the VxVM disks you want to use; You may need to select the **Force** checkbox next to the ASM disks you want to use if the disk group creation fails.

For detailed instructions, see the Oracle documentation.

## Verify the ASM setup

Verify that the database services for ASM are up and running after the installation.

### To verify the ASM installation

- 1 Change to the Oracle Clusterware/Grid Infrastructure home directory:

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# cd $CRS_HOME/bin
```

For Oracle RAC 11g Release 2:

```
# cd $GRID_HOME/bin
```

- 2 Verify the status of ASM on all the nodes in the cluster:

```
# ./srvctl status asm
```

ASM instance +ASM1 is running on node galaxy

The sample output shows that there is one ASM instance running on the local node.

## Configuring VCS service groups for database instances on ASM

This section describes how to configure the Oracle service group using the CLI for databases on ASM.

For sample service group illustration:

See “[sfrac05\\_main.cf file](#)” on page 681.

For a sample configuration file describing the configuration, see the file `sfrac05_main.cf` in the directory `/etc/VRTSvcs/conf/sample_rac/`.

The following procedure assumes that you have created the database.

**To configure the Oracle service group using the CLI**

- 1 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the service group to the VCS configuration:

```
# hagrps -add dbgrp
```

- 3 Modify the attributes of the service group:

```
# hagrps -modify dbgrp Parallel 1

# hagrps -modify dbgrp SystemList galaxy 0 nebula 1

# hagrps -modify dbgrp AutoStartList galaxy nebula
```

- 4 Add the CVMVolDg resource for the service group:

```
# hares -add oradata_voldg CVMVolDg dbgrp
```

- 5 Modify the attributes of the CVMVolDg resource for the service group:

```
# hares -modify oradata_voldg CVMGroup ora_asm_dg
# hares -modify oradata_voldg CVMActivation sw
# hares -modify oradata_voldg CVMVolume ora_asm_vol
```

- 6 Add the ASMDG resource for the service group:

```
# hares -add asmdg ASMDG dbgrp
```

## 7 Modify the attributes of the ASMDG resource for the service group.

---

**Note:** The \$ASM\_HOME variable refers to the ASM home directory.

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1, the ASM home directory is distinct from the ORACLE\_HOME directory.

---

For Oracle RAC 10g Release 2/Oracle RAC 11g Release 1:

```
# hares -modify asmdg DiskGroups ASM_RAC_DG
# hares -modify asmdg Home "$ASM_HOME"
# hares -local asmdg Sid
# hares -modify asmdg Sid "+ASM1" -sys galaxy
# hares -modify asmdg Sid "+ASM2" -sys nebula
# hares -modify asmdg Owner oracle
```

For Oracle RAC 11g Release 2:

```
# hares -modify asmdg DiskGroups ASM_RAC_DG
# hares -modify asmdg Home "$ASM_HOME"
# hares -local asmdg Sid
# hares -modify asmdg Sid "+ASM1" -sys galaxy
# hares -modify asmdg Sid "+ASM2" -sys nebula
# hares -modify asmdg Owner grid
```

## 8 Add the Oracle RAC database instance to the service group:

```
# hares -add asmdb Oracle dbgrp
```

## 9 Modify the attributes of the Oracle resource for the service group:

```
# hares -modify asmdb Owner oracle
# hares -local asmdb Sid
# hares -modify asmdb Sid oradb1 -sys galaxy
# hares -modify asmdb Sid oradb2 -sys nebula
# hares -modify asmdb Home "$ORACLE_HOME"
# hares -modify asmdb StartUpOpt SRVCTLSTART
# hares -modify asmdb ShutDownOpt SRVCTLSTOP
```

## 10 Set the dependencies between the ASMDG resource and the CVMVolDg resource for the Oracle service group:

```
# hares -link asmdg oradata_voldg
```

- 11 Set the dependencies between the Oracle resource and the ASMDG resource for the Oracle service group:

```
# hares -link asmdb asmdg
```

- 12 Create an online local firm dependency between the dbgrp service group and the cvm service group:

```
# hagrps -link dbgrp cvm online local firm
```

- 13 Enable the Oracle service group:

```
# hagrps -enableresources dbgrp
```

- 14 Change the cluster configuration to the read-only mode:

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

- 15 Bring the Oracle service group online on all the nodes:

```
# hagrps -online dbgrp -any
```

## Sample configuration file Veritas CVM and ASM main.cf file

A sample configuration file (sfrac05\_main.cf) illustrating ASM configuration in SF Oracle RAC is available at `/etc/VRTSvcs/conf/sample_rac/`.

For an illustration of the service group configuration:

See “[sfrac05\\_main.cf file](#)” on page 681.

# Creating a test database

This appendix includes the following topics:

- [About creating a test database](#)
- [Creating a database for Oracle](#)

## About creating a test database

A test database can be created and used for both testing and troubleshooting purposes.

The following optional procedures describe the methods for creating a test database.

## Creating a database for Oracle

Before you begin to create the database, ensure that the following prerequisites are met:

- The CRS daemons must be running.  
To verify the status of Oracle Clusterware, type the following command:  

```
# $CRS_HOME/bin/crs_stat
```
- All private IP addresses on each node must be up.  
Use the `ping` command to verify that all private IP addresses on each node are up.

Refer to your Oracle documentation for instructions on how to install the Oracle database.

You can create the database on one of the following types of storage:

- Shared raw volume  
See “[Creating the database storage on raw volumes](#)” on page 720.
- Cluster File System (CFS)  
See “[Creating the database storage on CFS](#)” on page 721.
- ASM  
See “[Creating database storage on ASM](#)” on page 713.

## Creating the database storage on raw volumes

You can create the database storage on shared raw volume.

### To create the database storage on shared raw volumes

- 1 Log in as root user.
- 2 On the master node, create a shared disk group:

```
# vxdg -s init oradatadg Disk_1
```

- 3 Create a volume in the shared group for each of the required tablespaces.  
Refer to the Oracle documentation to determine the tablespace requirements.

For example, type:

```
# vxassist -g oradatadg make VRT_galaxy 1000M
# vxassist -g oradatadg make VRT_spfile1 10M
.
.
```

- 4 Define the access mode and permissions for the volumes storing the Oracle data.

For each volume listed in \$ORACLE\_HOME/raw\_config, use the `vxedit(1M)` command:

```
# vxedit -g disk_group set group=group user=user mode=660 volume
```

For example:

```
# vxedit -g oradatadg set group=dba user=oracle mode=660 \
VRT_galaxy
```

In this example, `VRT_galaxy` is the name of one of the volumes. Repeat the command to define access mode and permissions for each volume in the `oradatadg`.

- 5 Create the database using the Oracle documentation.



## Creating the database storage on CFS

If you plan to use a cluster file system to store the Oracle database, use the following procedure to create the file system.

### To create the database storage on CFS

- 1 Create a disk group (for example, oradatadg):

```
# vxdg -s init oradatadg Disk_1
```

- 2 Create a single shared volume (for example, oradatavol) that is large enough to contain a file system for all tablespaces.

For example, assuming 6.8 GB is required for database storage, type:

```
# vxassist -g oradatadg make oradatavol 6800M
```

- 3 Start the volume in the disk group:

```
# vxvol -g oradatadg startall
```

- 4 Create a VxFS file system in this volume. From one node, type:

```
# mkfs -F vxfs /dev/vx/rdisk/oradatadg/oradatavol
```

- 5 Create a mount point for the shared file system:

```
# mkdir /oradata
```

- 6 From the same node, mount the file system:

```
# mount -F vxfs -o cluster /dev/vx/dsk/oradatadg/oradatavol \
/oradata
```

- 7 Set "oracle" as the owner of the file system, and set "755" as the permissions:

```
# chown oracle:oinstall /oradata
# chmod 755 /oradata
```

- 8 On the other node(s), complete step 5 and step 6.

- 9 Create the database using the Oracle documentation.



# High availability agent information

This appendix includes the following topics:

- [About agents](#)
- [CVMCluster agent](#)
- [CVMVxconfigd agent](#)
- [CVMVolDg agent](#)
- [CFSMount agent](#)
- [CFSfsckd agent](#)
- [PrivNIC agent](#)
- [MultiPrivNIC agent](#)
- [CSSD agent](#)
- [VCS agents for Oracle](#)
- [CRSResource agent](#)

## About agents

An agent is defined as a process that starts, stops, and monitors all configured resources of a type, and reports their status to Veritas Cluster Server (VCS). Agents have both entry points and attributes. Entry points are also known as agent functions and are referred to as "agent functions" throughout the document.

Attributes contain data about the agent. An attribute has a definition and a value. You change attribute values to configure resources, which are defined as the individual components that work together to provide application services to the public network. For example, a resource may be a physical component such as a disk or a network interface card, a software component such as Oracle or a Web server, or a configuration component such as an IP address or mounted file system.

Attributes are either optional or required, although sometimes the attributes that are optional in one configuration may be required in other configurations. Many optional attributes have predefined or default values, which you should change as required. A variety of internal use only attributes also exist. Do not modify these attributes—modifying them can lead to significant problems for your clusters. Attributes have type and dimension. Some attribute values can accept numbers, others can accept alphanumeric values or groups of alphanumeric values, while others are simple boolean on/off values.

The entry points and attributes for each SF Oracle RAC agent are described in this appendix.

## VCS agents included within SF Oracle RAC

SF Oracle RAC includes the following VCS agents:

- CVMCluster agent
- CVMVxconfigd agent
- CVMVolDg agent
- CFSSMount agent

An SF Oracle RAC installation automatically configures the CVMCluster resource and the CVMVxconfigd resource.

You must configure the CVMVolDg agent for each shared disk group. If the database uses cluster file systems, configure the CFSSMount agent for each volume in the disk group.

Use the information in this appendix about the entry points and attributes of the listed agents to make necessary configuration changes. For information on how to modify the VCS configuration:

See the *Veritas Cluster Server Administrator's Guide*

## VCS agents for Oracle included within SF Oracle RAC

SF Oracle RAC includes the following VCS agents for Oracle:

- Oracle agent

The Oracle agent monitors the database processes.

- **Netlsnr agent**

The Netlsnr agent brings the listener services online, monitors their status, and takes them offline.

- **PrivNIC agent**

The PrivNIC agent provides high availability to a single private IP address across LLT Ethernet interfaces for a system.

- **MultiPrivNIC agent**

The MultiPrivNIC agent provides high availability to multiple private IP addresses across LLT Ethernet interfaces for a system.

- **CSSD agent**

The CSSD (Cluster Synchronization Services daemon) agent provides the resources to monitor Oracle Clusterware. The agent ensures that the dependency of cssd on the OCR and the VOTE resources and the PrivNIC (optional) resource are satisfied.

- **ASMDG agent**

The ASMDG agent mounts and unmounts the ASM disk groups onto an ASM instance.

See the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide*

## CVMCluster agent

The CVMCluster agent controls system membership on the cluster port that is associated with Veritas Volume Manager (VxVM).

The CVMCluster agent performs the following functions:

- Joins a node to the CVM cluster port.
- Removes a node from the CVM cluster port.
- Monitors the node's cluster membership state.

## Entry points for CVMCluster agent

[Table G-1](#) describes the entry points used by the CVMCluster agent.

**Table G-1** CVMCluster agent entry points

| Entry Point | Description                                                                                                                               |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Online      | Joins a node to the CVM cluster port. Enables the Volume Manager cluster functionality by automatically importing the shared disk groups. |
| Offline     | Removes a node from the CVM cluster port.                                                                                                 |
| Monitor     | Monitors the node's CVM cluster membership state.                                                                                         |

Attribute definition for CVMCluster agent

Table G-2 describes the user-modifiable attributes of the CVMCluster resource type.

**Table G-2** CVMCluster agent attributes

| Attribute    | Dimension          | Description                                                                                                   |
|--------------|--------------------|---------------------------------------------------------------------------------------------------------------|
| CVMClustName | string-scalar      | Name of the cluster.                                                                                          |
| CVMNodeAddr  | string-association | List of host names and IP addresses.                                                                          |
| CVMNodeId    | string-association | Associative list. The first part names the system; the second part contains the LLT ID number for the system. |
| CVMTransport | string-scalar      | Specifies the cluster messaging mechanism.<br>Default = gab<br><b>Note:</b> Do not change this value.         |
| PortConfigd  | integer-scalar     | The port number that is used by CVM for vxconfigd-level communication.                                        |
| PortKmsgd    | integer-scalar     | The port number that is used by CVM for kernel-level communication.                                           |
| CVMTimeout   | integer-scalar     | Timeout in seconds used for CVM cluster reconfiguration. Default = 200                                        |

CVMCluster agent type definition

The following type definition is included in the file, `CVMTypes.cf`:

```

type CVMCluster (
    static int InfoTimeout = 0
    static int NumThreads = 1
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMTransport, CVMClustName,
        CVMNodeAddr, CVMNodeId, PortConfigd, PortKmsgd,
        CVMTimeout }
    NameRule = ""
    str CVMClustName
    str CVMNodeAddr{}
    str CVMNodeId{}
    str CVMTransport
    int PortConfigd
    int PortKmsgd
    int CVMTimeout
)

```

---

**Note:** The attributes `CVMNodeAddr`, `PortConfigd`, and `PortKmsgd` are not used in an SF Oracle RAC environment. GAB, the required cluster communication messaging mechanism, does not use them.

---

## CVMCluster agent sample configuration

The following is an example definition for the CVMCluster service group:

```

CVMCluster cvm_clus (
    Critical = 0
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

```

## CVMVxconfigd agent

The CVMVxconfigd agent starts and monitors the vxconfigd daemon. The vxconfigd daemon maintains disk and disk group configurations, communicates configuration changes to the kernel, and modifies the configuration information that is stored on disks. CVMVxconfigd must be present in the CVM service group.

The CVMVxconfigd agent is an OnOnly agent; the agent starts the resource when the cluster starts up and VCS restarts the resource when necessary. The Operations attribute specifies these default aspects of startup.

Symantec recommends starting the vxconfigd daemon with the `syslog` option, which enables logging of debug messages. Note that the SF Oracle RAC installation configures the `syslog` option for the CVMVxconfigd agent.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

## Entry points for CVMVxconfigd agent

[Table G-3](#) describes the entry points for the CVMVxconfigd agent.

**Table G-3** CVMVxconfigd entry points

| Entry Point         | Description                                                                                                                                                                                                                                                                               |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online              | Starts the vxconfigd daemon                                                                                                                                                                                                                                                               |
| Offline             | N/A                                                                                                                                                                                                                                                                                       |
| Monitor             | Monitors whether vxconfigd daemon is running                                                                                                                                                                                                                                              |
| imf_init            | Initializes the agent to interface with the AMF kernel module. This function runs when the agent starts up.                                                                                                                                                                               |
| imf_getnotification | Gets notification about the vxconfigd process state. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification. If the vxconfigd process fails, the function initiates a traditional CVMVxconfigd monitor entry point. |
| imf_register        | Registers or unregisters the vxconfigd process id (pid) with the AMF kernel module. This function runs after the resource goes into steady online state.                                                                                                                                  |

## Attribute definition for CVMVxconfigd agent

[Table G-4](#) describes the modifiable attributes of the CVMVxconfigd resource type.



**Table G-4**
CVMVxconfigd agent attribute

| Attribute        | Dimension | Description                                                                                                                                                     |
|------------------|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CVMVxconfigdArgs | keylist   | <p>List of the arguments that are sent to the <code>online</code> entry point.</p> <p>Symantec recommends always specifying the <code>syslog</code> option.</p> |

Table G-4 CVMVxconfigd agent attribute (continued)

| Attribute | Dimension           | Description |
|-----------|---------------------|-------------|
| IMF       | integer-association |             |

**Table G-4** CVMVxconfigd agent attribute (*continued*)

| Attribute | Dimension | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           |           | <p>This resource-type level attribute determines whether the CVMVxconfigd agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> <li>■ <b>Mode:</b> Define this attribute to enable or disable intelligent resource monitoring.</li> </ul> <p>Valid values are as follows:</p> <ul style="list-style-type: none"> <li>■ 0—Does not perform intelligent resource monitoring</li> <li>■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources</li> </ul> <p>Default: 0</p> <ul style="list-style-type: none"> <li>■ <b>MonitorFreq:</b> This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer.</li> </ul> <p>Default: 1</p> <p>You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"> <li>■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources</li> <li>■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources</li> </ul> <ul style="list-style-type: none"> <li>■ <b>RegisterRetryLimit:</b> If you enable intelligent resource monitoring, the agent invokes the oracle_imf_register agent function to register the resource with the AMF kernel driver. The value of the</li> </ul> |

Table G-4 CVMVxconfigd agent attribute (continued)

| Attribute | Dimension | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           |           | <p>RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes.</p> <p>Default: 3.</p> <p>For more details of IMF attribute for the agent type, refer to the <i>Veritas Cluster Server Administrator's Guide</i>.</p> |

## CVMVxconfigd agent type definition

The following type definition is included in the CVMTypes.cf file:

```
type CVMVxconfigd (  
    static int FaultOnMonitorTimeouts = 2  
    static int RestartLimit = 5  
    static str ArgList[] { CVMVxconfigdArgs }  
    static str Operations = OnOnly  
    keylist CVMVxconfigdArgs  
)
```

## CVMVxconfigd agent sample configuration

The following is an example definition for the CVMVxconfigd resource in the CVM service group:

```
CVMVxconfigd cvm_vxconfigd (  
    Critical = 0  
    CVMVxconfigdArgs = { syslog }  
)
```

# CVMVolDg agent

The CVMVolDg agent manages the CVM disk groups and CVM volumes and volume sets within the disk groups by performing the following functions:

- Imports the shared disk group from the CVM master node
- Starts the volumes and volume sets in the disk group
- Monitors the disk group, volumes, and volume sets
- Optionally, deports the disk group when the dependent applications are taken offline. The agent deports the disk group only if the appropriate attribute is set.

Configure the CVMVolDg agent for each disk group used by a Oracle service group. A disk group must be configured to only one Oracle service group. If cluster file systems are used for the database, configure the CFMount agent for each volume or volume set in the disk group.

## Entry points for CVMVolDg agent

[Table G-5](#) describes the entry points used by the CVMVolDg agent.

**Table G-5** CVMVolDg agent entry points

| Entry Point | Description                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online      | <p>Imports the shared disk group from the CVM master node, if the disk group is not already imported.</p> <p>Starts all volumes and volume sets in the shared disk group specified by the CVMVolume attribute.</p> <p>Sets the disk group activation mode to shared-write if the value of the CVMActivation attribute is sw. You can set the activation mode on both slave and master systems.</p> |
| Offline     | <p>Removes the temporary files created by the online entry point.</p> <p>If the CVMDeportOnOffline attribute is set to 1 and if the shared disk group does not contain open volumes on any node in the cluster, the disk group is deported from the CVM master node.</p>                                                                                                                           |

Table G-5 CVMVolDg agent entry points (continued)

| Entry Point | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor     | <p>Determines whether the disk group, the volumes, and the volume sets are online.</p> <p>The agent takes a volume set offline if the file system metadata volume of a volume set is discovered to be offline in a monitor cycle.</p> <p><b>Note:</b> If the CFSSMount resource goes offline and the file system on the volume set is unmounted, the agent retains the online state of the volume set even if the file system metadata volume in the volume set is offline. This is because the CVMVolDg agent is unable to determine whether or not the volumes that are offline are metadata volumes.</p> |
| Clean       | Removes the temporary files created by the online entry point.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Attribute definition for CVMVolDg agent

Table G-6 describes the user-modifiable attributes of the CVMVolDg resource type.

Table G-6 CVMVolDg agent attributes

| Attribute                | Dimension      | Description                                                                                                                                                                                                              |
|--------------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CVMDiskGroup (required)  | string-scalar  | Shared disk group name.                                                                                                                                                                                                  |
| CVMVolume (required)     | string-keylist | Name of shared volumes or volume sets. This list is used to check that the volumes or volume sets are in the correct state before allowing the resource to come online, and that the volumes remain in an enabled state. |
| CVMActivation (required) | string-scalar  | Activation mode for the disk group.<br><br>Default = sw (shared-write)<br><br>This is a localized attribute.                                                                                                             |

**Table G-6** CVMVolDg agent attributes (*continued*)

| Attribute                     | Dimension      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CVMVolumeIoTest(optional)     | string-keylist | List of volumes and volume sets that will be periodically polled to test availability. The polling is in the form of 4 KB reads every monitor cycle to a maximum of 10 of the volumes or volume sets in the list. For volume sets, reads are done on a maximum of 10 component volumes in each volume set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| CVMDeportOnOffline (optional) | integer-scalar | <p>Indicates whether or not the shared disk group must be deported when the last online CVMVolDg resource for a disk group is taken offline.</p> <p>The value 1 indicates that the agent will deport the shared disk group from the CVM master node, if not already deported, when the last online CVMVolDg resource for the disk group is taken offline.</p> <p>The value 0 indicates that the agent will not deport the shared disk group when the CVMVolDg resource is taken offline.</p> <p>The default value is set to 0.</p> <p><b>Note:</b> If multiple CVMVolDg resources are configured for a shared disk group, set the value of the attribute to either 1 or 0 for all of the resources.</p> <p>The CVM disk group is deported based on the order in which the CVMVolDg resources are taken offline. If the CVMVolDg resources in the disk group contain a mixed setting of 1 and 0 for the CVMDeportOnOffline attribute, the disk group is deported only if the attribute value is 1 for the last CVMVolDg resource taken offline. If the attribute value is 0 for the last CVMVolDg resource taken offline, the disk group is not deported.</p> <p>The deport operation fails if the shared disk group contains open volumes.</p> |

## CVMVolDg agent type definition

The CVMTypes.cf file includes the CVMVolDg type definition:

```
type CVMVolDg (
    static keylist RegList = { CVMActivation, CVMVolume }
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMDiskGroup, CVMVolume, CVMActivation,
    CVMVolumeIoTest, CVMDGAction, CVMDeportOnOffline }
    str CVMDiskGroup
    str CVMDGAction
    keylist CVMVolume
    str CVMActivation
    keylist CVMVolumeIoTest
    int CVMDeportOnOffline
    temp int voldg_stat
)
```

## CVMVolDg agent sample configuration

Each Oracle service group requires a CVMVolDg resource type to be defined. The following is a sample configuration:

```
CVMVolDg ora_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradata1, oradata2 }
    CVMActivation = sw
)
```

## CFSMount agent

The CFSMount agent brings online, takes offline, and monitors a cluster file system mount point.

The agent executable is located in /opt/VRTSvcs/bin/CFSMount/CFSMountAgent.

The CFSMount type definition is described in the /etc/VRTSvcs/conf/config/CFSTypes.cf file.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.



## Entry points for CFSMount agent

[Table G-7](#) provides the entry points for the CFSMount agent.

**Table G-7** CFSMount agent entry points

| Entry Point         | Description                                                                                                                                                                                                                      |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online              | Mounts a block device in cluster mode.                                                                                                                                                                                           |
| Offline             | Unmounts the file system, forcing unmount if necessary, and sets primary to secondary if necessary.                                                                                                                              |
| Monitor             | Determines if the file system is mounted. Checks mount status using the <code>fsclustadm</code> command.                                                                                                                         |
| Clean               | Generates a null operation for a cluster file system mount.                                                                                                                                                                      |
| imf_init            | Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for the agent. This function runs when the agent starts up.                                                                  |
| imf_getnotification | Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification. |
| imf_register        | Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).                                                       |

## Attribute definition for CFSMount agent

[Table G-8](#) lists user-modifiable attributes of the CFSMount Agent resource type.

**Table G-8** CFSMount Agent attributes

| Attribute   | Dimension      | Description                                                                                         |
|-------------|----------------|-----------------------------------------------------------------------------------------------------|
| MountPoint  | string-scalar  | Directory for the mount point.                                                                      |
| BlockDevice | string-scalar  | Block device for the mount point.                                                                   |
| NodeList    | string-keylist | List of nodes on which to mount. If NodeList is NULL, the agent uses the service group system list. |

**Table G-8** CFSMount Agent attributes *(continued)*

| Attribute | Dimension           | Description |
|-----------|---------------------|-------------|
| IMF       | integer-association |             |

**Table G-8** CFSMount Agent attributes (*continued*)

| Attribute | Dimension | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           |           | <p>Resource-type level attribute that determines whether the CFSMount agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> <li>■ <b>Mode:</b> Define this attribute to enable or disable intelligent resource monitoring.</li> </ul> <p>Valid values are as follows:</p> <ul style="list-style-type: none"> <li>■ 0—Does not perform intelligent resource monitoring</li> <li>■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources</li> <li>■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources</li> <li>■ 3—Performs intelligent resource monitoring for both online and for offline resources</li> </ul> <p>Default: 0</p> <ul style="list-style-type: none"> <li>■ <b>MonitorFreq:</b> This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer.</li> </ul> <p>Default: 1</p> <p>You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"> <li>■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources</li> </ul> |

Table G-8 CFSMount Agent attributes (continued)

| Attribute           | Dimension     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                     |               | <ul style="list-style-type: none"><li>■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources</li><li>■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the oracle_imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3.</li></ul> |
| MountOpt (optional) | string-scalar | <p>Options for the mount command. To create a valid MountOpt attribute string:</p> <ul style="list-style-type: none"><li>■ Use the VxFS type-specific options only.</li><li>■ Do not use the -o flag to specify the VxFS-specific options.</li><li>■ Do not use the -F vxfs file system type option.</li><li>■ Be aware the cluster option is not required.</li><li>■ Specify options in comma-separated list:</li></ul> <pre>ro ro,cluster blkclear,mincache=closesync</pre>                                                                                                                                                                                            |
| Policy (optional)   | string-scalar | <p>List of nodes to assume the primaryship of the cluster file system if the primary fails. If set to NULL or if none of the hosts specified in the list is active when the primary fails, a node is randomly selected from the set of active nodes to assume primaryship.</p>                                                                                                                                                                                                                                                                                                                                                                                           |

CFSMount agent type definition

The CFSTypes.cf file includes the CFSMount agent type definition:

```

type CFSSMount (
    static keylist RegList = { MountOpt, Policy, NodeList, ForceOff, SetPrimary }
    static keylist SupportedActions = { primary }
    static int FaultOnMonitorTimeouts = 1
    static int OnlineWaitLimit = 1
    static str ArgList[] = { MountPoint, BlockDevice, MountOpt, Primary, AMFMountType }
    str MountPoint
    str MountType
    str BlockDevice
    str MountOpt
    keylist NodeList
    keylist Policy
    temp str Primary
    str SetPrimary
    temp str RemountRes
    temp str AMFMountType
    str ForceOff
)

```

## CFSSMount agent sample configuration

Each Oracle service group requires a CFSSMount resource type to be defined:

```

CFSSMount ora_mount (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol1"
    Primary = nebula;
)

```

## CFSfsckd agent

The CFSfsckd agent starts, stops, and monitors the vxfsckd process. The CFSfsckd agent executable is `/opt/VRTSvcs/bin/CFSfsckd/CFSfsckdAgent`. The type definition is in the `/etc/VRTSvcs/conf/config/CFSTypes.cf` file. The configuration is added to the `main.cf` file after running the `cfscscluster config` command.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

## Entry points for CFSfsckd agent

Table G-9 describes the CFSfsckd agent entry points.

| Table G-9 CFSfsckd agent entry points |                                                                                                                                                                                                                                  |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Entry Points                          | Description                                                                                                                                                                                                                      |
| Online                                | Starts the vxfsckd process.                                                                                                                                                                                                      |
| Offline                               | Kills the vxfsckd process.                                                                                                                                                                                                       |
| Monitor                               | Checks whether the vxfsckd process is running.                                                                                                                                                                                   |
| Clean                                 | A null operation for a cluster file system mount.                                                                                                                                                                                |
| imf_init                              | Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for the agent. This function runs when the agent starts up.                                                                  |
| imf_getnotification                   | Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification. |
| imf_register                          | Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).                                                       |

## Attribute definition for CFSfsckd agent

Table G-10 lists user-modifiable attributes of the CFSfsckd Agent resource type.

Table G-10CFSfsckd Agent attributes

| Attribute | Dimension           | Description |
|-----------|---------------------|-------------|
| IMF       | integer-association |             |

Table G-10 CFSfsckd Agent attributes (continued)

| Attribute | Dimension | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           |           | <p>Resource-type level attribute that determines whether the CFSfsckd agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"><li>■ <b>Mode:</b> Define this attribute to enable or disable intelligent resource monitoring.</li></ul> <p>Valid values are as follows:</p> <ul style="list-style-type: none"><li>■ <b>0</b>—Does not perform intelligent resource monitoring</li><li>■ <b>1</b>—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources</li><li>■ <b>2</b>—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources</li><li>■ <b>3</b>—Performs intelligent resource monitoring for both online and for offline resources</li></ul> <p>Default: 0</p> <ul style="list-style-type: none"><li>■ <b>MonitorFreq:</b> This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer.</li></ul> <p>Default: 1</p> <p>You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"><li>■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources</li></ul> |



**Table G-10** CFSfsckd Agent attributes *(continued)*

| Attribute | Dimension | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           |           | <ul style="list-style-type: none"> <li>■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources</li> <li>■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the oracle_imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3.</li> </ul> |

## CFSfsckd agent type definition

The CFSfsckd type definition:

```
type CFSfsckd (
    static int RestartLimit = 1
    str ActivationMode{}
)
```

## CFSfsckd agent sample configuration

This is a sample of CFSfsckd configuration:

```
CFSfsckd vxfsckd (
)
```

# PrivNIC agent

The PrivNIC agent provides high availability to a single private IP address across LLT Ethernet interfaces for a system. Private IP addresses are required by Oracle Clusterware and the Oracle database to provide communication between the cluster nodes.

**Note:** The PrivNIC agent operates over LLT links. To use the agent, the Oracle Clusterware interconnects and the Oracle RAC database communication links must be configured as LLT links.

The PrivNIC agent relies on LLT to monitor the LLT Ethernet interfaces. It queries LLT for the number of visible nodes on each of the LLT Ethernet interfaces.

The PrivNIC agent provides a reliable alternative when operating system limitations prevent you from using NIC bonding to provide increased bandwidth using multiple network interfaces. In the event of a NIC failure or link failure, the agent fails over the private IP address from the failed link to the connected or available LLT link. If the preferred link becomes available, the IP address is failed back to the preferred link.

## Functions of the PrivNIC agent

[Table G-11](#) describes the PrivNIC agent's monitor entry point.

**Note:** Because the resource is persistent, only the monitor entry point is required.

**Table G-11** PrivNIC agent entry point

| Entry Point | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor     | The PrivNIC agent queries LLT to create a list of nodes visible on every LLT network interface. The PrivNIC agent then applies various filters to this list to arrive at a most desired failover decision and calculates a "winner" device on which to configure the IP address. The "winner" device is compared to the currently active device where the IP address is currently configured. If the active and "winner" devices are different, the PrivNIC agent initiates a failover to the "winner" device. |

## Attributes of the PrivNIC agent

[Table G-12](#) describes the user-modifiable attributes of the PrivNIC agent.

**Table G-12** Required attributes for PrivNIC agent

| Attribute | Dimension            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Device    | string - association | <p>Specifies the network interface device as shown by the <code>ifconfig</code> 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>At least one interface device must be specified.</p> <p>Example:</p> <pre>Device@galaxy = {bge1=0, bge2=1, bge3=2} Device@nebula = {bge1=0, bge2=1, bge3=2}</pre> |
| Address   | string-scalar        | <p>The numerical private IP address.</p> <p>Checks are performed to determine if this is a valid IP address.</p> <p>When configuring private IPv4 addresses for Oracle Clusterware, make sure that there are no leading zeroes in any of the octets that comprise the IP address, for example X.X.X.01 or X.X.0X.1 or X.0X.X.1 or 0X.X.X.1.</p> <p>Ensure that the IPv4 addresses have the format as "X.X.X.1".</p> <p>The following is an example of an IPv4 address:</p> <pre>Address = "192.168.12.1"</pre>                                                                                                                                 |
| NetMask   | string - association | <p>The numerical netmask for the private IP address. For example:</p> <pre>Address = "255.255.255.0"</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

### Optional attributes of the PrivNIC agent

**Table G-13** Optional attributes for PrivNIC agent

| Attribute     | Dimension            | Description                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DeviceTag     | string - association | Associates an LLT device "tag" with device. If an LLT device tag (as specified in the <code>/etc/llttab</code> file) differs from the name of the network interface as shown in "ifconfig," then DeviceTag must be specified for that interface.                                                                                                                                                                                  |
| GabPort       | string-scalar        | A single lower-case letter specifying the name of the GAB port to be used for filtering. "o" is the default. NULL disables GAB port filtering.<br><br>Example: <code>GabPort = "b"</code>                                                                                                                                                                                                                                         |
| UseVirtualIP  | integer-scalar       | The default is 0, which specifies that the agent use the physical interface for configuring the private IP address when possible.<br><br>The value 1 specifies that the agent always use the virtual interface for configuring the private IP address.<br><br>The value 2 (which includes the functionality of the value 1) specifies the agent should complain if the private IP address already exists on a physical interface. |
| UseSystemList | integer-scalar       | The value 1 specifies that the agent use the SystemList of the service group to filter the node list. Default = 0.                                                                                                                                                                                                                                                                                                                |
| ExcludeNode   | integer-vector       | List of nodes to permanently exclude from calculation.                                                                                                                                                                                                                                                                                                                                                                            |

### States of the PrivNIC agent

[Table G-14](#) lists the states of the PrivNIC agent.


**Table G-14** States of the PrivNIC agent

| State   | Description                                                                                                                       |
|---------|-----------------------------------------------------------------------------------------------------------------------------------|
| Online  | Indicates that the private IP address is available.                                                                               |
| Unknown | Indicates the inability to determine the state of the resource due to incorrect attribute settings or other configuration issues. |

## Sample service group configuration with the PrivNIC agent

Figure G-1 illustrates a basic service group configuration with the PrivNIC agent.

**Figure G-1** Basic service group configuration with the PrivNIC agent



This configuration shows the dependency of the CSSD resource on the private IP address configured as a PrivNIC resource along with the OCR and voting disk volume and mount point dependencies.

For sample deployment scenarios, see the appendix *SF Oracle RAC deployment scenarios*.

## Type definition of the PrivNIC resource

The following extract shows the type definition of the PrivNIC resource in the PrivNIC.cf file:

```

type PrivNIC (
    static str ArgList[] = { Device, DeviceTag, Address,
                           NetMask, UseVirtualIP, GabPort, UseSystemList,
                           ExcludeNode }
    static int OfflineMonitorInterval = 60
    static int MonitorTimeout = 300
    static str Operations = None

    str Device{}
    str DeviceTag{}
    str Address = ""
    str NetMask = ""
    int UseVirtualIP = 0
    str GabPort = "o"
    int UseSystemList = 0
    int ExcludeNode[]
)

```

## Sample configuration of the PrivNIC resource

The following extract from the configuration file illustrates the configuration of a PrivNIC resource.

```
group cvm (
    SystemList = { galaxy = 0, nebula = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula }
)

PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { qfe0 = 0, qfel = 1}
    Device@nebula = { qfe0 = 0, qfel = 1}
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    NetMask = "255.255.255.0"
)
```

For more examples, see the sample configuration files located at `/etc/VRTSvcS/conf/sample_rac/`.

## MultiPrivNIC agent

The MultiPrivNIC agent provides high availability to multiple private IP addresses across LLT Ethernet interfaces for a system. In the event of a NIC failure or link failure, the agent fails over the private IP address from the failed link to one of the available LLT links. To use the agent, the Oracle Clusterware interconnects and the Oracle RAC database communication links must be configured as LLT links.

The MultiPrivNIC agent is a reliable alternative in scenarios where operating system limitations prevent you from using NIC bonding to provide increased bandwidth and high availability using multiple network interfaces. Even if link aggregation solutions in the form of bonded NICs are implemented, the MultiPrivNIC agent can be used to provide additional protection against the failure of aggregated links by failing over the IP addresses to the available alternate links. These alternate links can be simple NIC interfaces or bonded NICs.

The SF Oracle RAC installer supports the configuration of MultiPrivNIC resources when you prepare to install Oracle RAC.

## Managing high availability of private interconnects

The MultiPrivNIC agent operates over LLT links and relies on LLT to monitor the cluster interfaces. It queries LLT to count and report the number of visible nodes on each of the LLT interfaces. In the event that a preferred link goes down, the IP address is failed over to the private link on which maximum number of peer nodes are visible. If multiple links see maximum nodes and if load-balancing is enabled, the agent considers the current traffic on all devices and calculates a "winner" device with lower traffic. If load balancing is not enabled, the IP address is failed over to the link with the lower network-id.

The failover decision for an IP address is made only when the link hosting the IP address fails. If the preferred link becomes available, the IP address is failed back to the preferred link regardless of whether load-balancing is enabled or disabled.

## Functions of the MultiPrivNIC agent

[Table G-15](#) describes the MultiPrivNIC agent's monitor entry point.

**Note:** Because the resource is persistent, only the monitor entry point is required.

**Table G-15** MultiPrivNIC agent entry point

| Entry point | Description                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor     | <p>The MultiPrivNIC agent queries LLT to create a list of the visible nodes on each LLT network interface.</p> <p>The agent applies various filters to this list and calculates a winner device on which to configure the IP address.</p> <p>If the active device on which the IP address is configured does not see the same number of nodes as the winner device, the agent fails over the IP address to the winner device.</p> |

## Attributes of the MultiPrivNIC agent

[Table G-16](#) below describes the user-modifiable attributes of the MultiPrivNIC resource type.

Table G-16 MultiPrivNIC agent attribute definitions

| Attribute | Dimension          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Device    | string-association | <p>The device attribute specifies the network interface displayed by the <code>ifconfig</code> command and the network ID associated with the interface.</p> <p>The network IDs of the interfaces connected to the same physical network must match. The interfaces specified in the MultiPrivNIC configuration should be exactly the same in name and total number as those which have been used for LLT configuration.</p> <p>An example of the device attribute is as follows:</p> <pre>Device@galaxy = {bge1=0, bge2=1, bge3=2} Device@nebula = {bge1=0, bge2=1, bge3=2}</pre>                                                                                                                                                                                                                                                                                                                     |
| Address   | string-association | <p>The numerical private IP address and its preferred device.</p> <p>The agent verifies that the IP addresses are valid addresses. When you configure private IP addresses, ensure that the addresses do not have a leading 0 in any of the octets that comprise the IP address. The IP address must be in the format "X.X.X.1".</p> <p>In the following example, 0 and 1 indicates the ID of the device on which the IP address is hosted. If the device is unavailable, the agent automatically reconfigures the IP address on one of the available devices. The available device is determined based on the <code>UseLoadBalance</code> attribute setting.</p> <p>An example of the address attribute is as follows:</p> <pre>Address @galaxy = { "192.168.12.1" =0,   "192.168.2.1" =0,   "192.168.3.1" =1 } Address @nebula = { "192.168.12.2" =0,   "192.168.2.2" =0,   "192.168.3.2" =1 }</pre> |



**Table G-16** MultiPrivNIC agent attribute definitions (*continued*)

| Attribute      | Dimension          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Netmask        | string-association | The netmask attribute is the numerical netmask for the private IP address.<br><br>For example, Address = "255.255.255.0".                                                                                                                                                                                                                                                                                                                                                                                                                        |
| UseLoadBalance |                    | A boolean value 0 or 1.<br><br>In the event that the preferred device is unavailable and multiple links see maximum nodes during failover: <ul style="list-style-type: none"><li>■ Setting the attribute to 1 fails over the IP address to the device with lower traffic.</li><li>■ Setting the attribute to 0 fails over the IP address to the device with lower network ID.</li></ul> <b>Note:</b> If the preferred device becomes available, the IP address is failed back to the preferred device regardless of the value of UseLoadBalance. |

## States of the MultiPrivNIC agent

[Table G-17](#) lists the states of the MultiPrivNIC agent.


**Table G-17** States of the MultiPrivNIC agent

| State   | Description                                                                                                                       |
|---------|-----------------------------------------------------------------------------------------------------------------------------------|
| Online  | Indicates that the private IP addresses are available.                                                                            |
| Unknown | Indicates the inability to determine the state of the resource due to incorrect attribute settings or other configuration issues. |

## Sample service group configuration with the MultiPrivNIC agent

[Figure G-2](#) illustrates a sample service group configuration with the MultiPrivNIC agent.

**Figure G-2** Basic service group configuration with MultiPrivNIC agent



The illustrated configuration shows the dependency of the CSSD resource on the private IP address configured as a MultiPrivNIC resource along with the OCR and voting disk volume and mount point dependencies.

For sample deployment scenarios, see the appendix *SF Oracle RAC deployment scenarios*.

## Type definition of the MultiPrivNIC resource

The following extract shows the type definition of the MultiPrivNIC resource in the `MultiPrivNIC.cf` file:

```

type MultiPrivNIC (
    static int MonitorTimeout = 300
    static int OfflineMonitorInterval = 60
    static str ArgList[] = { Device, DeviceTag, Address, NetMask,
                            UseVirtualIP, GabPort, UseSystemList,
                            ExcludeNode, UseLoadBalance }

    static str Operations = None
    str Device{}
    str DeviceTag{}
    str Address{}
    str NetMask
    int UseVirtualIP
    str GabPort = o
    int UseSystemList
    int ExcludeNode[]
    int UseLoadBalance
)

```

## Sample configuration of the MultiPrivNIC resource

The following extract from the configuration file illustrates the configuration of a MultiPrivNIC resource.

```
MultiPrivNIC multi_priv (  
    Critical = 0  
    Device @galaxy = { bge1 = 0, bge2 = 1 }  
    Device @nebula = { bge1 = 0, bge2 = 1 }  
    Address @galaxy = { "192.168.12.1" =0, "192.168.2.1" =0,  
                        "192.168.3.1" =1 }  
    Address @nebula = { "192.168.12.2" =0, "192.168.2.2" =0,  
                        "192.168.3.2" =1 }  
    NetMask = "255.255.255.0"  
    UseLoadBalance = 1  
)
```

For more examples, see the sample configuration files located at  
`/etc/VRTSvcs/conf/sample_rac/`.

In the above example, the interface next to the IP address is the preferred device for that particular IP address. If the same number of nodes are visible on each of the LLT interfaces, the IP addresses are configured on the preferred interfaces.

## CSSD agent

The CSSD agent starts, stops, and monitors Oracle Clusterware/Grid Infrastructure. It ensures that the OCR, the voting disk and the private IP address resources required by Oracle Clusterware/Grid Infrastructure are online before Oracle Clusterware/Grid Infrastructure starts. For this purpose, the cssd resource must be configured as a parent of the resources that manage the OCR, the voting disk, and the private IP address used by Oracle Clusterware/Grid Infrastructure. Using the CSSD agent in SF Oracle RAC installations ensures adequate handling of inter-dependencies and thus prevents the premature startup of Oracle Clusterware/Grid Infrastructure, which causes cluster failures.

For Oracle RAC 10g Release 2 and Oracle RAC 11g Release 1, during system startup, the Oracle Clusterware init scripts invoke the clsinfo script provided by Veritas software. The clsinfo script ensures that the dependent resources of cssd resource configured under VCS are online, thus ensuring that the OCR, the voting disk, and the private IP address resources are online before the cssd resource comes online. After the underlying resources are online, the CSSD agent starts Oracle Clusterware.

For Oracle RAC 11g Release 2, the automatic startup of Oracle Clusterware/Grid Infrastructure must be disabled when the system starts. This prevents the premature startup of Oracle Clusterware/Grid Infrastructure during system startup. Thus, VCS ensures that Oracle Clusterware/Grid Infrastructure is started using the CSSD agent only when all the dependent resources of the cssd resource configured under VCS are online.

During system shutdown, the agent stops Oracle Clusterware/Grid Infrastructure before the OCR and voting disk resources are taken offline by VCS. This ensures that Oracle Clusterware/Grid Infrastructure does not panic the nodes in the cluster due to unavailability of the required resources.

**Note:** It is mandatory to use CSSD agent in SF Oracle RAC installations. You must configure the CSSD agent after installing Oracle Clusterware/Grid Infrastructure.

## Functions of the CSSD agent

Table G-18 describes the functions of the CSSD agent.

Table G-18 CSSD agent functions

| Function | Description                              |
|----------|------------------------------------------|
| Online   | Starts Oracle Clusterware.               |
| Offline  | Stops Oracle Clusterware.                |
| Monitor  | Checks the status of Oracle Clusterware. |

## Attributes of the CSSD agent

Table G-19 lists the required attributes of the CSSD agent:

Table G-19 Required attributes for CSSD resource

| Attribute Name | Required Value                    |
|----------------|-----------------------------------|
| Critical       | 0                                 |
| StartProgram   | /opt/VRTSvcs/rac/bin/cssd-online  |
| StopProgram    | /opt/VRTSvcs/rac/bin/cssd-offline |
| CleanProgram   | /opt/VRTSvcs/rac/bin/cssd-clean   |
| MonitorProgram | /opt/VRTSvcs/rac/bin/cssd-monitor |

## States of the CSSD agent

[Table G-20](#) describes the states of the CSSD agent.


**Table G-20** CSSD agent states

| State   | Description                                                                                                                       |
|---------|-----------------------------------------------------------------------------------------------------------------------------------|
| Online  | Indicates that Oracle Clusterware is running.                                                                                     |
| Offline | Indicates that Oracle Clusterware is not running.                                                                                 |
| Unknown | Indicates the inability to determine the state of the resource due to incorrect attribute settings or other configuration issues. |

## Sample service group configurations with the CSSD agent

[Figure G-3](#) illustrates a basic service group configuration with the CSSD agent.

**Figure G-3** Basic service group configuration with the CSSD agent



In this basic configuration, the OCR and voting disk volumes/mount points and IP addresses are configured under the CSSD resource. This ensures that these resources are online before the CSSD agent starts Oracle Clusterware.

**Note:** Depending on whether the Oracle database is started by Oracle Clusterware or by the VCS agent for Oracle, you must configure the Oracle database mounts such that they are online before the database starts.

For more sample configurations:

See [“Sample configuration files”](#) on page 675.

## Type definition of the CSSD resource

The CSSD agent is an application agent. You can determine the name of the CSSD resource.

The following extract shows the type definition of the CSSD resource in the `types.cf` file.

```
type Application (
    static keylist SupportedActions =
    { "program.vfd", "user.vfd", "cksum.vfd", getcksum }
    static str ArgList[] =
    { User, StartProgram, StopProgram, CleanProgram,
      MonitorProgram, PidFiles, MonitorProcesses }
    static int ContainerOpts{} = { RunInContainer=1, PassCInfo=0 }
    str User
    str StartProgram
    str StopProgram
    str CleanProgram
    str MonitorProgram
    str PidFiles[]
    str MonitorProcesses[]
)
```

## Sample configuration of the CSSD resource

The following extract from the `main.cf` file illustrates a sample CSSD agent configuration:

```
Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline"
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
)
```

## VCS agents for Oracle

The VCS agents for Oracle include the following agents that work together to make Oracle highly available:

- The Oracle agent monitors the Oracle database processes.  
See [“Oracle agent functions”](#) on page 759.

See [“Resource type definition for the Oracle agent”](#) on page 764.

- The Netlsnr agent monitors the listener process.  
See [“Netlsnr agent functions”](#) on page 771.
- The ASMDG agent monitors the Oracle ASM disk groups.  
See [“ASMDG agent functions”](#) on page 776.  
See [“Resource type definition for the ASMDG agent”](#) on page 777.

Refer to the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide* for more details on the agent functions and the resource types.

## Oracle agent functions

The Oracle agent monitors the database processes.

[Table G-21](#) lists the Oracle agent functions.

**Table G-21** Oracle agent functions

| Agent operation | Description                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online          | <p>Starts the Oracle database.</p> <p>The agent uses the default startup option STARTUP_FORCE. For RAC clusters, you must manually change the value of the StartUpOpt attribute to SRVCTLSTART.</p> <p>If you set the option to SRVCTLSTART, the agent uses the following <code>srvctl</code> command to start the Oracle database:</p> <pre>srvctl start database -d database_name</pre> |
| Offline         | <p>Stops the Oracle database.</p> <p>The agent uses the default shutdown option IMMEDIATE. For RAC clusters, you must manually change the value of the ShutDownOpt attribute to SRVCTLSTOP.</p> <p>If you set the option to SRVCTLSTOP, the agent uses the following <code>srvctl</code> command to stop the Oracle database:</p> <pre>srvctl stop database -d database_name</pre>        |
| Monitor         | <p>Verifies the status of the Oracle processes. The Oracle agent provides two levels of monitoring: basic and detail.</p> <p>See <a href="#">“Monitor options for the Oracle agent”</a> on page 760.</p>                                                                                                                                                                                  |

**Table G-21** Oracle agent functions (*continued*)

| Agent operation            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oracle_imf_init            | Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for Oracle agent. This function runs when the agent starts up.                                                                                                                                                                                                                                                                                                                                                                                      |
| oracle_imf_getnotification | Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.                                                                                                                                                                                                                                                                                                                        |
| oracle_imf_register        | Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).                                                                                                                                                                                                                                                                                                                                                                              |
| Clean                      | <p>Forcibly stops the Oracle database.</p> <p>If you set the shutdown option to SRVCTLSTOP, the agent uses the following <code>srvctl</code> command:</p> <pre>srvctl stop database -d database_name</pre> <p>If the process does not respond to the <code>srvctl</code> command, then the agent does the following:</p> <ul style="list-style-type: none"> <li>■ Scans the process table for the processes that are associated with the configured instance</li> <li>■ Kills the processes that are associated with the configured instance</li> </ul> |
| Info                       | <p>Provides the static and dynamic information about the state of the database.</p> <p>See <a href="#">“Info entry point for SF Oracle RAC agent for Oracle”</a> on page 762.</p>                                                                                                                                                                                                                                                                                                                                                                       |
| Action                     | <p>Performs the predefined actions on a resource.</p> <p>See <a href="#">“Action entry point for SF Oracle RAC agent for Oracle”</a> on page 762.</p>                                                                                                                                                                                                                                                                                                                                                                                                   |

## Monitor options for the Oracle agent

The Oracle agent provides two levels of monitoring: basic and detail. By default, the agent does a basic monitoring.

The basic monitoring mode has the following options:

- Process check



■ Health check

The MonitorOption attribute of the Oracle resource determines whether the agent must perform basic monitoring in Process check or Health check mode.

[Table G-22](#) describes the basic monitoring options.

**Table G-22** Basic monitoring options

| Option         | Description                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 0<br>(Default) | <p>Process check</p> <p>The agent scans the process table for the ora_dbw0, ora_smon, ora_pmon, and ora_lgwr processes to verify that Oracle is running.</p> <p>In this mode, the agent also supports intelligent resource monitoring.</p>                                                                                                                                                   |
| 1              | <p>Health check (supported on Oracle 10g and later)</p> <p>The agent uses the Health Check APIs from Oracle to monitor the SGA and retrieve the information about the instance.</p> <p>If you want to use the Oracle agent's intentional offline functionality, you must enable Health check monitoring.</p> <p>The agent does not support intelligent resource monitoring in this mode.</p> |

Review the following considerations if you want to configure basic monitoring:

- Within a failover service group, when the administrator online's an Oracle resource on a node and if the Oracle instance is online on any other node within a cluster, the instance would come up. However, the database does not get mounted. In such circumstances, this failure is detected only by health check monitoring option of basic monitoring or detail monitoring. Detail monitoring updates the database table after detecting a failure whereas health check monitoring does not.
- If health check monitoring option of basic monitoring or detail monitoring is not configured, then such a conflict would go undetected.

In the detail monitoring mode, the agent performs a transaction on a test table in the database to ensure that Oracle functions properly. The agent uses this test table for internal purposes. Symantec recommends that you do not perform any other transaction on the test table. The DetailMonitor attribute of the Oracle resource determines whether the agent must perform detail monitoring.

## Info entry point for SF Oracle RAC agent for Oracle

The Veritas Cluster Server agent for Oracle supports the Info entry point, which provides static and dynamic information about the state of the database.

To invoke the Info entry point, type the following command:

```
# hares -value resource ResourceInfo [system] \
  [-clus cluster | -localclus]
```

The entry point retrieves the following static information:

- |                |              |                |
|----------------|--------------|----------------|
| ■ Version      | ■ InstanceNo | ■ InstanceName |
| ■ DatabaseName | ■ HostName   | ■ StartupTime  |
| ■ Parallel     | ■ Thread     | ■ InstanceRole |

The entry point retrieves the following dynamic information:

- |                            |                       |                  |
|----------------------------|-----------------------|------------------|
| ■ InstanceStatus           | ■ Logins              | ■ OpenMode       |
| ■ LogMode                  | ■ ShutdownPending     | ■ DatabaseStatus |
| ■ Shared Pool Percent free | ■ Buffer Hits Percent |                  |

You can add additional attributes by adding sql statements to the file /opt/VRTSagents/ha/bin/Oracle/resinfo.sql. For example:

```
select 'static:HostName:' || host_name from v$instance;
select 'dynamic:ShutdownPending:' || shutdown_pending from
v$instance;
```

The format of the selected record must be as follows:

```
attribute_type:userkey_name:userkey_value
```

The variable *attribute\_type* can take the value static and/or dynamic.

## Action entry point for SF Oracle RAC agent for Oracle

The Veritas Cluster Server agent for Oracle supports the Action entry point, which enables you to perform predefined actions on a resource.

To perform an action on a resource, type the following command:

```
# hares -action res token [-actionargs arg1 ...] \
  [-sys system] [-clus cluster]
```

You can also add custom actions for the agent.

For further information, refer to the *Veritas Cluster Server Agent Developer's Guide*.

See [Table G-24](#) on page 763.

[Table G-23](#) describes the agent's predefined actions.

**Table G-23** Predefined agent actions

| Action                  | Description                                                                                                                |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------|
| VRTS_GetInstanceName    | Retrieves the name of the configured instance. You can use this option for the Oracle and the Netlsnr resources.           |
| VRTS_GetRunningServices | Retrieves the list of processes that the agent monitors. You can use this option for the Oracle and the Netlsnr resources. |
| DBRestrict              | Changes the database session to enable the RESTRICTED mode.                                                                |
| DBUndoRestrict          | Changes the database session to disable the RESTRICTED mode.                                                               |
| DBSuspend               | Suspends a database.                                                                                                       |
| DBResume                | Resumes a suspended database.                                                                                              |
| DBTbspBackup            | Backs up a tablespace; <code>actionargs</code> contains name of the tablespace to be backed up.                            |

[Table G-24](#) lists the virtual fire drill actions of the Veritas Cluster Server agent for Oracle lets you run infrastructure checks and fix specific errors.

**Table G-24** Predefined virtual fire drill actions

| Virtual fire drill action | Description                                        |
|---------------------------|----------------------------------------------------|
| getid<br>(Oracle agent)   | Verifies that the Oracle Owner exists on the node. |

Table G-24      Predefined virtual fire drill actions (continued)

| Virtual fire drill action       | Description                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| home.vfd<br>(Oracle agent)      | Verifies the following: <ul style="list-style-type: none"><li>■ ORACLE_HOME is mounted on the node and corresponding entry is in the fstab.<br/>If the ORACLE_HOME is not mounted, the action entry point checks if any other resource has already mounted ORACLE_HOME.</li><li>■ Pfile is provided and it exists on the node.</li><li>■ Password file from \$ORACLE_HOME/dbs/orapw[SID] is present.</li></ul> |
| owner.vfd<br>(Oracle agent)     | Verifies the uid and gid of the Oracle Owner attribute. Checks if uid and gid of Owner attribute is the same on the node where the Oracle resource is currently ONLINE.                                                                                                                                                                                                                                        |
| pfile.vfd<br>(Oracle agent)     | Checks for the presence of pfile or spfile on the local disk. If both pfile and spfile are not present, the agent function exits. If the Oracle resource is online in the cluster, the agent function logs a message that the spfile must be on the shared storage because the Oracle resource is online.                                                                                                      |
| tnsadmin.vfd<br>(Netlsnr agent) | Checks if listener.ora file is present. If the listener.ora file is not present, it checks if ORACLE_HOME is mounted and displays appropriate messages.                                                                                                                                                                                                                                                        |

## Resource type definition for the Oracle agent

The Oracle agent of the Veritas Cluster Server agent for Oracle is represented by the Oracle resource type in SF Oracle RAC.

```
type Oracle (  
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/Oracle"  
    static keylist SupportedActions = { VRTS_GetInstanceName,  
        VRTS_GetRunningServices, DBRestrict, DBUndoRestrict, DBResume,  
        DBSuspend, DBTbspBackup, "home.vfd", "owner.vfd", "getid",  
        "pfile.vfd" }  
    static str ArgList[] = { Sid, Owner, Home, Pfile, StartUpOpt,  
        ShutDownOpt, DBAUser, DBAPword, EnvFile, AutoEndBkup,  
        DetailMonitor, User, Pword, Table, MonScript, AgentDebug,  
        Encoding, MonitorOption }  
    static str IMFRegList[] = { Home, Owner, Sid, MonitorOption }  
    static int ContainerOpts{} = { RunInContainer=1, PassCInfo=0 }  
    str Sid
```

```

str Owner
str Home
str Pfile
str StartUpOpt = STARTUP_FORCE
str ShutDownOpt = IMMEDIATE
str DBAUser
str DBAPword
str EnvFile
boolean AutoEndBkup = 1
int DetailMonitor = 0
str MonScript = "./bin/Oracle/SqlTest.pl"
str User
str Pword
str Table
boolean AgentDebug = 0
str Encoding
int MonitorOption = 0
static boolean IntentionalOffline = 0
)

```

## Attribute definition for the Oracle agent

Review the description of the Oracle agent attributes. The agent attributes are classified as required, optional, and internal.

[Table G-25](#) lists the required attributes. You must assign values to the required attributes.

**Table G-25** Required attributes for Oracle agent

| Required attributes | Type and dimension | Definition                                                                                                                                                                                     |
|---------------------|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Sid                 | string-scalar      | The variable \$ORACLE_SID that represents the Oracle instance. The Sid is considered case-sensitive by the Oracle agent and by the Oracle database server.                                     |
| Owner               | string-scalar      | The Oracle user who has privileges to start or stop the database instance.<br>The agent also supports LDAP users as Oracle user.                                                               |
| Home                | string-scalar      | The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home.<br><b>Note:</b> Do not append a slash (/) at the end of the path. |

Table G-26 lists the optional attributes for Oracle agent. You can configure the optional attributes if necessary.

Table G-26            Optional attributes for Oracle agent

| Optional Attributes | Type and Dimension | Definition                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DBAUser             | string-scalar      | The database user who has sysdba privileges to start or stop the database.                                                                                                                                                                                                                                                                                                                                        |
| DBAPword            | string-scalar      | Encrypted password for DBAUser.<br><br>Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the SF Oracle RAC Encrypt utility.                                                                                                                                                                                                                                     |
| StartUpOpt          | string-scalar      | Startup options for the Oracle instance. This attribute can take the following values:<br><br>■ STARTUP<br>■ STARTUP_FORCE<br>■ RESTRICTED<br>■ RECOVERDB<br>■ SRVCTLSTART<br>■ CUSTOM<br><br>Default is STARTUP_FORCE.                                                                                                                                                                                           |
| ShutDownOpt         | string-scalar      | Shut down options for the Oracle instance. This attribute can take the following values:<br><br>■ IMMEDIATE<br>■ TRANSACTIONAL<br>■ SRVCTLSTOP<br>■ CUSTOM<br><br>Default is IMMEDIATE.                                                                                                                                                                                                                           |
| EnvFile             | string-scalar      | The full path name of the file that is sourced by the entry point scripts. This file contains the environment variables set by the user for the Oracle database server environment such as LD_LIBRARY_PATH, NLS_DATE_FORMAT, and so on.<br><br>The syntax for the contents of the file depends on the login shell of Owner. File must be readable by Owner. The file must not contain any prompts for user input. |

**Table G-26** Optional attributes for Oracle agent (*continued*)

| Optional Attributes | Type and Dimension | Definition                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Pfile               | string-scalar      | <p>The name of the initialization parameter file with the complete path of the startup profile.</p> <p>You can also use the server parameter file. Create a one-line text initialization parameter file that contains only the SPFILE parameter. See the Oracle documentation for more information.</p>                                                                                                                       |
| AutoEndBkup         | integer-scalar     | <p>Setting the AutoEndBkup attribute to a non-zero value takes the datafiles in the database out of the backup mode, during Online.</p> <p>Default = 1</p>                                                                                                                                                                                                                                                                    |
| MonitorOption       | integer-scalar     | <p>Monitor options for the Oracle instance. This attribute can take values 0 or 1.</p> <ul style="list-style-type: none"><li>■ 0—Process check monitoring (recommended)</li><li>■ 1—Health check monitoring</li></ul> <p>The agent supports intelligent resource monitoring only when this attribute value is set to 0.</p> <p>Default = 0</p> <p>See <a href="#">“Monitor options for the Oracle agent”</a> on page 760.</p> |

**Table G-26** Optional attributes for Oracle agent (*continued*)

| Optional Attributes | Type and Dimension  | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IMF                 | integer-association | <p>This resource-type level attribute determines whether the Oracle agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> <li>■ <b>Mode:</b> Define this attribute to enable or disable intelligent resource monitoring.</li> </ul> <p>Valid values are as follows:</p> <ul style="list-style-type: none"> <li>■ 0—Does not perform intelligent resource monitoring</li> <li>■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources</li> <li>■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources</li> <li>■ 3—Performs intelligent resource monitoring for both online and for offline resources</li> </ul> <p>Default: 0</p> <ul style="list-style-type: none"> <li>■ <b>MonitorFreq:</b> This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer.</li> </ul> <p>Default: 1</p> <p>You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"> <li>■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources</li> <li>■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources</li> </ul> <ul style="list-style-type: none"> <li>■ <b>RegisterRetryLimit:</b> If you enable intelligent resource monitoring, the agent invokes the oracle_imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetryLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes.</li> </ul> <p>Default: 3.</p> |



**Table G-26** Optional attributes for Oracle agent (*continued*)

| Optional Attributes | Type and Dimension | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DetailMonitor       | integer-scalar     | <p>Determines whether or not detail monitoring is enabled. The attribute value 1 indicates that detail monitoring for Oracle is enabled.</p> <p>You must set the value of the LevelTwoMonitorFreq attribute to specify the frequency at which the agent must perform detailed monitoring.</p> <p>Default = 0</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| MonScript           | string-scalar      | <p>Pathname to the script provided for detail monitoring. The default (basic monitoring) is to monitor the database PIDs only.</p> <p><b>Note:</b> Detail monitoring is disabled if the value of the attribute MonScript is invalid or is set to an empty string.</p> <p>The pathname to the supplied detail monitor script is /opt/VRTSagents/ha/bin/Oracle/SqlTest.pl.</p> <p>MonScript also accepts a pathname relative to /opt/VRTSagents/ha. A relative pathname should start with "./", as in the path ./bin/Oracle/SqlTest.pl.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| User                | string-scalar      | Internal database user. Connects to the database for detail monitoring.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| LevelTwoMonitorFreq | integer-scalar     | <p>Specifies the frequency at which the agent for this resource type must perform second-level or detailed monitoring. You can also override the value of this attribute at resource-level.</p> <p>The value indicates the number of monitor cycles after which the agent will monitor Oracle in detail. For example, the value 5 indicates that the agent will monitor Oracle in detail every five online monitor intervals.</p> <p>If you manually upgraded to the SF Oracle RAC 5.1 SP1 agent, and if you had enabled detail monitoring in the previous version, then do the following:</p> <ul style="list-style-type: none"> <li>■ Set the value of the LevelTwoMonitorFreq attribute to the same value of that of the DetailMonitor attribute.</li> <li>■ Reset the value of the DetailMonitor attribute to 1</li> </ul> <p><b>Note:</b> If you set the AutoEndBkup attribute value to 0, then make sure that the LevelTwoMonitorFreq attribute value is 1 for detail monitoring.</p> <p>Default = 0</p> |
| Pword               | string-scalar      | <p>Encrypted password for internal database-user authentication.</p> <p>Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the SF Oracle RAC Encrypt utility.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Table G-26** Optional attributes for Oracle agent (*continued*)

| Optional Attributes | Type and Dimension | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Table               | string-scalar      | Table for update by User/Pword.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Encoding            | string-scalar      | <p>Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output.</p> <p>For example, if Oracle output is in "JAPANESE_JAPAN.JA16EUC," then "eucJP" is the Solaris value for Encoding. Refer to the Oracle and Solaris documentation for respective encoding values.</p> <p>Default is "".</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| IntentionalOffline  |                    | <p>This resource-type level attribute defines how VCS reacts when Oracle is intentionally stopped outside of VCS control.</p> <p>If you stop Oracle out of VCS control, the agent behavior is as follows:</p> <ul style="list-style-type: none"> <li>■ 0—The Oracle agent registers a fault and initiates the failover of the service group.</li> <li>■ 1—The Oracle agent takes the Oracle resource offline when Health check monitoring is enabled.</li> </ul> <p>If Health check monitoring is not enabled, the agent registers a fault and initiates the failover of the service group.</p> <p><b>Note:</b> If you want to use the intentional offline functionality of the agent, you must set the value of the MonitorOption attribute as 1 to enable Health check monitoring.</p> <p>See the <i>Veritas Cluster Server Administrator's Guide</i>.</p> |
| AgentDebug          | boolean-scalar     | <p>Additional debug messages are logged when this flag is set.</p> <p>Default = 0</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

[Table G-27](#) lists the internal attribute for Oracle agent. This attribute is for internal use only. Symantec recommends not to modify the value of this attribute.

**Table G-27** Internal attributes for Oracle agent

| Optional Attributes | Type and Dimension | Definition                                                                                                                                        |
|---------------------|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| AgentDirectory      | static-string      | <p>Specifies the location of binaries, scripts, and other files related to the Oracle agent.</p> <p>Default is /opt/VRTSagents/ha/bin/Oracle.</p> |

## Netlsnr agent functions

The Netlsnr agent is IMF-aware.

The listener is a server process that listens to incoming client connection requests and manages traffic to the database. The Netlsnr agent brings the listener services online, monitors their status, and takes them offline.

[Table G-28](#) lists the Netlsnr agent functions.

**Table G-28** Netlsnr agent functions

| Agent operation                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online                                   | Starts the listener process by using the following command:<br><br><code>lsnrctl start \$LISTENER</code>                                                                                                                                                                                                                                                                                                                                                                                             |
| Offline                                  | Stops the listener process by using the following command:<br><br><code>lsnrctl stop \$LISTENER</code><br><br>If the listener is configured with a password, the agent uses the password to stop the listener.                                                                                                                                                                                                                                                                                       |
| Monitor                                  | Verifies the status of the listener process.<br><br>The Netlsnr agent provides two levels of monitoring, basic and detail: <ul style="list-style-type: none"><li>■ In the basic monitoring mode, the agent scans the process table for the <code>tnslsnr</code> process to verify that the listener process is running.</li><li>■ In the detail monitoring mode, the agent uses the <code>lsnrctl status \$LISTENER</code> command to verify the status of the Listener process. (Default)</li></ul> |
| <code>netlsnr_imf_init</code>            | Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for Netlsnr agent. This function runs when the agent starts up.                                                                                                                                                                                                                                                                                                                                  |
| <code>netlsnr_imf_getnotification</code> | Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.                                                                                                                                                                                                                                                                     |
| <code>netlsnr_imf_register</code>        | Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).                                                                                                                                                                                                                                                                                                                           |
| Clean                                    | Scans the process table for <code>tnslsnr \$LISTENER</code> and kills it.                                                                                                                                                                                                                                                                                                                                                                                                                            |

Table G-28      Netlsnr agent functions (continued)

| Agent operation | Description                                                                                                                                    |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Action          | Performs the predefined actions on a resource.<br><br>See <a href="#">“Action entry point for SF Oracle RAC agent for Oracle”</a> on page 762. |

## Resource type definition for the Netlsnr agent

The Netlsnr agent of the Veritas Cluster Server agent for Oracle is represented by the Netlsnr resource type in SF Oracle RAC.

```
type Netlsnr (
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/Netlsnr"
    static keylist SupportedActions = { VRTS_GetInstanceName,
        VRTS_GetRunningServices, "tnsadmin.vfd" }
    static str ArgList[] = { Owner, Home, TnsAdmin, Listener,
        EnvFile, MonScript, LsnrPwd, AgentDebug, Encoding }
    static str IMFRegList[] = { Home, Owner, Listener }
    static int ContainerOpts{} = { RunInContainer=1, PassCInfo=0 }
    str Owner
    str Home
    str TnsAdmin
    str Listener
    str EnvFile
    str MonScript = "./bin/Netlsnr/LsnrTest.pl"
    str LsnrPwd
    boolean AgentDebug = 0
    str Encoding
    static boolean IntentionalOffline = 0
)
```

## Attribute definition for the Netlsnr agent

Review the description of the Netlsnr agent attributes. The agent attributes are classified as required, optional, and internal.

[Table G-29](#) lists the required attributes for Netlsnr agent. You must assign values to the required attributes.

**Table G-29** Required attributes for Netlsnr agent

| Required attributes | Type and dimension | Definition                                                                                                                                                                        |
|---------------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Owner               | string-scalar      | The Oracle user who has privileges to start or stop the listener process.<br>The agent also supports LDAP users as Oracle user.                                                   |
| Home                | string-scalar      | The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home.<br>Do not append a slash (/) at the end of the path. |

[Table G-30](#) lists the optional attributes for Netlsnr agent. You can configure the optional attributes if necessary.

**Table G-30** Optional attributes for Netlsnr agent

| Optional attributes | Type and dimension | Definition                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TnsAdmin            | string-scalar      | The \$TNS_ADMIN path to directory in which the Listener configuration file resides (listener.ora).<br>Default is /var/opt/oracle.                                                                                                                                                                                                                                                                   |
| Listener            | string-scalar      | Name of Listener. The name for Listener is considered case-insensitive by the Netlsnr agent and the Oracle database server.<br>Default is LISTENER.                                                                                                                                                                                                                                                 |
| LsnrPwd             | string-scalar      | The SF Oracle RAC encrypted password used to stop and monitor the listener. This password is set in the Listener configuration file.<br>Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the SF Oracle RAC Encrypt utility.                                                                                                                      |
| EnvFile             | string-scalar      | Specifies the full path name of the file that is sourced by the entry point scripts. This file contains the environment variables set by the user for the Oracle listener environment such as LD_LIBRARY_PATH and so on.<br>The syntax for the contents of the file depends on the login shell of Owner. This file must be readable by Owner. The file must not contain any prompts for user input. |

**Table G-30** Optional attributes for Netlsnr agent (*continued*)

| Optional attributes | Type and dimension  | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IMF                 | integer-association | <p>This resource-type level attribute determines whether the Netlsnr agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> <li>■ <b>Mode:</b> Define this attribute to enable or disable intelligent resource monitoring.</li> </ul> <p>Valid values are as follows:</p> <ul style="list-style-type: none"> <li>■ 0—Does not perform intelligent resource monitoring</li> <li>■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources</li> <li>■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources</li> <li>■ 3—Performs intelligent resource monitoring for both online and for offline resources</li> </ul> <p>Default: 0</p> <ul style="list-style-type: none"> <li>■ <b>MonitorFreq:</b> This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer.</li> </ul> <p>Default: 1</p> <p>You can set this attribute to a non-zero value in some cases where the agent requires to perform poll-based resource monitoring in addition to the intelligent resource monitoring.</p> <p>After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows:</p> <ul style="list-style-type: none"> <li>■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources</li> <li>■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources</li> </ul> <ul style="list-style-type: none"> <li>■ <b>RegisterRetryLimit:</b> If you enable intelligent resource monitoring, the agent invokes the <code>netlsnr_imf_register</code> agent function to register the resource with the AMF kernel driver. The value of the RegisterRetryLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes.</li> </ul> <p>Default: 3.</p> |

**Table G-30** Optional attributes for Netlsnr agent (*continued*)

| Optional attributes | Type and dimension | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MonScript           | string-scalar      | <p>Pathname to the script provided for detail monitoring. By default, the detail monitoring is enabled to monitor the listener process.</p> <p><b>Note:</b> If the value of the attribute MonScript is set to an empty string, the agent disables detail monitoring.</p> <p>The pathname to the supplied detail monitoring script is /opt/VRTSagents/ha/bin/Netlsnr/LsnrTest.pl.</p> <p>MonScript also accepts a pathname relative to /opt/VRTSagents/ha. A relative pathname should start with "./", as in the path ./bin/Netlsnr/LsnrTest.pl.</p> |
| LevelTwoMonitorFreq | integer-scalar     | <p>Specifies the frequency at which the agent for this resource type must perform second-level or detailed monitoring.</p> <p>If you enabled detail monitoring, then set the value of the LevelTwoMonitorFreq attribute.</p> <p>Default = 0</p>                                                                                                                                                                                                                                                                                                     |
| Encoding            | string-scalar      | <p>Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output.</p> <p>For example, if Oracle output is in "JAPANESE_JAPAN.JA16EUC," then "eucJP" is the Solaris value for Encoding. Refer to the Oracle and Solaris documentation for respective encoding values.</p> <p>Default is "".</p>                                                                                                                                                                                                            |
| IntentionalOffline  |                    | <p>For future use.</p> <p>Do not change the value of this attribute.</p> <p>Default = 0</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| AgentDebug          | boolean            | <p>Additional debug messages are logged when this flag is set.</p> <p>Default = 0</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Table G-31** lists the internal attribute for Netlsnr agent. This attribute is for internal use only. Symantec recommends not to modify the value of this attribute.

Table G-31 Internal attributes for Netlsnr agent

| Optional Attributes | Type and Dimension | Definition                                                                                                                                   |
|---------------------|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| AgentDirectory      | static-string      | Specifies the location of binaries, scripts, and other files related to the Netlsnr agent.<br><br>Default is /opt/VRTSagents/ha/bin/Netlsnr. |

ASMDG agent functions

The ASMDG agent mounts the ASM disk groups that the Oracle databases use, monitors the status, unmounts the ASM disk groups.

You must have specified the disk group names in the DiskGroups attribute of the ASMDG agent.

Table G-32 lists the ASMDG agent operations.

Table G-32 ASMDG agent operations

| Agent operation | Description                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online          | Mounts the specified Oracle ASM disk groups to an ASM instance by using the following SQL command:<br><br><code>alter diskgroup dg_name1, dg_name2 mount</code>                                                                                                                                                                                                                                                               |
| Offline         | Unmounts the specified Oracle ASM disk groups from an ASM instance by using the following SQL command:<br><br><code>alter diskgroup dg_name1, dg_name2 dismount</code><br><br><b>Note:</b> The following Oracle message appears in the VCS log when an ASM instance with no ASM disk groups mounted is shut down:<br>ORA-15100: invalid or missing diskgroup name                                                             |
| Monitor         | Verifies the status of the specified ASM disk groups.<br>The disk groups can be in one of the following states: <ul style="list-style-type: none"><li>■ mounted</li><li>■ dismounted</li><li>■ unknown</li><li>■ broken</li><li>■ connected</li></ul> If multiple ASM disk groups are configured for a resource, then the ASMDG agent returns the resource state considering the status of all the specified ASM disk groups. |



**Table G-32** ASMDG agent operations (*continued*)

| Agent operation | Description                                                                                                                                              |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Clean           | Forcibly unmounts the Oracle ASM disk groups by using the following SQL command:<br><br>alter diskgroup <i>dg_name1</i> , <i>dg_name2</i> dismount force |

## Resource type definition for the ASMDG agent

The ASMDG agent of the Veritas Cluster Server agent for Oracle is represented by the ASMDG resource type in SF Oracle RAC.

```
type ASMDG (
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/ASMDG"
    static str ArgList[] = { Sid, Owner, Home, DBAUser,
        DBAPword, DiskGroups, EnvFile, Encoding }
    static int ContainerOpts{} = { RunInContainer=1, PassCInfo=0 }
    str Sid
    str Owner
    str Home
    str DBAUser
    str DBAPword
    keylist DiskGroups
    str EnvFile
    str Encoding
)
```

## Attribute definition for the ASMDG agent

Review the description of the ASMDG agent attributes. The agent attributes are classified as required, optional, and internal.

[Table G-33](#) lists the required attributes. You must assign values to the required attributes.

**Table G-33** Required attributes for ASMDG agent

| Required attributes | Type and dimension | Definition                                                      |
|---------------------|--------------------|-----------------------------------------------------------------|
| DiskGroups          | keylist            | The ASM disk groups, where you store the Oracle database files. |
| Sid                 | string-scalar      | The variable \$ORACLE_SID that represents the ASM instance.     |

**Table G-33** Required attributes for ASMDG agent (*continued*)

| Required attributes | Type and dimension | Definition                                                                                                                                                                                             |
|---------------------|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Owner               | string-scalar      | The Oracle user who has privileges to mount or unmount the ASM disk group.<br><br>The agent also supports LDAP users as Oracle user.                                                                   |
| Home                | string-scalar      | The \$ORACLE_HOME path to Oracle ASM binaries and configuration files. For example, you could specify the path as /opt/ora_home.<br><br><b>Note:</b> Do not append a slash (/) at the end of the path. |

[Table G-34](#) lists the optional attributes for ASMDG agent. You can configure the optional attributes if necessary.

**Table G-34** Optional attributes for ASMDG agent

| Optional Attributes | Type and Dimension | Definition                                                                                                                                                                    |
|---------------------|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DBAUser             | string-scalar      | The ASM user who has sysasm privileges to start or stop the ASM instance. You can create ASM users for Oracle 11g R1 and later.                                               |
| DBAPword            | string-scalar      | Encrypted password for DBAUser.<br><br>Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the SF Oracle RAC Encrypt utility. |

**Table G-34** Optional attributes for ASMDG agent (*continued*)

| Optional Attributes | Type and Dimension | Definition                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EnvFile             | string-scalar      | The full path name of the file that is sourced by the entry point scripts. This file contains the environment variables set by the user for the Oracle database server environment such as LD_LIBRARY_PATH, NLS_DATE_FORMAT, and so on.<br><br>The syntax for the contents of the file depends on the login shell of Owner. File must be readable by Owner. The file must not contain any prompts for user input. |
| Encoding            | string-scalar      | Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. Default is "".                                                                                                                                                                                                                                                                                           |

[Table G-35](#) lists the internal attribute for ASMDG agent. This attribute is for internal use only. Symantec recommends not to modify the value of this attribute.

**Table G-35** Internal attributes for ASMDG agent

| Optional Attributes | Type and Dimension | Definition                                                                                                                                  |
|---------------------|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| AgentDirectory      | static-string      | Specifies the location of binaries, scripts, and other files related to the ASMDG agent.<br><br>Default is<br>/opt/VRTSagents/ha/bin/ASMDG. |

## CRSResource agent

The `CRSResource` agent provides an alternative mechanism for monitoring the Oracle database in the absence of the VCS Oracle agent. It checks the status of the Oracle Clusterware resources, which include the Oracle database instance, the listener, and the virtual IP address (VIP). The agent supports multiple database configurations and ensures that the Oracle database is online and available to an application when it starts.

## Functions of the CRSResource agent

Table G-36 describes the functions of the CRSResource agent.

Table G-36 CRSResource agent entry points

| Entry point | Description                                                                                                                                                                                                                                                              |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor     | <p>Checks the status of the following Oracle Clusterware resources:</p> <ul style="list-style-type: none"><li>■ Oracle database</li><li>■ Virtual IP address</li><li>■ Listener</li></ul> <p>The monitor script checks the status of the resources every 60 seconds.</p> |
| Online      | <p>A dummy entry point that uses a delay interval to ensure that the Oracle Clusterware resources come online before the monitor operation begins.</p>                                                                                                                   |

## States of the CRSResource agent

Table G-37 lists the states of the CRSResource agent.

Table G-37 States of the CRSResource agent

| State   | Description                                                                                                                               |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Online  | <p>Indicates that the Oracle database, VIP, and listener are running.</p>                                                                 |
| Faulted | <p>Indicates that the Oracle database, VIP, and listener are not running or are faulted.</p>                                              |
| Unknown | <p>Indicates the inability to determine the state of the resource due to incorrect attributes settings or other configuration issues.</p> |

## Attributes of the CRSResource agent

**Table G-38** Attributes of the CRSResource agent

| Attribute | Dimension     | Description                                                                                                                                                                                                           |
|-----------|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ResType   | string-scalar | Indicates the type of the Oracle Clusterware resource.<br><br>The types of resources are as follows: <ul style="list-style-type: none"><li>■ DB</li><li>■ VIP</li><li>■ Listener</li></ul>                            |
| DBHome    | string-scalar | The path to ORACLE_HOME containing the Oracle binaries and configuration files. For example, you could specify the path as /app/oracle/orahome.<br><br><b>Note:</b> Do not append a slash (/) at the end of the path. |
| CRSHome   | string-scalar | The path to CRS_HOME containing the Oracle Clusterware binaries.                                                                                                                                                      |
| DBName    | string-scalar | The name of the database on which the services are configured.<br><br>This attribute is optional if you are configuring CRSResource to monitor VIP or listener.                                                       |
| SID       | string-scalar | The name of the Oracle instance in the variable \$ORACLE_SID. The SID is case-sensitive.<br><br>This attribute is optional if you are configuring CRSResource to monitor VIP or listener.                             |
| Owner     | string-scalar | The name of the user for the Oracle database.<br><br>This attribute is optional if you are configuring CRSResource to monitor VIP or listener.                                                                        |

## VCS service group dependencies with the CRSResource agent

In a service group configuration with the CRSResource agent, Oracle Clusterware controls the database. An online local firm dependency exists between the groups—Application group, Oracle Clusterware group, and the CVM group.

Figure G-4 shows a schematic illustration of the service group dependencies.

**Figure G-4** Service group dependencies with CRSResource agent



In the configuration:

- When the system starts, the CVM group brings up the volume and mount points for the databases. The Oracle Clusterware group brings up the OCR and voting disk, configures the private IP address for Oracle Clusterware, and starts Oracle Clusterware. Oracle Clusterware starts the database and the application is brought online. CRSResource comes online when the Oracle Clusterware resources (database/VIP/listener) are started by Oracle Clusterware.

---

**Note:** When the system starts, all volumes and mount points **MUST** be online for the dependent service groups to be online.

---

- The oradata\_mnt and oradata\_voldg resources are configured as non-critical resources (critical=0) for managing failure scenarios. See [“How CRSResource agent handles failures”](#) on page 783.
- When CRSResource faults for any of the Oracle Clusterware resources, the application is brought offline.

The limitations of this configuration are as follows:

- The CFSSMount and CVMVolDg resources can not be set as critical resources in the group.  
If the mount points and volume disk groups for all the databases are configured as critical in a single service group, then failure of any of them results in the whole group being FAULTED or brought offline. To ensure that a resource failure does not affect other resources in the group, the attribute Critical is set to zero for the CFSSMount and CVMVolDg resources.  
However, if any of the database mounts fail to come online or a volume does not start, the whole service group fails to come online.
- CRSResource reports as FAULTED until Oracle Clusterware brings up the database instance, VIP, and listener. Even after Oracle Clusterware starts the database instance, VIP and listener, CRSResource remains in the FAULTED state for the OfflineMonitorInterval period. The status of CRSResource cannot be changed.

## How CRSResource agent handles failures

The CRSResource agent ensures that faults in the resources of an application do not adversely impact other applications running on the system. To isolate failures, the oradata\_mnt and oradata\_voldg resources are configured as non-critical resources (critical=0). This ensures that storage issues in either of these resources do not affect the other databases and the dependent application continues to be online.


---

**Note:** The resources are considered non-critical only for the purpose of managing failure scenarios.

---

[Figure G-5](#) illustrates a failure scenario.

**Figure G-5** How CRSResource agent handles failures



## Fault configurations with CRSResource agent

This section discusses scenarios that cause CRSResource to report FAULTED.

### Scenario 1: CRSResource fault at system startup

1. The CVM group brings online the volume and mount points for the databases when the system starts.
2. The Oracle Clusterware group brings up the OCR and voting disk, configures the private IP address for Oracle Clusterware and starts Oracle Clusterware.




Oracle Clusterware is unable to bring the database or other resources online within the start duration configured in the agent.

3. The agent starts the monitoring operation after the set duration.

If the Oracle Clusterware resources are not brought online yet, the CRSResource appears faulted.

Figure G-6 illustrates the scenario.

Figure G-6 CRSResource fault at system startup



### Scenario 2: CRSResource fault when resource is brought offline

Any or all of the Oracle Clusterware resources (database, listener, or VIP) are brought offline due to a fault or for administrative purposes.

CRSResource reports FAULTED until Oracle Clusterware brings the resources online.

## Resource type definition for the CRSResource agent

The following extract from the `CRSResource.cf` file describes the type definition for the CRSResource agent.

```
type CRSResource (
    static int MonitorTimeout = 300
    static int OfflineMonitorInterval = 60
    static str ArgList[] = { ResType, DBHome, CRSHome, DBName, SID, Owner }
    static str Operations = None
    str ResType
    str DBHome
    str CRSHome
    str DBName
    str SID
    str Owner
)
```

## Sample configuration for the CRSResource agent

The following sample configuration describes the DB, VIP, and Listener resource configurations for the CRSResource agent.

```
CRSResource crsr_oradb_db (
    ResType = DB
    DBHome = "/app/oracle/orahome"
    CRSHome = "/app/crshome"
    DBName = oradb
    SID @galaxy = oradb1
    SID @nebula = oradb2
    Owner = oracle
)

CRSResource crsr_oradb_vip (
    ResType = VIP
    DBHome = "/app/oracle/orahome"
    CRSHome = "/app/crshome"
)
```

```
CRSResource crsr_oradb_lsnr (  
    ResType = Listener  
    DBHome = "/app/oracle/orahome"  
    CRSHome = "/app/crshome"  
)
```



# SF Oracle RAC deployment scenarios

This appendix includes the following topics:

- [SF Oracle RAC cluster with VCS IPC and PrivNIC agent](#)
- [SF Oracle RAC cluster with UDP IPC and PrivNIC agent](#)
- [SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent](#)
- [SF Oracle RAC cluster with isolated Oracle traffic and MultiPrivNIC agent](#)
- [SF Oracle RAC cluster with NIC bonding, VCS IPC and PrivNIC agent](#)
- [SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent](#)
- [Configuration diagrams for setting up server-based I/O fencing](#)
- [SF Oracle RAC deployment scenarios in Oracle VM Server for SPARC environments](#)
- [Deploying Storage Foundation for Databases \(SFDB\) tools in a Storage Foundation for Oracle RAC environment](#)

## SF Oracle RAC cluster with VCS IPC and PrivNIC agent

This section illustrates the recommended configuration for the following scenario:

|                     |                                                                                                                 |
|---------------------|-----------------------------------------------------------------------------------------------------------------|
| Deployment scenario | Oracle RAC 10g database cache fusion traffic is distributed over multiple LLT links using VCS IPC over LMX/LLT. |
| Recommendation      | Use the PrivNIC agent.                                                                                          |

Sample main.cf  
configuration file

The following sample main.cf file describes the configuration:


`/etc/VRTSvcs/conf/sample_rac/sfrac01_main.cf`

In the illustrated configuration:

- Oracle uses VCS IPC over LMX/LLT for cache fusion.
- One private network IP address is used for Oracle Clusterware communication that takes place over one of the LLT private interconnect links.
- The CFS/CVM/VCS metadata and the Oracle database cache fusion traffic travels through LLT over the two physical links that are configured as LLT links.
- In the event of a NIC failure or link failure, the PrivNIC agent fails over the private network IP address from the failed link to the available LLT link. When the link is restored, the IP address is failed back to the original link.


[Figure H-1](#) illustrates the logical view of a two-node SF Oracle RAC cluster with VCS IPC and PrivNIC agent (Oracle RAC 10g).

**Figure H-1** SF Oracle RAC cluster with VCS IPC and PrivNIC agent (Oracle RAC 10g)



**Legends**

Public link (GigE)  
Private Interconnect (GigE)  
Oracle Clusterware Heartbeat  
Oracle DB Cache Fusion  
CFS /CVM / VCS Metadata



# SF Oracle RAC cluster with UDP IPC and PrivNIC agent

This section illustrates the recommended configuration for the following scenario:

|                                   |                                                                                                                             |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Deployment scenario               | Oracle RAC configured with UDP IPC for database cache fusion.                                                               |
| Recommendation                    | Use the PrivNIC agent.                                                                                                      |
| Sample main.cf configuration file | The following sample main.cf file describes the configuration:<br><code>/etc/VRTSvcs/conf/sample_rac/sfrac02_main.cf</code> |


In the illustrated configuration:

- A common IP address is used for Oracle Clusterware communication and Oracle database cache fusion.
- Oracle Clusterware communication and Oracle database cache fusion traffic flows over one of the LLT private interconnect links.
- The CFS/CVM/VCS metadata travels through LLT over the two physical links that are configured as LLT links.
- In the event of a link failure, the PrivNIC agent fails over the private network IP address from the failed link to the available LLT link. When the link is restored, the IP address is failed back to the original link.

Figure H-2 illustrates the logical view of a two-node SF Oracle RAC cluster with UDP IPC and PrivNIC agent.



Figure H-2 SF Oracle RAC cluster with UDP IPC and PrivNIC agent



# SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent

This section illustrates the recommended configuration for the following scenario:


|                                   |                                                                                                                             |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Deployment scenario               | Oracle RAC is configured with UDP IPC for database cache fusion.                                                            |
| Recommendation                    | Use the MultiPrivNIC agent.                                                                                                 |
| Sample main.cf configuration file | The following sample main.cf file describes the configuration:<br><code>/etc/VRTSvcs/conf/sample_rac/sfrac03_main.cf</code> |

In the illustrated configuration:

- One private IP address is used for each database for database cache fusion. One of the private IP addresses used for the database cache fusion is shared by Oracle Clusterware communication. The CFS/CVM/VCS metadata also travels through these links.
- In the event of a link failure, the MultiPrivNIC agent fails over the private network IP address from the failed link to the available LLT link. When the link is restored, the IP address is failed back to the original link.

Figure H-3 illustrates the logical view of a two-node SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent.

**Figure H-3** SF Oracle RAC cluster for multiple databases with UDP IPC and MultiPrivNIC agent



# SF Oracle RAC cluster with isolated Oracle traffic and MultiPrivNIC agent

This section illustrates the recommended configuration for the following scenario:


|                                   |                                                                                                                              |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| Deployment scenario               | Oracle RAC database cache fusion traffic is isolated from the CFS/CVM/VCS metadata.                                          |
| Recommendation                    | Use the MultiPrivNIC agent.                                                                                                  |
| Sample main.cf configuration file | The following sample main.cf file describes the configuration:<br><code>/etc/VRTSvcsc/conf/sample_rac/sfrac06_main.cf</code> |

In the illustrated configuration:

- The private network IP address used for database cache fusion is configured over a dedicated link for each database. These links are configured as low-priority LLT links.
- The private network IP address used for Oracle Clusterware communication is configured over a high-priority LLT link. This link is also used for the CFS/CVM/VCS metadata transfer.
- In the event of a link failure, the MultiPrivNIC agent fails over the private network IP address from the failed link to the available LLT link. When the link is restored, the IP address is failed back to the original link.

Figure H-4 illustrates the logical view of a two-node SF Oracle RAC cluster with Oracle traffic isolated from VCS / CVM / CFS traffic.

**Figure H-4** SF Oracle RAC cluster with isolated Oracle traffic and MultiPrivNIC agent



# SF Oracle RAC cluster with NIC bonding, VCS IPC and PrivNIC agent

This section illustrates the recommended configuration for the following scenario:


|                                   |                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Deployment scenario               | A bonded NIC interface is used along with another NIC interface to distribute Oracle RAC 10g database cache fusion traffic using VCS IPC over LMX/LLT.                                                                                                                                                                                       |
| Recommendation                    | Use the PrivNIC agent.                                                                                                                                                                                                                                                                                                                       |
| Sample main.cf configuration file | <p>The following sample main.cf file describes the configuration:</p> <p><code>/etc/VRTSvcs/conf/sample_rac/sfrac01_main.cf</code></p> <p><b>Note:</b> You must replace the bge1 interface in the sample file with the bonded NIC interface you use in the scenario and bge2 in the sample file with the NIC3 interface in the scenario.</p> |

In the illustrated configuration:

- Oracle uses VCS IPC over LMX/LLT for cache fusion.
- One common private network IP address is used for Oracle database cache fusion and Oracle Clusterware communication that takes place over the bonded NIC interface, configured as an LLT link.
- The Oracle Clusterware communication takes place over the other LLT private interconnect link.
- The Oracle database cache fusion as well as the CFS/CVM/VCS metadata travels through LLT over the three physical links.
- In the event of a bonded NIC interface failure, the PrivNIC agent fails over the private network IP address from the failed link to the other available LLT link. When the link is restored, the IP address is failed back to the original link.

Figure H-5 illustrates the logical view of a two-node SF Oracle RAC cluster with NIC bonding, VCS IPC, and PrivNIC agent (Oracle RAC 10g).

**Figure H-5** SF Oracle RAC cluster with NIC bonding, VCS IPC and PrivNIC agent (Oracle RAC 10g)



# SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent

This section illustrates the recommended configuration for the following scenario:

|                                   |                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Deployment scenario               | Oracle RAC with UDP IPC is configured to use a bonded NIC interface for distribution of Oracle database cache fusion traffic.<br><br>A second link is configured as a standby link.                                                                                                                                                   |
| Recommendation                    | Use the PrivNIC agent.                                                                                                                                                                                                                                                                                                                |
| Sample main.cf configuration file | The following sample main.cf file describes the configuration:<br><br><code>/etc/VRTSvcs/conf/sample_rac/sfrac02_main.cf</code><br><br><b>Note:</b> You must replace the bge1 interface in the sample file with the bonded NIC interface you use in the scenario and bge2 in the sample file with the NIC3 interface in the scenario. |


In the illustrated configuration:

- A common IP address is used for Oracle Clusterware communication and Oracle database cache fusion that is distributed over two underlying physical links of the bonded NIC interface. The bonded NIC interface is configured as a single LLT link.
- The CFS/CVM/VCS metadata also travels through the bonded link.
- In the event of a bonded link failure, the PrivNIC agent fails over the private network IP address from the failed link to the available standby LLT link. When the link is restored, the IP address is failed back to the original link.

Figure H-6 illustrates the logical view of a two-node SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent.



Figure H-6 SF Oracle RAC cluster with NIC bonding, UDP IPC, and PrivNIC agent



## Configuration diagrams for setting up server-based I/O fencing

The following CP server configuration diagrams can be used as guides when setting up CP server within your configuration:


- Two unique client clusters that are served by 3 CP servers:  
See [Figure H-7](#) on page 803.
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:

### Two unique client clusters served by 3 CP servers

[Figure H-7](#) displays a configuration where two unique client clusters are being served by 3 CP servers (coordination points). Each client cluster has its own unique user ID (UUID1 and UUID2).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen mechanism` set to `cps`.

**Figure H-7** Two unique client clusters served by 3 CP servers




## Client cluster served by highly available CPS and 2 SCSI-3 disks

**Figure H-8** displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfsnode` file on the client nodes, `vxfsnode` is set to customized with `vxfsnode` mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group vxencoordddg. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

**Figure H-8** Client cluster served by highly available CP server and 2 SCSI-3 disks




## Two node campus cluster served by remote CP server and 2 SCSI-3 disks

[Figure H-9](#) displays a configuration where a two node campus cluster is being served by one remote CP server and 2 local SCSI-3 LUN (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks (one from each site) are part of disk group `vxfencoorddg`. The third coordination point is a CP server on a single node VCS cluster.

Two node campus cluster served by remote CP server and 2 SCSI-3



## Legends

- Private Interconnects (GigE) —
- Public Links (GigE) —
- Dark Fiber —
- Connections —
- San 1 Connections —
- San 2 Connections —


## Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

[Figure H-10](#) displays a configuration where multiple client clusters are being served by one highly available CP server and 2 local SCSI-3 LUNS (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group `vxfencoorddg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

**Figure H-10** Multiple client clusters served by highly available CP server and 2 SCSI-3 disks





# SF Oracle RAC deployment scenarios in Oracle VM Server for SPARC environments

Oracle VM Server for SPARC (earlier known as Logical Domains) from Oracle is a technology that allocates resources such as processors, memory, disks or network devices to logical containers and manages them as logical domains within the physical host. The resulting logical domain has its own operating system and manages resources independently in its realm.

Oracle VM Servers can be set up to function in any of the following roles:

|                |                                                                                                                                                                                                                                                                                                                           |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Control domain | The Control domain is the physical host that has access to all physical resources on the system. The Logical Domains manager software, which is used for managing the guest and I/O domains, is installed on this domain. The Control domain can also serve as an I/O domain and provide services to other guest domains. |
| I/O domain     | The I/O domain has direct access to I/O devices. There can be a maximum of two I/O domains in an LDom setup. The domain is also called as the service domain when it provides I/O services to other guest domains.                                                                                                        |
| Guest domain   | The Guest domain uses the services delivered by the service domain to access physical devices. The Guest domain exists as an independent entity with virtual resources and own copy of operating system.                                                                                                                  |

LDom provides a cost-effective alternative architecture for deploying SF Oracle RAC. The same physical server can be used for multiple applications within various logical domains with optimal resource utilization. LDom is hardware-dependent and works with the latest SUN CoolThreads servers.

For detailed information, see the Oracle documentation.

## Sample configuration scenarios

The following server configuration is used for the sample scenarios presented in this document:

|           |                                                     |
|-----------|-----------------------------------------------------|
| Server    | SUN T2000 Server                                    |
| Processor | 24 T1 Based Processor (4Cores * 6 Threads per core) |
| Memory    | 8 GB                                                |

|                       |                                                                                                                                                                                                                                         |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PCI devices           | 2 NIC Cards + 1 HBA per BUS, all on board hard drives belong to one PCI bus                                                                                                                                                             |
| Firmware version      | SPARC-Enterprise-T2000 System Firmware 6.7.3                                                                                                                                                                                            |
| Operating system      | Solaris 10 Update 6                                                                                                                                                                                                                     |
| LDoms version         | LDoms Manager 1.1, SF Oracle RAC 5.1, Oracle 10g (10.2.0.4)                                                                                                                                                                             |
| Database version      | Oracle RAC 10g Release 2 (10.2.0.4)<br><br>For the latest information on supported Oracle database versions, see the following TechNote:<br><a href="http://www.symantec.com/docs/TECH44807">http://www.symantec.com/docs/TECH44807</a> |
| SF Oracle RAC version | SF Oracle RAC 5.0 MP3<br><br>SF Oracle RAC 5.1                                                                                                                                                                                          |

For more details of the firmware levels required for LDoms manager 1.1, visit the Oracle documentation.

The sample deployment scenarios are as follows:

- SF Oracle RAC with Oracle RAC database on I/O domains of two hosts  
See “[SF Oracle RAC with Oracle RAC database on I/O domains of two hosts](#)” on page 811.
- SF Oracle RAC with Oracle RAC database on Guest domains of two hosts  
See “[SF Oracle RAC with Oracle RAC database on guest domains of two hosts](#)” on page 812.
- SF Oracle RAC with Oracle RAC database on Guest domains of a single host  
See “[SF Oracle RAC with Oracle RAC database on guest domains of single host](#)” on page 814.
- SF Oracle RAC with Oracle RAC database on Guest domain and I/O domain of a single host  
See “[SF Oracle RAC with Oracle RAC database on I/O domain and guest domain of single host](#)” on page 816.

## Preparing to deploy SF Oracle RAC in Logical Domain environments

Complete the following tasks before you deploy SF Oracle RAC in LDom environments:

1. Install LDoms Manager software.
2. Configure the physical host as a primary domain.  
  
The primary domain will serve as the control/IO/service domain by default.

3. Create new LDom configurations.

## SF Oracle RAC with Oracle RAC database on I/O domains of two hosts

This section describes the tasks required to set up I/O domains on two hosts for deploying SF Oracle RAC.


The benefits of this approach are as follows:

- The computing resources on the hosts are available to other LDomS.
- Direct access to storage ensures better database performance.

The disadvantage of this approach is that if the I/O load is extensive, the I/O latency on other LDomS will be high.

Figure H-11 illustrates the scenario.

**Figure H-11** SF Oracle RAC with Oracle RAC database on I/O domains of two hosts



### To set up I/O domains on two hosts for deploying SF Oracle RAC

- 1 Complete the preparatory steps for setting up an LDom environment.  
See [“Preparing to deploy SF Oracle RAC in Logical Domain environments”](#) on page 810.
- 2 Create the secondary I/O domain.
- 3 Provision a PCI bus from the primary domain to the secondary I/O domain.
- 4 Create virtual network service in the primary domain.  
The virtual interface connected to the service will be used by the secondary I/O domain as its public interface. The interfaces that remain on the secondary I/O domain will be used for LLT heartbeat.
- 5 Create virtual disk service on the primary domain.  
The virtual disk connected to the service will be used by the secondary I/O domain as its root disk.
- 6 Install Solaris operating system on the secondary I/O domain using Jumpstart method.  
After OS installation, all devices present on the provisioned PCI bus are visible on the secondary I/O domain.
- 7 Repeat steps 2 to 6 on the second host.
- 8 Install and configure SF Oracle RAC.
- 9 Set up Oracle RAC database.

## SF Oracle RAC with Oracle RAC database on guest domains of two hosts

This section describes the tasks required to set up guest domains on two hosts for deploying SF Oracle RAC.


The benefits of this approach are as follows:

- The configuration provides a completely virtualized domain.
- The hardware resources can be effectively utilized with other LDoms.

The disadvantage of this approach is that since there is no direct disk access from the virtual domain to the physical disk, there could be a minor lag on disk access times.

[Figure H-12](#) illustrates the scenario.

**Figure H-12** SF Oracle RAC with Oracle RAC database on guest domains of two hosts



### To set up guest domains on two hosts for deploying SF Oracle RAC

- 1 Complete the preparatory steps for setting up an LDom environment.  
See [“Preparing to deploy SF Oracle RAC in Logical Domain environments”](#) on page 810.
- 2 Create virtual disk and network services on the primary domain.  
The virtual services are bound to the virtual devices created for the guest LDom. A guest LDom accesses the physical hardware through the virtual devices connected to the virtual services. Assign individual services for each disk from storage. This helps to maintain sequence of the disks in the guest domain.
- 3 Create the guest domain.

- 4 Create the virtual devices on the guest domain and bind it to the virtual services created on the primary domain.
- 5 Install Solaris operating system on the guest domain using Jumpstart method.
- 6 Repeat steps 2 to 5 on the second host.
- 7 Install and configure SF Oracle RAC.
- 8 Set up Oracle RAC database.

## SF Oracle RAC with Oracle RAC database on guest domains of single host

This section describes the tasks required to set up guest domains on a single host for deploying SF Oracle RAC.

---

**Note:** This setup is recommended for use as a four-node cluster by using an additional physical host with the same configuration.

---

The benefits of this approach are as follows:


- Reduction in the number of physical servers used makes it a very cost-effective setup.
- The setup is easy to create and maintain. It is also flexible and portable.
- Many guest LDomS from multiple systems can be joined together to form a bigger cluster.
- If the primary domain reboots, only the guest LDom attached to it is affected. The guest LDom attached to the secondary service domain continues to be available. Please note that shutting down the primary domain halts all domains.

The disadvantages of this approach are as follows:

- Hardware failures act as a single point of failure, bringing down all associated domains.
- Almost all physical resources are consumed in the creation of this setup on a T2000 Server.

[Figure H-13](#) illustrates the scenario.

**Figure H-13** SF Oracle RAC with Oracle RAC database on guest domains of single host



### To set up guest domains on a single host for deploying SF Oracle RAC

- 1 Complete the preparatory steps for setting up an LDom environment.  
See [“Preparing to deploy SF Oracle RAC in Logical Domain environments”](#) on page 810.
- 2 Create a split PCI configuration on a T2000 Server.  
Each bus has 2NICs and 1HBA.

**3** Create one primary domain and one secondary domain.

The primary domain must be configured to be used as a control and I/O Domain (primary service domain). The secondary domain must be initially configured as an I/O domain. Next, configure the secondary I/O domain to be used as a secondary service domain. The primary domain is used to create and assign the services for this secondary service domain.

Both the primary and secondary Service domains see the storage through its respective paths.

**4** Create the disk and network services on the primary service domain and assign it to Guest LDom A.

**5** Create and assign a different set of disk and network services for the secondary service domain using the primary service domain. This ensures that services for guest LDom B are provided by the secondary service domain.

**6** Install Solaris operating system on the secondary service domain and the guest LDoms using Jumpstart method.

**7** Install and configure SF Oracle RAC.

---

**Note:** Set the public link to be used as a low priority heartbeat link due to unavailability of NICs.

---

**8** Set up Oracle RAC database.

## SF Oracle RAC with Oracle RAC database on I/O domain and guest domain of single host

This section describes the tasks required to set up guest domains on single host for deploying SF Oracle RAC.

---

**Note:** This setup is recommended for use as a four-node cluster by using an additional physical host with the same configuration.

---

The benefits of this approach are as follows:

- guest LDoms can be added to the host at any time unlike in Scenario 3 where all resources are utilized leaving no room for additional LDoms.
- This setup results in better disk performance as compared to Scenario 3.




- If the primary domain reboots, only the guest LDom attached to it is affected. The guest LDom attached to the secondary service domain continues to be available. Please note that shutting down the primary domain halts all domains.

The disadvantage of this approach is that hardware failures act as a single point of failure, bringing down all associated domains.

Figure H-14 illustrates the scenario.

**Figure H-14** SF Oracle RAC with Oracle RAC database on I/O domain and guest domain of single host



### To set up guest domains of single host for deploying SF Oracle RAC

- 1 Complete the preparatory steps for setting up an LDom environment.  
See [“Preparing to deploy SF Oracle RAC in Logical Domain environments”](#) on page 810.
- 2 Create a split PCI configuration on a T2000 Server.  
Each bus has 2NICs and 1HBA.

- 3 Create the secondary I/O domain.
- 4 Provision a PCI bus from the primary domain to the secondary I/O domain.
- 5 Create virtual disk service on the primary domain. The virtual disk connected to the service is used by the secondary I/O domain as its root disk.
- 6 Create the disk and network services on the primary service domain and assign it to the guest domain.
- 7 Install Solaris operating system on the secondary service domain and the guest domain using Jumpstart method.
- 8 Install and configure SF Oracle RAC.  
Perform the following steps after configuring SF Oracle RAC:
  - Set the public link to be used as a low priority heartbeat link due to unavailability of NICs.
  - Run the command "oifcfg" after installation of Oracle Clusterware. This is required to resolve the issue of different interface names on the guest LDom and the secondary I/O LDom.
  - Configure MultiPrivNic manually.
- 9 Set up Oracle RAC database.


## Deploying Storage Foundation for Databases (SFDB) tools in a Storage Foundation for Oracle RAC environment

If you are deploying the SFDB tools with Storage Foundation for Oracle RAC (multiple instance Oracle) your setup configuration will reflect the following conditions:

- A highly available parallel cluster with a multiple instances of Oracle is set up on galaxy and nebula with SF for Oracle RAC.
- The database is online on galaxy, nebula, and mercury.
- The datafiles are mounted and shared on galaxy, nebula, and mercury.
- The SFDB tools is mounted and shared on galaxy, nebula, and mercury.
- You can run the SFDB tools commands on galaxy, nebula, and mercury.
- Clustered ODM is supported for this configuration.

In the figure below the repository directory resides in the Oracle mount points.

**Figure H-15**      Deploying the database repository with Storage Foundation



For an SF Oracle RAC configuration, the systems are online in parallel and do not use failover mechanisms within the cluster.



# Glossary

|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Agent</b>                        | A process that starts, stops, and monitors all configured resources of a type, and reports their status to VCS.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Authentication Broker</b>        | The Veritas Security Services component that serves, one level beneath the root broker, as an intermediate registration authority and a certification authority. The authentication broker can authenticate clients, such as users or services, and grant them a certificate that will become part of the Veritas credential. An authentication broker cannot, however, authenticate other brokers. That task must be performed by the root broker.                                                               |
| <b>Cluster</b>                      | A cluster is one or more computers that are linked together for the purpose of multiprocessing and high availability. The term is used synonymously with VCS cluster, meaning one or more computers that are part of the same GAB membership.                                                                                                                                                                                                                                                                     |
| <b>CVM (cluster volume manager)</b> | The cluster functionality of Veritas Volume Manager.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Disaster Recovery</b>            | Administrators with clusters in physically disparate areas can set the policy for migrating applications from one location to another if clusters in one geographic area become unavailable due to an unforeseen event. Disaster recovery requires heartbeating and replication.                                                                                                                                                                                                                                  |
| <b>disk array</b>                   | A collection of disks logically arranged into an object. Arrays tend to provide benefits such as redundancy or improved performance.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>DMP (Dynamic Multi-Pathing)</b>  | A feature designed to provide greater reliability and performance by using path failover and load balancing for multiported disk arrays connected to host systems through multiple paths. DMP detects the various paths to a disk using a mechanism that is specific to each supported array type. DMP can also differentiate between different enclosures of a supported array type that are connected to the same host system.                                                                                  |
| <b>SmartTier</b>                    | A feature with which administrators of multi-volume VxFS file systems can manage the placement of files on individual volumes in a volume set by defining placement policies that control both initial file location and the circumstances under which existing files are relocated. These placement policies cause the files to which they apply to be created and extended on specific subsets of a file system's volume set, known as placement classes. The files are relocated to volumes in other placement |

classes when they meet specified naming, timing, access rate, and storage capacity-related conditions.

**Failover** A failover occurs when a service group faults and is migrated to another system.

**GAB (Group Atomic Broadcast)** A communication mechanism of the VCS engine that manages cluster membership, monitors heartbeat communication, and distributes information throughout the cluster.

**HA (high availability)** The concept of configuring the product to be highly available against system failure on a clustered network using Veritas Cluster Server (VCS).

**IP address** An identifier for a computer or other device on a TCP/IP network, written as four eight-bit numbers separated by periods. Messages and other data are routed on the network according to their destination IP addresses.

See also virtual IP address

**Jeopardy** A node is in jeopardy when it is missing one of the two required heartbeat connections. When a node is running with one heartbeat only (in jeopardy), VCS does not restart the applications on a new node. This action of disabling failover is a safety mechanism that prevents data corruption.

**latency** For file systems, this typically refers to the amount of time it takes a given file system operation to return to the user.

**LLT (Low Latency Transport)** A communication mechanism of the VCS engine that provides kernel-to-kernel communications and monitors network communications.

**logical volume** A simple volume that resides on an extended partition on a basic disk and is limited to the space within the extended partitions. A logical volume can be formatted and assigned a drive letter, and it can be subdivided into logical drives.

See also LUN

**LUN** A LUN, or logical unit, can either correspond to a single physical disk, or to a collection of disks that are exported as a single logical entity, or virtual disk, by a device driver or by an intelligent disk array's hardware. VxVM and other software modules may be capable of automatically discovering the special characteristics of LUNs, or you can use disk tags to define new storage attributes. Disk tags are administered by using the `vxdisk` command or the graphical user interface.

**main.cf** The file in which the cluster configuration is stored.

**mirroring** A form of storage redundancy in which two or more identical copies of data are maintained on separate volumes. (Each duplicate copy is known as a mirror.) Also RAID Level 1.

**Node** The physical host or system on which applications and service groups reside. When systems are linked by VCS, they become nodes in a cluster.

|                                   |                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>resources</b>                  | Individual components that work together to provide application services to the public network. A resource may be a physical component such as a disk group or network interface card, a software component such as a database server or a Web server, or a configuration component such as an IP address or mounted file system.                                      |
| <b>Resource Dependency</b>        | A dependency between resources is indicated by the keyword "requires" between two resource names. This indicates the second resource (the child) must be online before the first resource (the parent) can be brought online. Conversely, the parent must be offline before the child can be taken offline. Also, faults of the children are propagated to the parent. |
| <b>Resource Types</b>             | Each resource in a cluster is identified by a unique name and classified according to its type. VCS includes a set of pre-defined resource types for storage, networking, and application services.                                                                                                                                                                    |
| <b>root broker</b>                | The first authentication broker, which has a self-signed certificate. The root broker has a single private domain that holds only the names of brokers that shall be considered valid.                                                                                                                                                                                 |
| <b>SAN (storage area network)</b> | A networking paradigm that provides easily reconfigurable connectivity between any subset of computers, disk storage and interconnecting hardware such as switches, hubs and bridges.                                                                                                                                                                                  |
| <b>Service Group</b>              | A service group is a collection of resources working together to provide application services to clients. It typically includes multiple resources, hardware- and software-based, working together to provide a single service.                                                                                                                                        |
| <b>Service Group Dependency</b>   | A service group dependency provides a mechanism by which two service groups can be linked by a dependency rule, similar to the way resources are linked.                                                                                                                                                                                                               |
| <b>Shared Storage</b>             | Storage devices that are connected to and used by two or more systems.                                                                                                                                                                                                                                                                                                 |
| <b>shared volume</b>              | A volume that belongs to a shared disk group and is open on more than one node at the same time.                                                                                                                                                                                                                                                                       |
| <b>SNMP Notification</b>          | Simple Network Management Protocol (SNMP) developed to manage nodes on an IP network.                                                                                                                                                                                                                                                                                  |
| <b>State</b>                      | The current activity status of a resource, group or system. Resource states are given relative to both systems.                                                                                                                                                                                                                                                        |
| <b>Storage Checkpoint</b>         | A facility that provides a consistent and stable view of a file system or database image and keeps track of modified data blocks since the last Storage Checkpoint.                                                                                                                                                                                                    |
| <b>System</b>                     | The physical system on which applications and service groups reside. When a system is linked by VCS, it becomes a node in a cluster.<br><br>See Node                                                                                                                                                                                                                   |
| <b>types.cf</b>                   | A file that describes standard resource types to the VCS engine; specifically, the data required to control a specific resource.                                                                                                                                                                                                                                       |

|                                        |                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>VCS (Veritas Cluster Server)</b>    | An open systems clustering solution designed to eliminate planned and unplanned downtime, simplify server consolidation, and allow the effective management of a wide range of applications in multiplatform environments.                                                                                                 |
| <b>Virtual IP Address</b>              | A unique IP address associated with the cluster. It may be brought up on any system in the cluster, along with the other resources of the service group. This address, also known as the IP alias, should not be confused with the base IP address, which is the IP address that corresponds to the host name of a system. |
| <b>VxFS (Veritas File System)</b>      | A component of the Veritas Storage Foundation product suite that provides high performance and online management capabilities to facilitate the creation and maintenance of file systems. A file system is a collection of directories organized into a structure that enables you to locate and store files.              |
| <b>VxVM (Veritas Volume Manager)</b>   | A Symantec product installed on storage clients that enables management of physical disks as logical devices. It enhances data storage management by controlling space allocation, performance, data availability, device installation, and system monitoring of private and shared systems.                               |
| <b>VVR (Veritas Volume Replicator)</b> | A data replication tool designed to contribute to an effective disaster recovery plan.                                                                                                                                                                                                                                     |



# Index

## Symbols

\$CRS\_HOME  
    setting up 360  
\$ORACLE\_HOME  
    setting up 360

## A

adding  
    users 156  
agent operations  
    Netlsnr agent 771  
    Oracle agent 759  
agents  
    about 723  
    CFSfsckd 741  
    CFSMount 736, 741  
    CSSD 755  
    CVMCluster 725  
    CVMVolDg 732  
    CVMVxconfigd 727  
    log files 415  
    MultiPrivNIC 750  
    of VCS 724  
    of VCS for Oracle 724  
    PrivNIC 745  
ASM  
    about 707  
    configuring with CVM  
        about 710  
        configuring disk group 713  
        creating shared raw volumes 714  
        installing ASM 711  
    main.cf 718  
    SF Oracle RAC configuration 708  
    verifying setup 715  
ASMDG agent  
    attribute definitions 777  
    resource type 777  
ASMDG agent attributes  
    AgentDirectory 777  
    DBAPword 777

ASMDG agent attributes *(continued)*

    DBAUser 777  
    DiskGroups 777  
    Encoding 777  
    EnvFile 777  
    Home 777  
    Owner 777  
    Sid 777  
attribute definitions  
    ASMDG agent 777  
    Netlsnr agent 772  
    Oracle agent 765  
attributes  
    about agent attributes 724  
    CFSMount agent 737, 742  
    CVMCluster agent 726  
    CVMVolDg agent 726, 734  
    CVMVxconfigd agent 728  
    MultiPrivNIC agent 751  
    PrivNIC agent  
        optional 748  
        required 746  
    UseFence 180  
Automatic Storage Management. *See* ASM

## B

basic monitoring 760  
    health check 760  
    process 760

## C

cables  
    cross-over Ethernet 495  
CFS  
    creating databases 721  
    stopping applications 618  
    unmounting file systems 618  
CFSfsckd agent 741  
    attributes 742  
CFSMount agent 736, 741  
    attributes 737

- CFSMount agent *(continued)*
  - entry points 737
  - sample configuration 740–741
  - type definition 740
- CFSTypes.cf 740
- cluster
  - removing a node from 540
  - verifying operation 221
- cluster configuration 148
- Cluster File System. *See* CFS
- Cluster Manager 37
- Cluster Volume Manager. *See* CVM
- clusters
  - basic setup 38
  - four-node illustration 38
- commands
  - gcoconfig 571
  - hastatus 221
  - hasys 222
  - lltconfig 670
  - lltstat 218
  - vradmin 586
  - vxassist 579, 581, 721
  - vxdisksetup (initializing disks) 169
  - vxedit 720
  - vxvol 579, 721
- configuration
  - required information
    - for global clusters 652
    - for I/O fencing 650
    - for secure clusters 651
    - for VCS 649
- configuration files
  - removing 627
  - See also* main.cf samples
- configuring
  - rsh 701
  - ssh 701–702
- Configuring global clusters
  - for SF Oracle RAC 564
- configuring VCS
  - adding users 156
  - event notification 156, 158
  - global clusters 160
  - secure mode 152
- coordinator disks
  - setting up 178
- CP server database 33
- CP server user privileges 33

- CSSD agent 755
  - attributes 756
- CVM
  - configuration worksheet 659
  - CVMTypes.cf file 726
  - installation worksheet 659–660
  - upgrading protocol version 305
- CVMCluster agent 725
  - attributes 726
  - entry points 725
  - sample configuration 727
  - type definition 726
- CVMTypes.cf
  - definition, CVMCluster agent 726
  - definition, CVMVolDg agent 736
  - definition, CVMVxconfigd agent 732
- CVMVolDg agent 732
  - attributes 734
  - entry points 733
  - sample configuration 736
  - type definition 736
- CVMVxconfigd agent 727
  - attributes 728
  - CVMTypes.cf 732
  - entry points 728
  - sample configuration 732
  - type definition 732

## D

- databases
  - creating 719
  - creating tablespaces
    - on CFS 721
    - on raw volumes 720
- detail monitoring 760
- disks
  - adding and initializing 169
  - coordinator 178
  - testing with vxfcntlthdw 170
  - verifying node access 172

## E

- environment variables
  - MANPATH 89
- Ethernet controllers 495

## F

- files. *See* configuration files

**G**

- GAB
  - port memberships 220
- GAB ports 221
- gabtab file
  - verifying after installation 670
- gcoconfig command 571
- Global Cluster Option (GCO)
  - overview 43
- global clusters
  - about 35
  - configuration 160
    - using VVR 571
  - configuration examples 589
  - configuring VCS service groups 571
  - illustration of dependencies 587
  - migration 601–602
  - overview 564
  - replicating database volumes 587
  - required information
    - for configuration 652
  - requirements 562
- groups
  - log owner 588
  - RVG 589

**H**

- hastatus -summary command 221
- hasys -display command 222
- health check APIs 760
- health check monitoring 760
- hubs
  - independent 495

**I**

- I/O fencing
  - checking disks 170
  - required information
    - for configuration 650
  - setting up 178
  - shared storage 170
- installation
  - of Oracle Clusterware 383
  - of Oracle Grid Infrastructure 383
  - of Oracle RAC 327
  - preparation 645
- installation worksheets 645

## installing

- Root Broker 123

## installsfrac

- installing Clusterware 383
- installing Oracle Grid Infrastructure 383
- installing Oracle RAC 330
- installing SF Oracle RAC 94
- upgrading SF Oracle RAC 238

**J**

- Java Console 37

**L**

## links

- private network 670

## Live Upgrade

- preparing 266
- upgrade paths 263
- upgrading Solaris on alternate boot disk 270

## LLT

- interconnects 89
- verifying 218

## lltconfig command 670

## llthosts file

- verifying after installation 670

## lltstat command 218

## llttab file

- verifying after installation 670

**M**

## main.cf

- after SF Oracle RAC installation 669

## main.cf files 693

## main.cf samples

- for CSSD resource 758

## MANPATH environment variable 89

## media speed 89

- optimizing 89

## migration, of Oracle 417

## monitoring

- basic 760
- detail 760

## MultiPrivNIC

- for UDP 356

## MultiPrivNIC agent 750

- attributes 751
- entry point 751
- type definition 754

- MultiPrivNIC.cf 754
  - sample configuration 755

## N

- Netlsnr agent
  - attribute definitions 772
  - operations 771
  - resource type 772
- Netlsnr agent attributes
  - AgentDebug 772
  - AgentDirectory 772
  - EnvFile 772
  - Home 772
  - Listener 772
  - LsnrPwd 772
  - MonScript 772
  - Owner 772
  - TnsAdmin 772
- nodes
  - adding Oracle RAC nodes
    - configuring GAB 502
    - configuring LLT 502
    - configuring VCSMM 502
    - configuring VXFEN 502
    - starting Volume Manager 501
  - preparing Oracle RAC nodes
    - configuring CVM 511
    - configuring private IP addresses 523
    - preparing \$CRS\_HOME 525
    - preparing \$ORACLE\_HOME 525
  - removing Oracle RAC 11g nodes
    - editing VCS configuration files 541
    - GAB configuration 542
    - LLT configuration 541
    - modifying VCS configuration 542
  - removing Oracle RAC nodes
    - Oracle procedure 540
  - removing Oracle RACnodes
    - workflow 539
- nodes, removing 539

## O

- OCR
  - creating
    - on raw volumes 340
  - creating file system
    - using CLI 341

- operating systems
  - required Solaris patches 52
- operations
  - Netlsnr agent 771
  - Oracle agent 759
- optimizing
  - media speed 89
- options
  - SF Oracle RAC configuration 145
- Oracle
  - agent log file 415
  - creating databases 720–721
  - stopping instances 614
  - supported versions 53
- Oracle 11g
  - migration from Oracle RAC 10g
    - installing Oracle RAC 10g 420
    - post-upgrade tasks 421
- Oracle 11g Release 2
  - migration from Oracle RAC 10g
    - of existing database 421
    - pre-upgrade tasks 418
- Oracle 11gRelease 2
  - migration from Oracle 9i 2
    - pre-upgrade tasks 419
- Oracle agent
  - attribute definitions 765
  - operations 759
  - resource type 764
- Oracle agent attributes
  - AgentDebug 765
  - AgentDirectory 765
  - AutoEndBkup 765
  - DBAPword 765
  - DBAUser 765
  - DetailMonitor 765
  - EnvFile 765
  - Home 765
  - IMF 765
  - MonitorOption 765
  - MonScript 765
  - Owner 765
  - Pfile 765
  - Pword 765
  - ShutDownOpt 765
  - Sid 765
  - StartUpOpt 765
  - Table 765
  - User 765

**Oracle Grid Infrastructure**

installing 383

**Oracle RAC**

database replication 585

**installation**

installing database 389

workflow 381

**manual installation**

installing database binaries 391

installing Oracle Clusterware 386

installing Oracle Grid Infrastructure 386

linking ODM library 405

**post-installation**

adding UDP IP addresses 407

**pre-installation**

creating OCR 340

creating Oracle disk groups 334

creating Oracle users and groups 331

creating voting disk volumes 340

identifying public virtual IP addresses 329

setting system parameters 330

setting up Oracle user equivalence 376

setting up storage 62

starting installsfrac 330

workflow 327

**Oracle RAC 11g**

agent log file 415

configuring service groups

preventing automatic starting of

database 409

**installation**

configuring service groups 410

setting environment variables 383

verifying 392

**post-installation**

adding patches 395

creating database 406

relinking SF Oracle RAC libraries 400

**pre-installation**

configuring MultiPrivNIC 356

preparing \$CRS\_HOME 360

preparing \$ORACLE\_HOME 360

setting environment variables 383

**Oracle RAC 11g Clusterware**

creating file system 363, 525

verifying installation 392

**Oracle RAC 11g Release 1 Clusterware**

creating file system 368, 530

**P****patches**

adding 395

for Solaris 52

**patchsets**

adding 395

**PATH variable**

VCS commands 217

**ports**

GAB 221

**preparing**

Live Upgrade 266

**primary site**

creating SRL volume 578

setting up 565

setting up replication objects 579

VCS configuration 592, 595

**PrivNIC agent 745**

attributes 746

entry point 746

sample configuration 750

type definition 749

**PrivNIC.cf 749**

process monitoring 760

**R****raw volumes**

creating

Oracle DBCA procedure 714

**remote shell 702**

removing a node from a cluster 540

**replication**

automatic synchronization 585

configuring on both sites 570

full synchronization with Checkpoint 586

modifying VCS configuration 588

options 563

setting up primary site 565

setting up secondary site 567

supported hardware 562

supported software 562

supported technologies 562

using VVR 578

verifying status 587

**resource type**

ASMDG 777

Netlsnr 772

Oracle 764

## resources

- CVMVolDg 589
- RVGSharedPri 589

## response file

- about 425
- syntax 428

## Root Broker 36

- installing 123

## rsh

- configuration 701

**S**

## SCSI-3 persistent reservations

- verifying 178

## secondary site

- configuring replication 581
- creating SRL volume 581
- setting up 567
- setting up disk groups 583
- setting up RLINKs 582

## secure clusters

- adding users 651
- required information
  - for configuration 651

## secure shell 702

## service groups

- configuring 410
- VCS, for global clusters 571

## setting umask, before installing 87

## SF Oracle RAC

- about 27
- coordinator disks 178
- high-level view 38
- illustration of global cluster dependencies 587
- takeover 602

## SF Oracle RAC configuration

- about 145
- configuring clusters 148
- creating configuration files 161
- of components 146
- options 145
- preparation
  - worksheets 645
- starting processes 161
- verifying 220

## SF Oracle RAC installation

- on alternate root 115
- pre-installation tasks
  - mounting product disc 88

SF Oracle RAC installation *(continued)*pre-installation tasks *(continued)*

- setting MANPATH 89
- setting umask 87
- setting up shared storage 88
- synchronizing time settings 87
- verifying systems 90

## preinstallation information 45

## preparation

- worksheets 645

## requirements

- hardware 46
- software 53

## using installsfrac 94

## verifying

- cluster operations 217
- GAB operations 217
- LLT operations 217

## SF Oracle RAC uninstallation

## preparation

- stopping applications, CFS 618
- stopping applications, VxFS 620
- stopping Oracle instances 614
- stopping VCS 619
- uninstalling 617
- unmounting CFS file systems 618
- unmounting VxFS file systems 620

## removing configuration files 627

## removing packages 626

## using uninstallsfrac program 626

## workflow 611

## SF Oracle RAC upgrade

## post-upgrade tasks

- relinking Oracle RAC libraries 302
- upgrading CVM protocol version 305
- upgrading disk group version 305

## preparation 231, 258

## restoring configuration files 239

## stopping cluster resources 231, 258

## stopping Oracle RAC 231, 258

## upgrading licenses 283

## using installsfrac 238

## shared storage

- setting up 88

## SMTP email notification 156

## SNMP trap notification 158

## ssh

- configuration 701
- configuring 702

- Symantec Product Authentication Service 36, 123, 152
- synchronizing time settings, before installing 87
- system communication using remote shell
  - secure shell 702
- system state attribute value 221

## U

- uninstallation
  - of SF Oracle RAC 621
- uninstallsrac program
  - removing packages 626
- upgrade paths
  - Live Upgrade 263
- users
  - creating 331
  - for &ProductNameShort;, adding 651
- using Live Upgrade 263

## V

- VCS
  - agent log file 415
  - command directory path variable 217
  - configuration, for database volume replication 587
  - configuring service groups 571
  - notifications 35
  - required information
    - for configuration 649
  - stopping 619
- VCS configuration
  - for replication 588
- VCS Global cluster option. *See* GCO
- VCS notifications
  - SMTP notification 35
  - SNMP notification 35
- Veritas File System
  - stopping applications 620
  - unmounting 620
- Veritas Operations Manager 36
- Veritas Volume Replicator. *See* VVR
- virtual fire drill 763
- voting disk
  - creating
    - on raw volumes 340
- voting disks
  - creating file system
    - using CLI 341

- vradm command 586

## VVR

- about 35
- configuration examples 589
- configuration worksheet 659
- configuring global clusters 571
- configuring on both sites 570
- database volume replication
  - configuring VCS 587
- defining heartbeat cluster objects 572
- defining remote clusters 572
- global cluster overview 577
- installation worksheet 659–660
- Oracle RAC database replication 585
- primary site
  - creating SRL volume 578
  - setting up replication objects 579
- replication
  - using automatic synchronization 585
  - using full synchronization with Checkpoint 586
- replication agents 563
- secondary site
  - configuring replication 581
  - creating SRL volume 581
  - setting up disk groups 583
  - setting up RLINKs 582
- setting up primary site 565
- setting up replication 578
- setting up secondary site 567
- types of replication 563
- VCS configuration
  - CVMoIDG resource 589
  - log owner group 588
  - Oracle RAC database service group 589
  - primary site 592
  - RVG group 589
  - RVGSharedPri resource 589
  - secondary site 595
  - verifying replication 587
- vxassist command 579, 581, 721
- vxdisksetup command 169
- VxFS. *See* Veritas File System
- vxvol command 579, 721

## W

- worksheets
  - for &ProductNameShort; 646
  - for CVM 659

worksheets (*continued*)  
  for replication  
    SRDF 660  
    VVR 659  
  for VVR 659